

Contrastive Self-Supervised Learning for Skeleton Action Recognition

Xuehao Gao

Yang Yang *

Shaoyi Du[†]

Xi'an Jiaotong University, China

GAOXUEHAO.XJTU@GMAIL.COM

YYANG@MAIL.XJTU.EDU.CN

DUSHAOYI@XJTU.EDU.CN

Abstract

Learning discriminative features plays a significant role in action recognition. Many attempts have been made to train deep neural networks by their labeled data. However, in previous networks, the view or distance variations can cause the intra-class differences even larger than inter-class differences. In this work, we propose a new contrastive self-supervised learning method for action recognition of unlabeled skeletal videos. Through contrastive representation learning by adequate compositions of viewpoints and distances, the self-supervised net selects discriminative features which have invariance motion semantics for action recognition. We hope this attempt can be helpful for the unsupervised learning study of skeleton-based action recognition.

1. Introduction

Action understanding is a fundamental study in the computer vision field. Compared with RGB images, body joint time-series (skeletons) are effective descriptors of actions, which are robust against the background and lighting changes. Since each joint is easily identified by a 2D or 3D position vector, skeleton sequence becomes a high-level and abstract representation of an action. Although recent methods have achieved remarkable progress with the development of deep neural networks, most methods rely on strong supervision on action labels. Large-scale data annotation is laborious and expensive, even impractical for complex data such as videos. Furthermore, annotation is a challenging problem by itself, since some action classes are ambiguous, which are up to the interpretation of each annotator for a given sequence. Thus, unsupervised methods that do not use labeled data are necessary for skeleton action recognition.

Learning effective motion representations without human supervision is a long-standing and challenging problem. Existing approaches [Su et al. \(2020\)](#) are mostly based on the encoder-decoder architecture. Specifically, given a skeleton action sequence as the encoder input, the decoder predicts the encoder's input sequence. But this idea needs to re-generate each frame in a sequence, which is computationally expensive and often ignores the semantic association between the encoder input and the decoder output. In that case, the generality of the learned representation is limited. Furthermore, the auto-encoder will not be able to achieve efficient performances without particular training strategies [Su et al. \(2020\)](#). Inspired by the recent successful discriminative approaches in the latent space [Bachman et al. \(2019\)](#); [Dosovitskiy et al. \(2014\)](#); [Chen et al. \(2020\)](#), here we try

* X.Gao and Y.Yang contributed equally to this work.

[†] Corresponding author.

to obtain effective feature representations of skeleton sequences under a simple contrast learning framework.

Previous approaches [Zhang et al. \(2019\)](#); [Li et al. \(2018\)](#) have proven that one of the key challenges in action recognition lies in the large variety of action representation when motions are captured from different viewpoints. But, in reality, the same action can be easily recognized by two observers when they stand at different viewpoints or distances. This observation inspires us to obtain effective feature representations under a contrastive self-supervised learning framework. Accordingly, we define contrastive prediction tasks by the composition of multiple data augmentation operations (*i.e.* view variations). And then we maximize the agreement between different augmented views of the same data example via a contrastive loss in the latent space. Through contrastive representation learning by adequate compositions of viewpoints and distances, the self-supervised net selects discriminative features which have invariant motion semantics for action recognition.

2. Related Work

Numerous approaches have been introduced particularly for human action recognition. Most of them are supervised where an annotated set of actions and labels should be provided for training. In an unsupervised setup, the problem of action recognition is much more challenging and only a few methods have been proposed. In this section, we give a brief review on the unsupervised method and the skeleton action recognition by variant of viewpoint and distance.

2.1. Unsupervised Learning in Skeleton Action Recognition

The unsupervised setup has advantages for action recognition since it does not require the labeling of sequences. When additional action appears, the unsupervised networks do not need re-training. This kind of method typically aims to obtain an effective feature representation by predicting or regenerating future frames of input sequences. Srivastava et al. [Srivastava et al. \(2015\)](#) proposed a recurrent-based sequence (Seq2Seq) model as an autoencoder to learn the representation of a video. Zheng et al. [Zheng et al. \(2018\)](#) proposed a GAN encoder-decoder (LongTGAN). The decoder attempts to re-generate the input sequence and the discriminator is used to discriminate whether the re-generation is accurate. Su et al. [Su et al. \(2020\)](#) make the decoder and the encoder self-organize their hidden states into a feature space which clusters similar movements into the same cluster and distinct movements into distant clusters. The feature representation used for action recognition is taken from the final state of the encoder’s hidden representation. However, these typical encoder-decoder-based unsupervised methods are computationally expensive and rely on particular training strategies. We explore to learn robust and efficient features representation based on a simpler yet effective idea in the contrastive learning framework.

