

# Contrastive Learning from Extremely Augmented Skeleton Sequences for Self-Supervised Action Recognition

Tianyu Guo<sup>1</sup>, Hong Liu<sup>1\*</sup>, Zhan Chen<sup>1</sup>, Mengyuan Liu<sup>2</sup>, Tao Wang<sup>1</sup>, Runwei Ding<sup>1</sup>

<sup>1</sup> Key Laboratory of Machine Perception, Peking University, Shenzhen Graduate School, China

<sup>2</sup> The School of Intelligent Systems Engineering, Sun Yat-sen University, China

{levigty, taowang}@stu.pku.edu.cn, {hongliu, zhanchen\_cz, dingrunwei}@pku.edu.cn, nkliuyifang@gmail.com

## Abstract

In recent years, self-supervised representation learning for skeleton-based action recognition has been developed with the advance of contrastive learning methods. The existing contrastive learning methods use normal augmentations to construct similar positive samples, which limits the ability to explore novel movement patterns. In this paper, to make better use of the movement patterns introduced by extreme augmentations, a Contrastive Learning framework utilizing Abundant Information Mining for self-supervised action Representation (AimCLR) is proposed. First, the extreme augmentations and the Energy-based Attention-guided Drop Module (EADM) are proposed to obtain diverse positive samples, which bring novel movement patterns to improve the universality of the learned representations. Second, since directly using extreme augmentations may not be able to boost the performance due to the drastic changes in original identity, the Dual Distributional Divergence Minimization Loss (D<sup>3</sup>M Loss) is proposed to minimize the distribution divergence in a more gentle way. Third, the Nearest Neighbors Mining (NNM) is proposed to further expand positive samples to make the abundant information mining process more reasonable. Exhaustive experiments on NTU RGB+D 60, PKU-MMD, NTU RGB+D 120 datasets have verified that our AimCLR can significantly perform favorably against state-of-the-art methods under a variety of evaluation protocols with observed higher quality action representations. Our code is available at <https://github.com/Levigty/AimCLR>.

## Introduction

On account that action recognition has very broad application in many fields such as human-computer interaction, video content analysis, and smart surveillance, it has always been a popular research topic in the field of computer vision. Due to the development of depth sensors (Zhang 2012) and the human pose estimation algorithms (Cao et al. 2019; Fang et al. 2017), skeleton-based action recognition has gradually become a significant branch of action recognition.

In the past few years, most of the existing skeleton-based action recognition methods are based on the supervised learning framework. Whether it is a CNN-based method (Du, Fu, and Wang 2015; Ke et al. 2017; Liu, Liu, and Chen

2017), RNN-based method (Du, Wang, and Wang 2015; Song et al. 2018; Zhang et al. 2019), or GCN-based method (Yan, Xiong, and Lin 2018; Shi et al. 2019; Si et al. 2019; Chen et al. 2021), numerous labeled data is used to learn the action representation. Fully supervised action recognition methods are inevitably data-driven, but the cost of labeling large-scale datasets is particularly high. Therefore, more and more researchers intend to use unlabeled skeleton data for learning human action representation.

Recently, several works (Zheng et al. 2018; Su, Liu, and Shlizerman 2020; Lin et al. 2020) focus on designing pretext tasks for self-supervised methods to learn action representations from unlabeled skeleton data. With the development of contrastive self-supervised learning and its ability to make feature representations have better discrimination, several works (Rao et al. 2021; Li et al. 2021) directly rely on the contrastive learning framework, using normal augmentations to construct similar positive samples. However, those carefully designed augmentations limit the model to further explore the novel movement patterns exposed by other augmentations and there are still several significant motivations that need to be carefully considered:

**1) Stronger data augmentations could benefit representation learning.** In SkeletonCLR (Li et al. 2021), it just uses two data augmentations *Shear* and *Crop*. Nevertheless, studies (Tian et al. 2020; Wang and Qi 2021) have shown that data augmentation design is crucial, and the abundant semantic information introduced by stronger data augmentations can significantly improve the generalizability of learned representations and eventually bridge the gap with the fully supervised methods.

**2) Directly using stronger augmentations could deteriorate the performance.** Stronger data augmentations bring novel movement patterns while the augmented skeleton sequence may not keep the identity of the original sequence. Therefore, directly using extreme augmentations may not necessarily be able to boost the performance due to the drastic changes in original identity. Thus, additional efforts are needed to explore the role of stronger augmentations.

**3) How to force the model to learn more features.** Simply relying on the contrastive learning framework can not force the model to study more features well. Studies (Pan et al. 2021; Cheng et al. 2020) have shown that the drop mechanism can be used for contrastive learning and

\*Corresponding Author: hongliu@pku.edu.cn

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

can effectively solve the problem of over-fitting. Currently, the drop mechanism is not well exploited in self-supervised skeleton-based action recognition.

**4) How to better expand the positive set to make the learning process more reasonable.** In contrastive learning, two different augmented samples from the same sample are considered as positive samples, while samples in the memory bank are all treated as negative samples. However, the samples in the memory bank are not necessarily all negative samples which makes the learning process unreasonable to a certain extent.

To this end, a contrastive learning framework utilizing abundant information mining for self-supervised action representation (AimCLR) is proposed. Specifically, the framework of AimCLR is shown in Figure 1. Different from traditional contrastive learning methods (Rao et al. 2021; Li et al. 2021) which directly use normally augmented view, a novel framework proposed in our work is based on the extreme augmentations and the drop mechanism which obtain diverse positive samples and bring abundant spatio-temporal information. Then the Dual Distributional Divergence Minimization Loss ( $D^3M$  Loss) is proposed to minimize the distribution divergence between the normally augmented view and the extremely augmented views. Furthermore, the Nearest Neighbors Mining (NNM) is used to expand the positive set to make the learning process more reasonable.

In summary, we have made the following contributions:

- Compared with the traditional contrastive learning method using similar augmented pairs, AimCLR is proposed to use more extreme augmentations and more reasonable abundant information mining which greatly improve the effect of contrastive learning.
- Specifically, the extreme augmentations and the Energy-based Attention-guided Drop Module (EADM) are proposed to introduce novel movement patterns to force the model to learn more general representations. Then the  $D^3M$  Loss is proposed to gently learn from the introduced movement patterns. In order to alleviate the irrationality of the positive set, we further propose the Nearest Neighbors Mining (NNM) strategy.
- With the multi-stream fusion scheme, our 3s-AimCLR achieves state-of-the-art performances under a variety of evaluation protocols such as KNN evaluation, linear evaluation, semi-supervised evaluation, and finetune evaluation protocol on three benchmark datasets.