2.2. Variant of Viewpoints and Distance in Skeleton Action Recognition

For humans, it will not affect us to recognize the action class from arbitrary observing viewpoints and distances. Human actions can be captured from arbitrary camera setups, while it is a challenging problem for recognition algorithms. Researchers have paid much attention to this issue and proposed various view-invariant approaches. Zhang et al. [Zhang et al. \(2019\)](#) proposed a view-adaptive idea to leverage content-dependent and view-adaptation models to automatically learn and

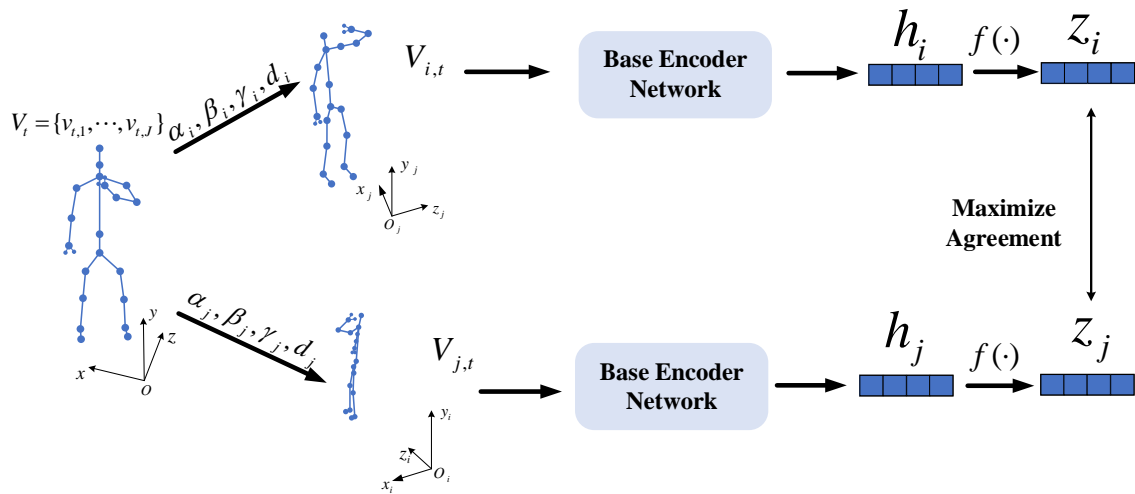


Figure 1: The framework for contrastive learning of skeleton sequence using the semantic invariance under the variant of distance and viewpoint. Two separate transformations of viewpoint and distance are applied to a given skeleton sequence to obtain a correlated pair. A base encoder network and a projection function $f(\cdot)$ are trained to maximize agreement by using a contrastive loss.

determine the suitable viewpoint. And for each sequence, the skeletons are transformed into representations under those views. This improvement has proved that viewpoints are crucial for the model to learn features representation in skeleton-based action recognition. Furthermore, the distance of a subject to the camera may influence the scale of the skeletons. In order to learn the robust and effective features representation in an unannotated skeleton sequence, we design a simple contrastive learning method based on the idea of semantic invariance under the variant of distance and viewpoint.

3. Methodology

We learn our feature representations by maximizing agreement between different viewpoints and distance setups of the same skeleton sequence via the contrastive loss in the latent space. As illustrated in Figure 1, we just using the base encoder network (e.g. *ResNet* He et al. (2016)) and representation vector h for following action recognition task.

3.1. View Augmentation for Contrastive Representation Learning

The raw 3D skeleton is recorded corresponding to the camera coordinate system (global coordinate system), with the origin located at the position of the camera sensor. Note that the input skeleton V_t to our system as in Figure 1 is the skeleton representation under this initial camera coordinate system. To be insensitive to the initial position of action, for each sequence, we translate the global coordinate system to the body center of the first frame as our new global coordinate system. One can choose to observe the action from stochastic positions. Thanks to the availability of the 3D skeletons

captured from a fixed view, it is possible to set up a movable virtual camera and observe the action from new observation viewpoints. The given skeleton can be transformed into a representation under the movable virtual camera coordinate system, which is also referred to as the observation coordinate system.