## Related Work

**Supervised Skeleton-based Action Recognition.** Early skeleton-based action recognition methods are usually based on hand-crafted features (Wang et al. 2012; Vemulapalli, Arate, and Chellappa 2014; Vemulapalli and Chellappa 2016). With the rapid development of deep learning in recent years, some methods (Du, Wang, and Wang 2015; Song et al. 2018; Zhang et al. 2019) use RNN to process skeleton data. Meanwhile, several methods convert the 3D skeleton sequence into an image representation and have achieved good results based on CNN (Du, Fu, and Wang 2015; Ke et al. 2017; Liu, Liu, and Chen 2017). In recent years, with the introduction

of graph convolutional networks, a variety of GCN-based methods (Shi et al. 2019; Si et al. 2019; Chen et al. 2021) have emerged on the basis of ST-GCN (Yan, Xiong, and Lin 2018) to better model the spatio-temporal structure relationship. In this paper, we adopt the widely-used ST-GCN as the encoder to extract the skeleton features.

### Contrastive Self-Supervised Representation Learning.

Some contrastive learning methods (Zhang, Isola, and Efros 2016; Pathak et al. 2016; Gidaris, Singh, and Komodakis 2018) focus on designing various novel pretext tasks to find the pattern information hidden in the unlabeled data. MoCo and MoCov2 (He et al. 2020; Chen et al. 2020b) promotes contrastive self-supervised learning through a queue-based memory bank and momentum update mechanism. SimCLR (Chen et al. 2020a) uses a much larger batch size to compute the embeddings in real-time, and uses a multi-layer perceptron (MLP) to further improve the performance of self-supervised representation learning. Current work CLSA (Wang and Qi 2021) shows that strong augmentations are beneficial to the performance of downstream tasks and it expects to learn from strongly augmented samples. The development of contrastive self-supervised representation learning also laid the foundation for our AimCLR.

### Self-supervised Skeleton-based Action Recognition.

LongT GAN (Zheng et al. 2018) proposes to use the encoder-decoder to regenerate the input sequence to obtain useful feature representation. P&C (Su, Liu, and Shlizerman 2020) proposes a training strategy to weaken the decoder, forcing the encoder to learn more discriminative features. Yang et al. (2021b) design a novel skeleton cloud colorization technique to learn skeleton representations. AS-CAL (Rao et al. 2021) and SkeletonCLR (Li et al. 2021) use momentum encoder for contrastive learning with single-stream skeleton sequence while CrosSCLR (Li et al. 2021) proposes cross-stream knowledge mining strategy to improve the performance and ISC (Thoker, Doughty, and Snoek 2021) proposes inter-skeleton contrastive learning to learn from multiple different input skeleton representations. In order to learn more general features,  $MS^2L$  (Lin et al. 2020) introduces multiple self-supervised tasks to learn more general representations. However, the currently existing methods rarely explore the gains that abundant spatio-temporal information brings to the task of action recognition. Therefore, a more concise and general framework needs to be proposed.

## AimCLR

### SkeletonCLR Overview

SkeletonCLR (Li et al. 2021) is based on the recent advanced practice MoCov2 (Chen et al. 2020b) to learn single-stream 3D action representations. The pipeline of the SkeletonCLR is shown in the bottom blue part of Figure 1. Given an encoded query  $\hat{z}$  and encoded key  $z$ , the batch embeddings of  $z$  are stored in first-in-first-out memory bank  $\mathbf{M} = \{m_i\}_{i=1}^M$  to get rid of redundant computation. It serves as negative samples for the next training steps. Then the InfoNCE loss (Oord, Li, and Vinyals 2018) can be written as:

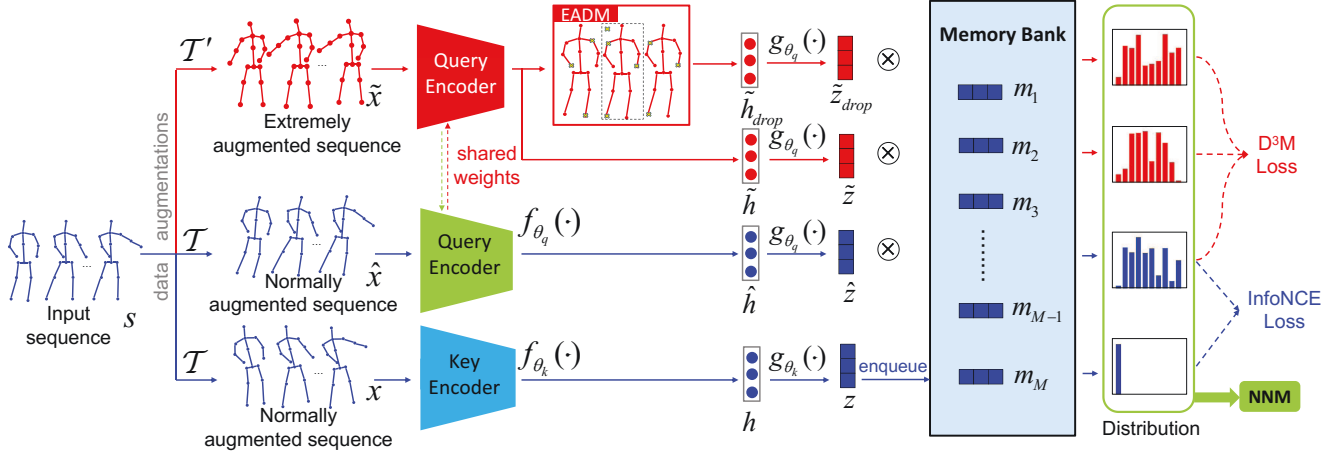


Figure 1: The pipeline of the proposed AimCLR. Through the extreme augmentations  $\mathcal{T}'$  and normal augmentations  $\mathcal{T}$ ,  $x$ ,  $\hat{x}$  and  $\tilde{x}$  are obtained from the input sequence  $s$ . The query encoder and an MLP extract  $\hat{z}$  and  $\tilde{z}$  while the query encoder with EADM and an MLP is used to obtain the  $\tilde{z}_{drop}$ . The momentum updated key encoder and an MLP is used to obtain  $z$ ,  $z$  are stored in the memory bank in each training step, serving as negative samples for the next training steps. While using InfoNCE loss, we also propose D<sup>3</sup>M Loss to minimize the distribution divergence of  $\hat{z}$ ,  $\tilde{z}$ , and  $\tilde{z}_{drop}$ . Furthermore, we propose NNM to expand the positive set to make the learning process more reasonable.

$$\mathcal{L}_{Info} = -\log \frac{\exp(\hat{z} \cdot z / \tau)}{\exp(\hat{z} \cdot z / \tau) + \sum_{i=1}^M \exp(\hat{z} \cdot m_i / \tau)}, \quad (1)$$

where  $\tau$  is the temperature hyper-parameter, and dot product  $\hat{z} \cdot z$  is to compute their similarity where  $z$ ,  $\hat{z}$  are normalized.

After computing the InfoNCE loss in Eq. (1), the query encoder is updated via gradients while the key encoder is updated as a moving-average of the query encoder. We denote the parameters of the query encoder as  $\theta_q$  and those of the key encoder as  $\theta_k$ . Then the key encoder is updated as:

$$\theta_k \leftarrow m\theta_k + (1 - m)\theta_q, \quad (2)$$

where  $m \in [0, 1)$  is a momentum coefficient. The momentum encoder is updated slowly based on the encoder change, which ensures stable key representations.

## Data Augmentations

In contrastive learning, augmentations for positive samples bring semantic information for the encoder to learn. However, those carefully designed augmentations limit the encoder to further explore the novel patterns exposed by other augmentations. Therefore, we aim to introduce novel movement patterns through stronger data augmentation, thereby benefiting contrastive learning.

1) *Normal Augmentations  $\mathcal{T}$* . One spatial augmentation *Shear* and one temporal augmentation *Crop* are used as the normal augmentations like SkeletonCLR.

2) *Extreme Augmentations  $\mathcal{T}'$* . We introduce four spatial augmentations: *Shear*, *Spatial Flip*, *Rotate*, *Axis Mask* and two temporal augmentations: *Crop*, *Temporal Flip* and two spatio-temporal augmentations: *Gaussian Noise* and *Gaussian Blur*. We use the combination of all the 8 augmentations (2 normal and 6 other augmentations) as ‘‘Extreme Augmentations’’ to finally get one extremely augmented sequence.

On account that the combination of extreme augmentations is complicated, we hope to explore a more general framework in which extreme augmentations can introduce more novel movement patterns than normal augmentations.

## Energy-based Attention-guided Drop Module

For a feature learned by the encoder, we hope that even if some important features are discarded, different actions can be distinguished. Studies (Pan et al. 2021; Cheng et al. 2020) have shown that the drop mechanism can be used for contrastive learning, and can effectively solve the problem of over-fitting. It inspires us to calculate the spatio-temporal attention map to drop several important features, which could force the model to learn more features and obtain more general and robust feature representations.

Actually, there are lots of modules (Hu, Shen, and Sun 2018; Woo et al. 2018; Lee, Kim, and Nam 2019) proposed to calculate the attention maps. In order not to introduce additional parameters, we adopt the parameter-free attention module Simam (Yang et al. 2021a) to calculate the attention map. Formally,  $t$  and  $x_i$  are the target neuron and other neurons in a single channel of the input feature  $X \in \mathbb{R}^{C \times T \times V}$  where  $C$  denotes the number of channels,  $T$  denotes the temporal dimension and  $V$  denotes the spatial dimension. The minimal energy  $e_t$  of target neuron  $t$  can be computed with the following:

$$e_t = \frac{4(\hat{\sigma}^2 + \lambda)}{(t - \hat{\mu})^2 + 2\hat{\sigma}^2 + 2\lambda}, \quad (3)$$

where  $\hat{\mu} = \frac{1}{N} \sum_{i=1}^N x_i$ ,  $\hat{\sigma}^2 = \frac{1}{N} \sum_{i=1}^N (x_i - \hat{\mu})^2$ ,  $\lambda$  is a hyper-parameter, and  $N = T \times V$  is the number of neurons on the channel. Eq. (3) indicates that the lower energy  $e_t$ , the neuron  $t$  is more distinctive from surround neurons, and more important for visual processing. Therefore, the importance

---

**Algorithm 1: Energy-based attention-guided drop module.**

---

**Require:** a GCN feature  $X$ : the dimension is  $\mathbb{R}^{C \times T \times V}$ ;  
 $keep\_margin$ : control the importance margin to drop.

**Ensure:** Features after processing.

- 1: Compute the attention map  $\tilde{\alpha}$ ;
  - 2: Generate the spatial attention map  $M_s$  and the temporal attention map  $M_t$  using  $\tilde{\alpha}$  and  $keep\_margin$ ;
  - 3: Apply the spatial mask:  $X = X \times M_s$ ;
  - 4: Normalize the feature:  
 $X = X \times count(M_s) / count\_ones(M_s)$ ;
  - 5: Apply the temporal mask:  $X = X \times M_t$ ;
  - 6: Normalize the feature:  
 $X = X \times count(M_t) / count\_ones(M_t)$ ;
  - 7: **return**  $X$ ;
- 

of each neuron can be obtained by  $1/e_t$ . Then we can obtain the attention map  $\tilde{\alpha}$  by that  $\tilde{\alpha} = \text{sigmoid}(1/\mathbf{E})$ , where  $\mathbf{E}$  groups all  $e_t$ . After that, we use the attention map  $\tilde{\alpha}$  to drop some important features using Algorithm 1.

### Dual Distributional Divergence Minimization

As shown in Figure 1, for the input sequence  $s$ , we apply normal augmentations  $\mathcal{T}$  and extreme augmentations  $\mathcal{T}'$  to obtain  $x, \hat{x}$  and  $\tilde{x}$ . The query encoder  $f_{\theta_q}$  is applied to extract features:  $\hat{h} = f_{\theta_q}(\hat{x})$  and  $\tilde{h} = f_{\theta_q}(\tilde{x})$ .  $\tilde{h}_{drop}$  is the dropped features after EADM. An MLP head  $g_{\theta_q}$  is applied to project the feature to a lower dimension space:  $\hat{z} = g_{\theta_q}(\hat{h})$ ,  $\tilde{z} = g_{\theta_q}(\tilde{h})$ ,  $\tilde{z}_{drop} = g_{\theta_q}(\tilde{h}_{drop})$ . The key encoder  $f_{\theta_k}$  and  $g_{\theta_k}$  are the momentum updated version of  $f_{\theta_q}$  and  $g_{\theta_q}$ .