Given a skeleton sequence \mathcal{S} with T frames, under the global coordinate system, we denote the set of joints in the t -th frame as $\mathbf{V}_t = \{\mathbf{v}_{t,1}, \dots, \mathbf{v}_{t,J}\}$. For the t -th frame, we assume the movable virtual camera is placed at a stochastic viewpoint and distance, with the corresponding observation coordinate system obtained from a translation by $\mathbf{d}_t \in \mathbb{R}^3$, and a rotation of $\alpha_t, \beta_t, \gamma_t$ radians anticlockwise around the x -axis, y -axis, and z -axis, respectively, of the global coordinate system. Therefore, the representation of the j -th skeleton joint $\mathbf{v}'_{t,j} = [x'_{t,j}, y'_{t,j}, z'_{t,j}]^T$ of the t -th frame under this observation coordinate system \mathcal{O} is

$$\mathbf{v}'_{t,j} = [x'_{t,j}, y'_{t,j}, z'_{t,j}]^T = \mathbf{R}_t \times (\mathbf{v}_{t,j} - \mathbf{d}_t). \quad (1)$$

\mathbf{R}_t can be represented as

$$\mathbf{R}_t = \mathbf{R}_{t,\alpha}^x \times \mathbf{R}_{t,\beta}^y \times \mathbf{R}_{t,\gamma}^z, \quad (2)$$

where $\mathbf{R}_{t,\beta}^y$ denotes the coordinate transformation for rotating the original coordinate system around the y -axis by β_t radians anticlockwise, which is defined as

$$\mathbf{R}_{t,\beta}^y = \begin{bmatrix} \cos(\beta_t) & \sin(\beta_t) & 0 \\ -\sin(\beta_t) & \cos(\beta_t) & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (3)$$

Similarly, $\mathbf{R}_{t,\alpha}^x$ and $\mathbf{R}_{t,\gamma}^z$ denote the coordinate transforms for rotating the original coordinate system around the x -axis by α_t radians, and around the z -axis by γ_t radians anticlockwise, respectively. Note that all the skeleton joints in the t -th frame share the same transform parameters, *i.e.*, $\alpha_t, \beta_t, \gamma_t, \mathbf{d}_t$, considering that the changing of viewpoints and distance is a rigid transformation. Given these transformation parameters, the skeleton representation $\mathbf{V}_t = \{\mathbf{v}_{t,1}, \dots, \mathbf{v}_{t,J}\}$ under the new observation coordinate can be obtained from Eq.1. Besides, the viewpoint and distance can be varied for different frames or on a sequence level. The key problem becomes how to obtain an effective feature representations in skeleton sequence under the contrast learning framework by adequate compositions of viewpoints and distances.

3.2. The Contrastive Learning Framework

Firstly, given a skeleton sequence \mathbf{V}_t that is captured at an initial position and orientation, we adopt two stochastic transforms setups on it to provide a positive pair, denoted $\mathbf{V}_{i,t}$ and $\mathbf{V}_{j,t}$. Secondly, a base encoder network that extracts representation vectors \mathbf{h} from transformed skeleton sequences. Similar to other typical contrastive learning methods, this simple framework design conveniently decouples the predictive task from other components such as the base encoder network architecture, and this base encoder network is shared by both branches of our framework. The base encoder network serves as the skeleton-based motion features extractor. Here, we can adopt the commonly used *ResNet* He et al. (2016) scheme to extract spatial and temporal features. Finally, projection function $f(\cdot)$ will map representations to space where contrastive loss is applied. We opt for simplicity and use 2-layer MLP as a projection function. The base encoder network and projection

function are trained to maximize the agreement by using contrastive loss. After that, we just use the pretrained base encoder network and the representation vector \mathbf{h} for downstream tasks.

We randomly sample a minibatch of N examples and define the contrastive prediction task on pairs of view-augmented examples derived from the minibatch, resulting in $2N$ data points. For a given positive pair, we treat other $2(N - 1)$ as negative samples. The contrastive loss function is used to maximize the agreement between positive pair, meanwhile minimizing the agreement between negative pairs. Here, we adopt a common formulation of loss function used in the previous contrastive learning work. Then the loss function of a positive pair of examples (i, j) is defined as

$$\ell_{i,j} = -\log \frac{\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_j) / \tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_k) / \tau)}, \quad (4)$$

where $\text{sim}(\mathbf{z}_i, \mathbf{z}_j) = \mathbf{z}_i^\top \mathbf{z}_j / \|\mathbf{z}_i\| \|\mathbf{z}_j\|$ denotes the dot product between ℓ_2 normalized \mathbf{z}_i and \mathbf{z}_j (*i.e.* cosine similarity). $\mathbb{1}_{[k \neq i]}$ is an indicator function evaluating to 1 iff $k \neq i$ and τ denotes a temperature parameter. After that the final loss function is

$$\mathcal{L} = \frac{1}{2N} \sum_{k=1}^N [\ell(2k - 1, 2k) + \ell(2k, 2k - 1)]. \quad (5)$$

As illustrated in Eq.5, the final loss is calculated across all positive pairs, both (i, j) and (j, i) .