The memory bank  $\mathbf{M} = \{m_i\}_{i=1}^M$  of  $M$  negative samples is a first-in-first-out queue updated per iteration by  $z$ . After each inference step,  $z$  will enqueue while the earliest embedding in  $\mathbf{M}$  will dequeue.  $\mathbf{M}$  provides numerous negative embeddings while the new calculated  $z$  is the positive embedding. Thus, we can obtain a conditional distribution:

$$p(m_i | \hat{z}) = \frac{\exp(\hat{z} \cdot m_i / \tau)}{\exp(\hat{z} \cdot z / \tau) + \sum_{i=1}^M \exp(\hat{z} \cdot m_i / \tau)}, \quad (4)$$

which encodes the likelihood of the query  $\hat{z}$  being assigned to the embedding  $m_i$  in the memory bank  $\mathbf{M}$ . Similarly, we can also have the likelihood of positive pairs for the query  $\hat{z}$  being assigned to its positive counterpart  $z$ :

$$p(z | \hat{z}) = \frac{\exp(\hat{z} \cdot z / \tau)}{\exp(\hat{z} \cdot z / \tau) + \sum_{i=1}^M \exp(\hat{z} \cdot m_i / \tau)}. \quad (5)$$

The InfoNCE loss in Eq. (1) can be rewritten in another form:

$$\mathcal{L}_{Info} = -q(z | \hat{z}) \log p(z | \hat{z}) - \sum_{i=1}^M q(m_i | \hat{z}) \log p(m_i | \hat{z}), \quad (6)$$

where  $q(z | \hat{z})$  is the ideal distribution of the likelihood,  $p(z | \hat{z})$  is the distribution learned by the network. To avoid

the unknown ideal distribution exploration, InfoNCE loss regards  $\hat{z}$  as a one-hot distribution, where positive pairs have  $q(z | \hat{z}) = 1$  and negative pairs satisfy  $q(m_i | \hat{z}) = 0 (i \in [1, M])$ . It means that InfoNCE loss maximizes the agreement of two different augmented sequences' representations from the same sequence while minimizing the agreement with other negative sequences. To explore the novel movement patterns from the extreme augmentations, a straightforward approach is directly using the extremely augmented sequence as query and the normally augmented sequence as key in InfoNCE loss. However, compared to the normally augmented sequence, the extremely augmented sequence may not keep the identity of the original sequence due to the dramatic changes in movement patterns, leading to performance degradation.

In addition, it's almost impossible to obtain the ideal likelihood distribution. Fortunately, CLSA (Wang and Qi 2021) found that the normally augmented query and the extremely augmented query share similar distribution for a randomly initialized network. It inspires us that the distribution of normally augmented query over memory bank can be used to supervise that of the extremely augmented query. It avoids directly using one-hot distribution for extremely augmented views and is able to explore the novel patterns exposed by the extreme augmentations.

Similar to Eq. (4) and Eq. (5), we obtain a conditional distribution for  $\tilde{z}$  based on its positive samples and negative samples:  $p(z | \tilde{z})$  and  $p(m_i | \tilde{z})$ . The conditional distribution  $p(m_i | \tilde{z}_{drop})$  and  $p(z | \tilde{z}_{drop})$  for  $\tilde{z}_{drop}$  is calculated in the same way. Then, we propose to minimize the following distributional divergence between the normally augmented view and the extremely augmented view such that:

$$\mathcal{L}_{d1} = -p(z | \tilde{z}) \log p(z | \tilde{z}) - \sum_{i=1}^M p(m_i | \tilde{z}) \log p(m_i | \tilde{z}). \quad (7)$$

Similarly, the distributional divergence between the normally augmented view and the dropped extremely augmented view is minimized such that:

$$\mathcal{L}_{d2} = -p(z | \tilde{z}) \log p(z | \tilde{z}_{drop}) - \sum_{i=1}^M p(m_i | \tilde{z}) \log p(m_i | \tilde{z}_{drop}), \quad (8)$$

Therefore, the proposed D<sup>3</sup>M loss can be formulated as  $\mathcal{L}_D = 1/2(\mathcal{L}_{d1} + \mathcal{L}_{d2})$ .

### Nearest Neighbors Mining

Traditional contrastive learning methods regard the normally augmented samples from the same sample as positive samples and all samples in the memory bank as negative samples. However, the samples in the memory bank are not necessarily all negative samples (Dwibedi et al. 2021). Therefore, we hope that the nearest neighbors of query  $\hat{z}$ ,  $\tilde{z}$ , and  $\tilde{z}_{drop}$  over the memory bank  $\mathbf{M}$  should be considered as positive samples to expand the positive set.

Specifically,  $N_+^n$  is the index set of the nearest top-k neighbors that are most similar to the normally augmented query  $\hat{z}$  in the memory bank  $\mathbf{M}$ . Similarly, we could also have the  $N_+^e$  and  $N_+^d$  to represent the index set of the nearest top-k neighbors of the extremely augmented query  $\tilde{z}$  and the dropped extremely augmented query  $\tilde{z}_{drop}$ . Thus, we can set the nearest top-k neighbors as positive samples to make the learning process more reasonable:

$$\mathcal{L}_N = -\log \frac{\exp(\hat{z} \cdot z/\tau)}{\exp(\hat{z} \cdot z/\tau) + \sum_{i=1}^M \exp(\hat{z} \cdot m_i/\tau)} - \sum_{j \in N_+} \log \frac{\exp(\hat{z} \cdot m_j/\tau)}{\exp(\hat{z} \cdot z/\tau) + \sum_{i=1}^M \exp(\hat{z} \cdot m_i/\tau)}. \quad (9)$$

In this equation,  $N_+ = N_+^n \cup N_+^e \cup N_+^d$ . Compared to Eq. (1), Eq. (9) will lead to a more regular space by pulling close more high-confidence positive samples.