4. Experimental Protocol

Here, we lay out the protocol for our empirical studies, which aims to understand different design choices in our framework and improve representation quality by combining our findings.

4.1. Dataset and Metrics

Following the above description in the data augmentation, skeleton action samples will be rotated in three-dimensional space. Thus some datasets (e.g. Kinetics Skeleton 400 dataset) are inappropriate to evaluation since their each skeleton joints are recorded along with their 2D spatial coordinates. For learning feature representations without labels by the base encoder networks, we conduct unsupervised pretraining on NTU-RGB+D 60 and NTU-RGB+D 120 dataset. To evaluate the learned representations, we follow the widely used linear evaluation protocol, where a linear classifier is trained on the representation vector obtained from the frozen base neural network. The test accuracy is used as a proxy for the representation quality.

4.1.1. NTU-RGB+D 60 DATASET

The NTU RGB+D 60 [Shahroudy et al. \(2016\)](#) is a large-scale action recognition dataset collected by three Kinect cameras simultaneously. It contains 56,578 skeleton sequences and 60 action classes performed by 40 different subjects. Each human skeleton is represented by 25 joints with 3D coordinates. The authors recommend evaluating the model performance under two settings: (1) Cross-Subject(X-Sub), where half of 40 subjects are used for training and the rest for testing. (2) Cross-View(X-View), where the sequences captured by two of three cameras are used for training, and those captured by the last cameras are used for testing.

4.1.2. NTU-RGB+D 120 DATASET

This dataset [Liu et al. \(2019\)](#) is an extension of NTU-RGB+D 60 and currently the largest dataset for 3D action recognition. It contains 114,480 skeleton sequences and 120 action classes performed by 106 different subjects. It provides two types of evaluating setting: (1) Cross-Subject(X-Sub), where half of 106 subjects are used for training and the rest for testing. (2) Cross-Setup(X-Set), where the sequences captured by half of cameras are used for training and the rest for testing.

4.2. Default Setting

Benefitting from the experiences of several previous contrast learning methods, we realize that contrastive learning benefits from the larger batch size and more training steps compared to supervised learning. Thus, we will conduct training experiments with slightly larger batch sizes and steps purposefully. However, training with a large batch size may be unstable when using standard SGD/Momentum with linear learning rate scaling. To stabilize the training, we use the LARS optimizer for all batch sizes. Besides, in distributed training with data parallelism, the mean and variance of Batch Normalization (BN) are typically aggregated locally per device. In contrastive learning, as positive pairs are computed in the same device, the model can exploit the local information leakage to improve the prediction accuracy without improving representations. We address this issue by aggregating BN mean and variance overall devices during the training.

4.3. Variant of Viewpoints and Distance for Contrastive Representation Learning

Previous work has realized that viewpoint is crucial for features representation in supervised methods. Furthermore, the distance of a subject to the camera does influence the scale of the skeletons. To exploit this semantic invariance under the variants of distance and viewpoint, we conduct a contrastive learning framework for learning robust feature representations in an unannotated skeleton sequence. In the following experiments, we will adopt variants of viewpoints and distance to define the contrastive prediction task in a systematic way. We consider that frame-wise transformation is inferior to the sequence level transformation because the former loses more information, *e.g.* the motion across frames. Thus, we will conduct transformations on the sequence level in following experiments. To comprehensively study the impact of viewpoint and distance, we define positions of raw 3D skeletons are initial state, and provide several viewpoints and distance transformations, respectively. For each viewpoint setup, α, β, γ are randomly selected from the degree set (*i.e.* $30^\circ, 60^\circ, 90^\circ, 180^\circ$). For each distance setup, d is a random times the unit distance where the alternative multiplier is $1 \sim 5\times$. To understand the effects of individual transformation of the skeleton and the importance of transformations composition, the experiment will be repeated multiple times:

- on stochastic viewpoints and distance setups with alternative α, β, γ, d options.
- on stochastic viewpoints setups with random α, β pair and a fixed γ .
- on stochastic viewpoints setups with random α, γ pair and a fixed β .
- on stochastic viewpoints setups with random β, γ pair and a fixed α .
- on stochastic distance setups with a fixed viewpoints setup.

Table 1: Ablation for models trained with compositions of data augmentation.

Viewpoint Augmentation	Distance Augmentation	Accuracy	
		Top-1	Top-5
✓		35.0%	68.0%
	✓	23.1%	45.6%
✓	✓	45.0%	78.3%

Table 2: NTU-RGB+D 60 X-Sub accuracies of linear classifiers trained on representations learned with different viewpoints and distance setups.

Accuracy \ Distance \ Viewpoint	1×	2×	3×	4×	5×
[-30°, 30°]	43.7%	44.1%	45.0%	45.5%	46.8%
[-60°, 60°]	44.0%	44.5%	45.5%	45.0%	47.0%
[-90°, 90°]	45.7%	46.0%	46.7%	46.9%	46.5%
[-180°, 180°]	47.3%	47.9%	48.0%	47.4%	47.2%

- with base encoder architectures of different complexity (*i.e.* models with varied depth and width).

5. Experimental Results

5.1. Ablation study

To verify the effectiveness of each component of our method, we conduct extensive ablation studies following the above experimental protocol. Unless stated, the performance reported in the ablation study is based on 20 epochs pre-training on the NTU-RGB+D 60 X-Sub dataset.

Effects of Data Augmentation. To understand the effects of individual augmentation of the skeleton and the importance of augmentations composition, we pre-train the model with different augmentation choices. As shown in the Table 1, we observe that no single transformation suffices to learn good representations, and both viewpoint and distance augmentations are indispensable. Specifically, using both viewpoint and distance augmentations yields 45.0% top-1 accuracy. To further understand the effects of various viewpoint and distance setups on the final representation, we

Table 3: Ablation for models trained with different choices of rotation coordinate axis and angle range for viewpoint augmentation.

Accuracy \ Angle Range \ Coordinate Axis	[-30°,30°]	[-60°,60]	[-90°,90°]	[-180°,180°]
x	35.0%	36.1%	37.2%	37.9%
x, y	38.4%	39.7%	39.4%	38.4%
x, y, z	39.2%	41.5%	40.6%	39.7%

Table 4: NTU60 X-Sub accuracies of linear classifiers trained on representations learned with different encoder architectures.

Encoder Architecture	# Parameters	Accuracy
ResNet-18	12.2 M	40.2%
ResNet-50	26.5 M	47.1%
ResNet-101	45.1 M	52.3%

Table 5: Comparison with hand-crafted, supervised and unsupervised methods on NTU60 dataset. * represents depth image based methods.

Method	X-View Accuracy(%)	X-Sub Accuracy(%)
Hand-Crafted Methods		
HON4D* Oreifej and Liu (2013)	7.3	30.6
Super Normal Vector* Yang and Tian (2014)	13.6	31.8
HOG ² * Ohn-Bar and Trivedi (2013)	13.6	31.8
Skeletal Quads Evangelidis et al. (2014)	41.1	38.6
HOPC Rahmani et al. (2014)	52.8	50.1
Supervised Methods		
H-RNN Du et al. (2015b)	64.0	59.1
Deep RNN Shahroudy et al. (2016)	64.1	56.3
Part-aware LSTM Shahroudy et al. (2016)	70.3	62.9
Unsupervised Methods		
Shuffle & Learn* Misra et al. (2016)	40.9	46.2
LongT GAN Zheng et al. (2018)	48.1	39.1
P&C FW Su et al. (2020)	61.1	50.6
Ours	62.1	52.3

show the accuracy performance pre-trained on each specific setup in the Table 2. We can see that moderate data augmentation helps the encoder to learn more task-relevant information. Additional noise introduced by the harder data augmentation may lead to degraded performance. Observed from the Table 3 which shown the influence of rotation axis and rotation angle range in the view-point augmentation, small rotation with multiple angles of a skeleton sequence can provide the optimal multiview for learning action-specific features in contrastive representation.

Effects of Base Encoder. The simple design choice of our method is to conveniently decouple the defined predictive task from other components such as the neural network architecture of the base encoder. Here, we use ResNet [He et al. \(2016\)](#) as an alternative encoder to extract representations. In addition, a skeletal sequence is transformed into an image before inputting this CNN-based encoder, similar that in [Du et al. \(2015a\)](#). Observed from the Table 4, we can see that the best result obtain from the ResNet-101. Thus, the performance of contrastive learning benefits from bigger models.