**Two-stage Training Strategy.** In the early training stage, the model is not stable and strong enough to provide reliable nearest neighbors. Thus, we perform two-stage training for AimCLR: In the first stage, the model is trained with the loss function:  $\mathcal{L}_1 = \alpha \mathcal{L}_{Info} + \beta \mathcal{L}_D$ . Then in the second training stage, the loss function is  $\mathcal{L}_2 = \alpha \mathcal{L}_N + \beta \mathcal{L}_D$  to start mining the nearest neighbors. Here,  $\alpha$  and  $\beta$  are the coefficient to balance the loss. Though other values may achieve better results, we use  $\alpha = \beta = 1$  to make AimCLR more general.

## Experiments

### Dataset

**PKU-MMD Dataset** (Liu et al. 2020): It contains almost 20,000 action sequences covering 51 action classes. It consists of two subsets. Part I is an easier version for action recognition, while part II is more challenging with more noise caused by view variation. We conduct experiments under the cross-subject protocol on the two subsets.

**NTU RGB+D 60 Dataset** (Shahroudy et al. 2016): The dataset contains 56,578 action sequences and 60 action classes. There are two evaluation protocols: cross-subject (xsub) and cross-view (xview). In xsub, half of the subjects are used as training sets, and the rest are used as test sets. In xview, the samples of camera 2 and 3 are used for training while the samples of camera 1 are used for testing.

**NTU RGB+D 120 Dataset** (Liu et al. 2019): It is NTU RGB+D 60 based extension, whose scale is up to 120 action classes and 113,945 sequences. There are two evaluation protocols: cross-subject (xsub) and cross-setup (xset). In xsub, actions performed by 53 subjects are for training and the others are for testing. In xset, all 32 setups are separated as half for training and the other half for testing.

### Experimental Settings

All the experiments are conducted on the PyTorch (Paszke et al. 2019) framework. For data pre-processing, we follow SkeletonCLR and CrosSCLR (Li et al. 2021) for a fair comparison. The mini-batch size is set to 128.

**Self-supervised Pretext training.** ST-GCN is adopted as the encoder. For contrastive settings, we follow that in SkeletonCLR. Specifically, the feature dimension is 128, the size

w/ NA	w/ EA	w/ EADM	w/ NNM	NTU-60(%)	
				xsub	xview
✓				75.0	79.8
	✓			71.3	77.8
✓	✓			77.4	82.5
✓	✓	✓		78.2	82.8
✓	✓	✓	✓	<b>78.9</b>	<b>83.8</b>

Table 1: Ablation study results on NTU-60 dataset.

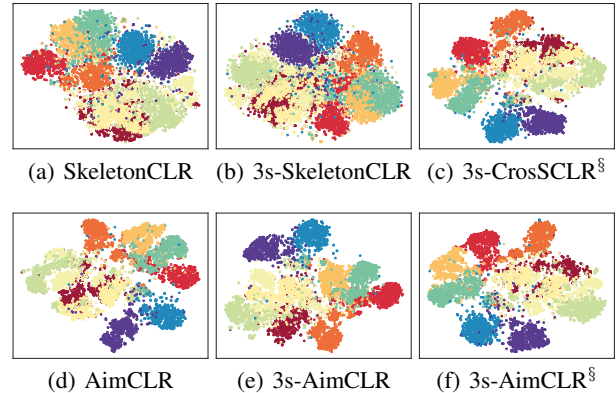


Figure 2: The t-SNE visualization of embeddings on NTU-60 xsub. “§” means using cross-stream knowledge mining strategy proposed in 3s-CrosSCLR (Li et al. 2021).

of the memory bank is 32768, the momentum coefficient  $m$  is set to 0.999, and the temperature hyper-parameter  $\tau$  is set to 0.07. For optimization, we use SGD with momentum (0.9) and weight decay (0.0001). The model is trained for 300 epochs with a learning rate of 0.1 (decreases to 0.01 at epoch 250). In CrosSCLR’s cross-training strategy, it already explored that if the value of  $k$  is too large, it may cause mining errors then cause performance degradation. Thus, we set  $k = 1$  in the nearest neighbors mining for fair comparisons. For the two-stage training strategy mentioned in Section , the encoder is trained with  $\mathcal{L}_1$  in the first 150 epochs while trained with  $\mathcal{L}_2$  in the remaining epochs. We also generate three streams of skeleton sequences, i.e., joint, bone, and motion like SkeletonCLR. For all the reported results of three streams, we use the weights of [0.6, 0.6, 0.4] for weighted fusion like other multi-stream GCN methods.

**KNN Evaluation Protocol.** A K-Nearest Neighbor (KNN) classifier is used on the features of the trained encoder. It can also reflect the quality of the features learned by the encoder.

**Linear Evaluation Protocol.** The models are verified by linear evaluation for the action recognition task. Specifically, we train a linear classifier (a fully-connected layer followed by a softmax layer) supervised with encoder fixed.

**Semi-supervised Evaluation Protocol.** We pre-train the encoder with all data and then finetuning the whole model with only 1% or 10% randomly selected labeled data.

**Finetune Protocol.** We append a linear classifier to the trained encoder and then train the whole model to compare

Method	Stream	NTU-60(%)				PKU(%)		NTU-120(%)			
		xsub		xview		part I		xsub		xset	
		acc.	$\Delta$	acc.	$\Delta$	acc.	$\Delta$	acc.	$\Delta$	acc.	$\Delta$
SkeletonCLR (CVPR 21)	joint	68.3		76.4		80.9		56.8		55.9	
<b>AimCLR (Ours)</b>	joint	<b>74.3</b>	$\uparrow 6.0$	<b>79.7</b>	$\uparrow 3.3$	<b>83.4</b>	$\uparrow 2.5$	<b>63.4</b>	$\uparrow 6.6$	<b>63.4</b>	$\uparrow 7.5$
SkeletonCLR (CVPR 21)	motion	53.3		50.8		63.4		39.6		40.2	
<b>AimCLR (Ours)</b>	motion	<b>66.8</b>	$\uparrow 13.5$	<b>70.6</b>	$\uparrow 19.8$	<b>72.0</b>	$\uparrow 8.6$	<b>57.3</b>	$\uparrow 17.7$	<b>54.4</b>	$\uparrow 14.2$
SkeletonCLR (CVPR 21)	bone	69.4		67.4		72.6		48.4		52.0	
<b>AimCLR (Ours)</b>	bone	<b>73.2</b>	$\uparrow 3.8$	<b>77.0</b>	$\uparrow 9.6$	<b>82.0</b>	$\uparrow 9.4$	<b>62.9</b>	$\uparrow 14.5$	<b>63.4</b>	$\uparrow 11.4$
3s-SkeletonCLR (CVPR 21)	joint+motion+bone	75.0		79.8		85.3		60.7		62.6	
<b>3s-AimCLR (Ours)</b>	joint+motion+bone	<b>78.9</b>	$\uparrow 3.9$	<b>83.8</b>	$\uparrow 4.0$	<b>87.8</b>	$\uparrow 2.5$	<b>68.2</b>	$\uparrow 7.5$	<b>68.8</b>	$\uparrow 6.2$

Table 2: Linear evaluation results compared with SkeletonCLR on NTU-60, PKU-MMD, and NTU-120 dataset. “ $\Delta$ ” represents the gain compared to SkeletonCLR using the same stream data. “3s” means three stream fusion.

Method	100ep	150ep	200ep	300ep
3s-SkeletonCLR(repro.)	71.3	73.8	74.1	74.1
3s-CrosSCLR(repro.) <sup>§</sup>	70.0	72.8	76.0	77.2
<b>3s-AimCLR (Ours)<sup>§</sup></b>	75.4	76.0	78.2	78.6
<b>3s-AimCLR (Ours)</b>	<b>76.5</b>	<b>77.4</b>	<b>78.3</b>	<b>78.9</b>

Table 3: Linear evaluation results on NTU-60 xsub for different epochs.

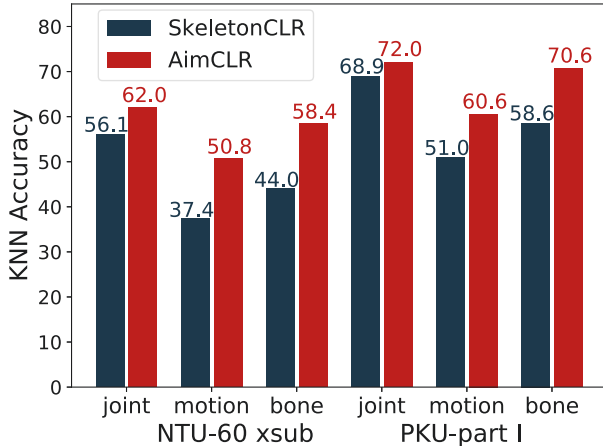


Figure 3: Comparison of KNN accuracy of SkeletonCLR (Li et al. 2021) and our AimCLR.

it with fully supervised methods.

### Ablation Study

We conduct ablation studies on different datasets to verify the effectiveness of different components of our method.

**The effectiveness of the data augmentation.** As shown in Table 1, 3s-SkeletonCLR (Li et al. 2021) uses the normal augmentations (w/ NA) and achieves the accuracy of 75.0% and 79.8% on xsub and xview respectively. While simply replacing the normal augmentations with extreme augmentations (w/ EA), the accuracies drop on both xsub and xview. It also illustrates that directly using extreme augmentations

Method	NTU-60(%)	
	xsub	xview
<i>Single-stream:</i>		
LongT GAN (AAAI 18)	39.1	48.1
MS <sup>2</sup> L (ACM MM 20)	52.6	-
AS-CAL (Information Sciences 21)	58.5	64.8
P&C (CVPR 20)	50.7	76.3
SeBiReNet (ECCV 20)	-	<b>79.7</b>
SkeletonCLR (CVPR 21)	68.3	76.4
<b>AimCLR (Ours)</b>	<b>74.3</b>	<b>79.7</b>
<i>Three-stream:</i>		
3s-SkeletonCLR (CVPR 21)	75.0	79.8
3s-Colorization (ICCV 21)	75.2	83.1
3s-CrosSCLR (CVPR 21) <sup>§</sup>	77.8	83.4
<b>3s-AimCLR (Ours)<sup>§</sup></b>	78.6	82.6
<b>3s-AimCLR (Ours)</b>	<b>78.9</b>	<b>83.8</b>

Table 4: Linear evaluation results on NTU-60 dataset.

may not necessarily be able to boost the performance due to the drastic changes in original identity. While both EA and NA are used, i.e., when  $\mathcal{L}_{d1}$  loss in Eq. (7) comes into play, the accuracy is improved by 2.4% and 2.7%.

**The effectiveness of the EADM and NNM.** From Table 1, it is worth noting that when EADM is introduced, the accuracy on xsub and xview are improved by 0.8% and 0.3%, respectively. Notably, our 3s-AimCLR achieves the highest accuracy when NNM is further introduced. It also shows that the proposed EADM and NNM can make the encoder further learn more robust and suitable features for downstream tasks.

**The effectiveness of the AimCLR.** We conduct experiments on three datasets to verify the performance of our AimCLR and the SkeletonCLR. As can be seen from Table 2, for the three different streams of the three datasets, AimCLR performs much better than SkeletonCLR. The gains in motion and bone stream are significant. For the fusion results, our 3s-AimCLR far exceeds 3s-SkeletonCLR on the three datasets. In addition, it can be seen from Table 3 that our 3s-AimCLR is always better than 3s-CrosSCLR and 3s-SkeletonCLR under the same training epochs no matter us-



Method	part I(%)	part II(%)
<i>Supervised:</i>		
ST-GCN (AAAI 18)	84.1	48.2
VA-LSTM (TPAMI 19)	84.1	50.0
<i>Self-supervised:</i>		
LongT GAN (AAAI 18)	67.7	26.0
MS <sup>2</sup> L (ACM MM 20)	64.9	27.6
3s-CrosSCLR (CVPR 21) <sup>§</sup>	84.9	21.2
ISC (ACM MM 21)	80.9	36.0
<b>3s-AimCLR (Ours)<sup>§</sup></b>	<b>87.4</b>	<b>39.5</b>
<b>3s-AimCLR (Ours)</b>	<b>87.8</b>	<b>38.5</b>

Table 5: Linear evaluation results on PKU-MMD dataset.

Method	xsub(%)	xset(%)
P&C (CVPR 20)	42.7	41.7
AS-CAL (Information Sciences 21)	48.6	49.2
3s-CrosSCLR (CVPR 21) <sup>§</sup>	67.9	66.7
ISC (ACM MM 21)	67.9	67.1
<b>3s-AimCLR (Ours)<sup>§</sup></b>	<b>68.0</b>	<b>68.7</b>
<b>3s-AimCLR (Ours)</b>	<b>68.2</b>	<b>68.8</b>

Table 6: Linear evaluation results on NTU-120 dataset.

ing cross-stream knowledge mining strategy or not. The result of 3s-AimCLR at 100 epochs is even better than the result of 3s-SkeletonCLR at 300 epochs.