Table 6: Comparison with supervised and unsupervised methods on NTU120 dataset.

Method	X-Set Accuracy(%)	X-Sub Accuracy(%)
Supervised Methods		
Part-Aware LSTM Shahroudy et al. (2016)	26.3	25.5
Soft RNN Hu et al. (2018)	44.9	36.3
ST-LSTM+Trust Gate Liu et al. (2016)	57.9	55.7
Unsupervised Methods		
P&C FW Su et al. (2020)	42.7	41.7
Ours	44.3	43.5

5.2. Performance Comparison

We compare our method with previous relevant unsupervised learning methods, supervised methods, and hand-crafted methods on NTU60 and NTU120 datasets. To make the comparisons fair and convincing, for an unsupervised learning method P&C [Su et al. \(2020\)](#), we implement it on the same linear evaluation protocol instead of KNN evaluations.

As shown in Table 5, our method enjoys evident advantages over existing unsupervised methods on the NTU-RGB+D 60 dataset. Compared with re-generative learning methods (e.g. P&C [Su et al. \(2020\)](#)), our method not only has a simpler framework, but also yields better performance. Similarly, our method also surpasses P&C on the X-Set and X-Sub tests of NTU120 dataset with the 1.6% and 1.8 % accuracy improvement respectively, as shown in Table 6.

6. Acknowledgments

This work was supported by the National Key Research and Development Program of China under Grant No. 2018AAA0102500 and the National Natural Science Foundation of China under Grant No. 61971343. We would like to thank Maosen Li for his feedback on the draft.

References

- Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *arXiv preprint arXiv:1906.00910*, 2019.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. Citeseer, 2014.
- Yong Du, Yun Fu, and Liang Wang. Skeleton based action recognition with convolutional neural network. In *3rd IAPR Asian Conference on Pattern Recognition, ACPR 2015*, pages 579–583. IEEE, 2015a.
- Yong Du, Wei Wang, and Liang Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1110–1118, 2015b.

- Georgios Evangelidis, Gurkirt Singh, and Radu Horaud. Skeletal quads: Human action recognition using joint quadruples. In *2014 22nd International Conference on Pattern Recognition*, pages 4513–4518. IEEE, 2014.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Jian-Fang Hu, Wei-Shi Zheng, Lianyang Ma, Gang Wang, Jianhuang Lai, and Jianguo Zhang. Early action prediction by soft regression. *IEEE transactions on pattern analysis and machine intelligence*, 41(11): 2568–2583, 2018.
- Junnan Li, Yongkang Wong, Qi Zhao, and Mohan S Kankanhalli. Unsupervised learning of view-invariant action representations. *arXiv preprint arXiv:1809.01844*, 2018.
- Jun Liu, Amir Shahroudy, Dong Xu, and Gang Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In *European conference on computer vision*, pages 816–833. Springer, 2016.
- Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C Kot. 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, 2019.
- Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision*, pages 527–544. Springer, 2016.
- Eshed Ohn-Bar and Mohan Trivedi. Joint angles similarities and hog2 for action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition workshops*, pages 465–470, 2013.
- Omar Oreifej and Zicheng Liu. Hon4d: Histogram of oriented 4d normals for activity recognition from depth sequences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 716–723, 2013.
- Hossein Rahmani, Arif Mahmood, Du Q Huynh, and Ajmal Mian. Hopc: Histogram of oriented principal components of 3d pointclouds for action recognition. In *European Conference on Computer Vision*, pages 742–757. Springer, 2014.
- Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1010–1019, 2016.
- Nitish Srivastava, Elman Mansimov, and Ruslan Salakhudinov. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, pages 843–852. PMLR, 2015.
- Kun Su, Xiulong Liu, and Eli Shlizerman. Predict & cluster: Unsupervised skeleton based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9631–9640, 2020.
- Xiaodong Yang and YingLi Tian. Super normal vector for activity recognition using depth sequences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 804–811, 2014.
- Pengfei Zhang, Cuiling Lan, Junliang Xing, Wenjun Zeng, Jianru Xue, and Nanning Zheng. View adaptive neural networks for high performance skeleton-based human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1963–1978, 2019.

Nenggan Zheng, Jun Wen, Risheng Liu, Liangqu Long, Jianhua Dai, and Zhefeng Gong. Unsupervised representation learning with long-term dynamics for skeleton based action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.