**Qualitative Results.** We apply t-SNE (Van der Maaten and Hinton 2008) with fix settings to show the embedding distribution in Figure 2. The reported t-SNE results are fair comparisons under the same training epoch with the same randomly selected 10 class samples. From the visual results, we can draw a conclusion that 3s-AimCLR could cluster the embeddings of the same class closer than 3s-SkeletonCLR. While using the cross-stream knowledge mining strategy, our 3s-AimCLR<sup>§</sup> can also make the action classes that overlapped seriously more distinguishable compared with 3s-CrosSCLR<sup>§</sup>. We also calculated the NMI (Normalized Mutual Information) to support our claims: (c)3s-CrosSCLR<sup>§</sup>: 0.6668, (f)3s-AimCLR<sup>§</sup>: 0.6793. The huge gains under multiple evaluation protocols in Section also corroborate the authenticity of the t-SNE results.

### Comparison with State-of-the-art

We compare the proposed method with prior related methods under a variety of evaluation protocols.

**KNN Evaluation Results.** Notably, the KNN classifier does not require learning extra weights compared with the linear classifier. From Figure 3, we can see that our AimCLR is better than SkeletonCLR on the two datasets under the KNN classifier. The obvious gains also show that the features learned by our AimCLR are more discriminative.

**Linear Evaluation Results on NTU-60.** As shown in Table 4, for a single stream (i.e., joint stream), our AimCLR outperforms all other methods (Zheng et al. 2018; Lin et al. 2020; Rao et al. 2021; Su, Liu, and Shlizerman 2020; Nie, Liu, and Liu 2020; Li et al. 2021). For the performance of the

Method	PKU-MMD(%)		NTU-60(%)	
	part I	part II	xsub	xview
<i>1% labeled data:</i>				
LongT GAN (AAAI 18)	35.8	12.4	35.2	-
MS <sup>2</sup> L (ACM MM 20)	36.4	13.0	33.1	-
ISC (ACM MM 21)	37.7	-	35.7	38.1
3s-CrosSCLR (CVPR 21)	49.7	10.2	51.1	50.0
3s-Colorization (ICCV 21)	-	-	48.3	52.5
<b>3s-AimCLR (Ours)</b>	<b>57.5</b>	<b>15.1</b>	<b>54.8</b>	<b>54.3</b>
<i>10% labeled data:</i>				
LongT GAN (AAAI 18)	69.5	25.7	62.0	-
MS <sup>2</sup> L (ACM MM 20)	70.3	26.1	65.2	-
ISC (ACM MM 21)	72.1	-	65.9	72.5
3s-CrosSCLR (CVPR 21)	82.9	28.6	74.4	77.8
3s-Colorization (ICCV 21)	-	-	71.7	78.9
<b>3s-AimCLR (Ours)</b>	<b>86.1</b>	<b>33.4</b>	<b>78.2</b>	<b>81.6</b>

Table 7: Semi-supervised evaluation results on PKU-MMD dataset and NTU-60 dataset.

Method	NTU-60(%)		NTU-120(%)	
	xsub	xview	xsub	xset
SkeletonCLR (CVPR 21) <sup>†</sup>	82.2	88.9	73.6	75.3
<b>AimCLR (Ours)<sup>†</sup></b>	<b>83.0</b>	<b>89.2</b>	<b>76.4</b>	<b>76.7</b>
3s-ST-GCN (AAAI 18)	85.2	91.4	77.2	77.1
3s-CrosSCLR (CVPR 21)	86.2	92.5	<b>80.5</b>	80.4
<b>3s-AimCLR (Ours)</b>	<b>86.9</b>	<b>92.8</b>	80.1	<b>80.9</b>

Table 8: Finetuned results on NTU-60 and NTU-120 dataset. “†” means using the bone stream data.

3-streams, our 3s-AimCLR leads 3s-SkeletonCLR 3.9% and 4.0% under the xsub and xview protocols, respectively. It is worth mentioning that regardless of whether 3s-AimCLR uses a cross-stream knowledge mining strategy, the results are better than the 3s-CrosSCLR and 3s-Colorization (Yang et al. 2021b). It also indicates that even without the knowledge mining between streams, 3s-AimCLR has the ability to learn better feature representations.

**Linear Evaluation Results on PKU-MMD.** As shown in Table 5, our 3s-AimCLR is ahead of the existing self-supervised methods in both part I and part II of this dataset. Part II is more challenging with more skeleton noise caused by the view variation. Notably, 3s-CrosSCLR suffers on part II while our 3s-AimCLR performs well. It also proves that our 3s-AimCLR has a strong ability to cope with movement patterns caused by skeleton noise.

**Linear Evaluation Results on NTU-120.** As shown in Table 6, our 3s-AimCLR defeats the other self-supervised methods on NTU-120. Our fusion results outperforms the advanced ISC (68.2% vs 67.9% on xsub and 68.8% vs 67.1% on xset). This shows that our 3s-AimCLR is also competitive on multi-class large-scale datasets.

**Semi-supervised Evaluation Results.** From Table 7, even only with a small labeled subset, our 3s-AimCLR performs better than the state-of-the-art consistently for all configurations. The results of using 1% and 10% labeled data

far exceed ISC, 3s-CrosSCLR, and 3s-Colorization. It also proves that the novel movement patterns brought by extreme augmentations have a huge impact when there is only a small amount of labeled data.

**Finetuned Evaluation Results.** For fair comparisons, the ST-GCN used in the methods of Table 8 all have the same structure and parameters. For a single bone stream, the finetuned results of our AimCLR are better than that of SkeletonCLR. What’s more, the finetuned 3s-AimCLR also outperforms the 3s-CrosSCLR and the supervised 3s-ST-GCN, indicating the effectiveness of our method.

## Conclusion

In this paper, AimCLR is proposed to explore the novel movement patterns brought by extreme augmentations. Specifically, the extreme augmentations and the energy-based attention-guided drop module are proposed to bring novel movement patterns to improve the universality of the learned representations. The  $D^3M$  Loss is proposed to minimize the distribution divergence in a more gentle way. In order to alleviate the irrationality of the positive set, the nearest neighbors mining strategy is further proposed to make the learning process more reasonable. Experiments show that 3s-AimCLR significantly performs favorably against state-of-the-art methods under a variety of evaluation protocols with observed higher quality action representations. In the future, we hope to explore more effective combinations of extreme augmentations and more effective ways to interact information between different streams.

## Acknowledgments

This work is supported by National Key R&D Program of China (No. 2020AAA0108904), Science and Technology Plan of Shenzhen (No. JCYJ20200109140410340).

## References

Cao, Z.; Hidalgo, G.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2019. OpenPose: Realtime multi-person 2D pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 43(1): 172–186.

Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020a. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning (ICML)*, 1597–1607.

Chen, X.; Fan, H.; Girshick, R.; and He, K. 2020b. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*.

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 *AAAI Conference on Artificial Intelligence*, volume 35, 1113–1122.

Cheng, K.; Zhang, Y.; Cao, C.; Shi, L.; Cheng, J.; and Lu, H. 2020. Decoupling GCN with dropgraph module for skeleton-based action recognition. In *European Conference on Computer Vision (ECCV)*, 536–553.

Du, Y.; Fu, Y.; and Wang, L. 2015. Skeleton based action recognition with convolutional neural network. In *Asian Conference on Pattern Recognition (ACPR)*, 579–583.

Du, Y.; Wang, W.; and Wang, L. 2015. Hierarchical recurrent neural network for skeleton based action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1110–1118.

Dwibedi, D.; Aytar, Y.; Tompson, J.; Sermanet, P.; and Zisserman, A. 2021. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. In *IEEE International Conference on Computer Vision (ICCV)*, 9588–9597.

Fang, H.-S.; Xie, S.; Tai, Y.-W.; and Lu, C. 2017. RMPE: Regional multi-person pose estimation. In *IEEE International Conference on Computer Vision (ICCV)*, 2334–2343.

Gidaris, S.; Singh, P.; and Komodakis, N. 2018. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*.

He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 9729–9738.

Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 7132–7141.

Ke, Q.; Bennamoun, M.; An, S.; Sohel, F.; and Boussaid, F. 2017. A new representation of skeleton sequences for 3D action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3288–3297.

Lee, H.; Kim, H.-E.; and Nam, H. 2019. SRM: A style-based recalibration module for convolutional neural networks. In *IEEE International Conference on Computer Vision (ICCV)*, 1854–1862.

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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4741–4750.

Lin, L.; Song, S.; Yang, W.; and Liu, J. 2020. MS<sup>2</sup>L: Multi-task self-supervised learning for skeleton based action recognition. In *ACM International Conference on Multimedia (ACM MM)*, 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*, 16(2): 1–24.

Liu, M.; Liu, H.; and Chen, C. 2017. Enhanced skeleton visualization for view invariant human action recognition. *Pattern Recognition*, 68: 346–362.

Nie, Q.; Liu, Z.; and Liu, Y. 2020. Unsupervised 3D human pose representation with viewpoint and pose disentanglement. In *European Conference on Computer Vision (ECCV)*, 102–118.



- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Pan, T.; Song, Y.; Yang, T.; Jiang, W.; and Liu, W. 2021. VideoMoCo: Contrastive video representation learning with temporally adversarial examples. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 11205–11214.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems (NeurIPS)*, 32: 8026–8037.
- Pathak, D.; Krahenbuhl, P.; Donahue, J.; Darrell, T.; and Efros, A. A. 2016. Context encoders: Feature learning by inpainting. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2536–2544.
- 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1010–1019.
- Shi, L.; Zhang, Y.; Cheng, J.; and Lu, H. 2019. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 12026–12035.
- Si, C.; Chen, W.; Wang, W.; Wang, L.; and Tan, T. 2019. An attention enhanced graph convolutional LSTM network for skeleton-based action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1227–1236.
- Song, S.; Lan, C.; Xing, J.; Zeng, W.; and Liu, J. 2018. Spatio-temporal attention-based LSTM networks for 3D action recognition and detection. *IEEE Transactions on Image Processing (TIP)*, 27(7): 3459–3471.
- Su, K.; Liu, X.; and Shlizerman, E. 2020. Predict & Cluster: Unsupervised skeleton based action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 9631–9640.
- Thoker, F. M.; Doughty, H.; and Snoek, C. G. 2021. Skeleton-contrastive 3D action representation learning. In *ACM International Conference on Multimedia (ACM MM)*.
- Tian, Y.; Sun, C.; Poole, B.; Krishnan, D.; Schmid, C.; and Isola, P. 2020. What makes for good views for contrastive learning? *arXiv preprint arXiv:2005.10243*.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(11).
- Vemulapalli, R.; Arrate, F.; and Chellappa, R. 2014. Human action recognition by representing 3D skeletons as points in a lie group. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 588–595.
- Vemulapalli, R.; and Chellappa, R. 2016. Rolling rotation-s for recognizing human actions from 3D skeletal data. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4471–4479.
- Wang, J.; Liu, Z.; Wu, Y.; and Yuan, J. 2012. Mining action-let ensemble for action recognition with depth cameras. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1290–1297.
- Wang, X.; and Qi, G.-J. 2021. Contrastive learning with stronger augmentations. *arXiv preprint arXiv:2104.07713*.
- Woo, S.; Park, J.; Lee, J.-Y.; and Kweon, I. S. 2018. CBAM: Convolutional block attention module. In *European Conference on Computer Vision (ECCV)*, 3–19.
- Yan, S.; Xiong, Y.; and Lin, D. 2018. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI Conference on Artificial Intelligence*, volume 32.
- Yang, L.; Zhang, R.-Y.; Li, L.; and Xie, X. 2021a. SimAM: A simple, parameter-free attention module for convolutional neural networks. In *International Conference on Machine Learning (ICML)*, 11863–11874.
- Yang, S.; Liu, J.; Lu, S.; Er, M. H.; and Kot, A. C. 2021b. Skeleton cloud colorization for unsupervised 3D action representation learning. In *IEEE International Conference on Computer Vision (ICCV)*.
- Zhang, P.; Lan, C.; Xing, J.; Zeng, W.; Xue, J.; and Zheng, N. 2019. View adaptive neural networks for high performance skeleton-based human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 41(8): 1963–1978.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *European Conference on Computer Vision (ECCV)*, 649–666.
- 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 *AAAI Conference on Artificial Intelligence*, volume 32.