

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Bootstrapped Representation Learning for Skeleton-Based Action Recognition

Olivier Moliner^{1,2} Sangxia Huang² Kalle Åström¹

¹Centre for Mathematical Sciences, Lund University, Lund, Sweden

²R&D Center Lund Laboratory, Sony Europe B.V., Lund, Sweden

{olivier.moliner, karl.astrom}@math.lth.se, sangxia.huang@sony.com

Abstract

In this work, we study self-supervised representation learning for 3D skeleton-based action recognition. We extend Bootstrap Your Own Latent (BYOL) for representation learning on skeleton sequence data and propose a new data augmentation strategy including two asymmetric transformation pipelines. We also introduce a multi-viewpoint sampling method that leverages multiple viewing angles of the same action captured by different cameras. In the semisupervised setting, we show that the performance can be further improved by knowledge distillation from wider networks, leveraging once more the unlabeled samples. We conduct extensive experiments on the NTU-60, NTU-120 and PKU-MMD datasets to demonstrate the performance of our proposed method. Our method consistently outperforms the current state of the art on linear evaluation, semisupervised and transfer learning benchmarks.

1. Introduction

Action recognition is an essential task in computer vision with many real-world applications, such as surveillance and assisted living, human-robot interaction, video retrieval and autonomous driving. Recent advances in sensors [42] and human pose estimation methods [2,47] make it possible to perform action recognition on skeleton data instead of RGB images, paving the way for light-weight action recognition algorithms that are robust to different background and light-ing conditions [8,22,27,32,37,40,41,57]. However, training these algorithms in a fully-supervised manner requires large datasets of 3D skeleton data with accurate annotations, which are time-consuming and costly to prepare.

To address this issue, self-supervised methods have been proposed to learn action representations from skeleton data without human-provided labels. The learned representations can subsequently be fine-tuned on a smaller set of labeled data to obtain an action recognition model. Early methods have been focusing on learning representations by solving reconstruction or prediction problems [19, 46, 61], while more recent works adopt a contrastive learning framework [26, 29, 38].

Contrastive learning has been used recently to reach remarkable performance in the image domain and is closing the performance gap with supervised pre-training when transferred to downstream tasks [4-6, 15, 34, 50]. These methods learn representations by mapping different augmented views of the same input sample (positive pairs) closer while pushing augmented views of different input samples (negative pairs) apart. Contrastive methods require a large number of negative samples to achieve good performance, necessitating large batch sizes [4, 5] or memory banks of negative samples [6, 15]. More recently, Bootstrap Your Own Latent (BYOL) [13] showed that explicit negative pairs are not required to learn transferable visual representations. Instead, it uses two networks - an online network and a target network - to encode two augmented views of the same input sample. The online network is trained to predict the output of the target network, and the target network is updated with an exponential moving average of the online network.

In this work, we adapt BYOL for skeleton-based action recognition. As is the case with contrastive learning, data augmentation is an essential part of this method as it guides the network to learn relevant features in the absence of labels. We propose new data augmentation approaches for action sequences tailored to make the learned representations robust to semantically-irrelevant variations. We introduce two asymmetric augmentation pipelines to reduce the distribution shift between self-supervised pre-training and supervised fine-tuning. We also propose a multi-viewpoint sampling method to leverage readily available positive pairs to produce distinct views of the same sample. Several datasets feature the same action sequence captured simultaneously from different angles by different cameras. A naive adaptation of BYOL would treat these recordings of the same action sequence as unrelated samples. Under multi-viewpoint sampling, these are treated as positive pairs, which encourages learning view-invariant representations.

We evaluate the learned representations on linear evalu-



Figure 1. **Overview of our proposed method.** We sample two random viewpoints \mathbf{s}_{i,c_1} , \mathbf{s}_{i,c_2} of the same action sequence captured from different viewing angles and transform them with two asymmetric augmentation pipelines. The resulting views are passed to the online and target networks. The online network is trained to predict the output of the target network. The gradients are only propagated through the online network, while the target network's parameters are updated with an exponential moving average of the online network's parameters.

ation, semi-supervised learning and transfer learning tasks. We show that representations learned with our method outperform the current state of the art on all benchmarks. Inspired by [5], we achieve further performance improvement at the semi-supervised tasks using knowledge distillation from wider networks.

To summarize, the main contributions of this paper are as follows:

- We adapt BYOL to learn representations for skeletonbased action recognition.
- We propose a data augmentation strategy for skeleton sequences including asymmetric transformation pipelines and multi-viewpoint sampling.
- We conduct extensive experiments on three large-scale datasets for skeleton-based action recognition, NTU-60, NTU-120 and PKU-MMD, to demonstrate the performance of our proposed method. Our method consistently outperforms the current state of the art on linear evaluation, semi-supervised and transfer learning benchmarks.

2. Related Work

Skeleton-based Action recognition. Early skeleton-based action recognition methods relied on hand-crafted features [52,54,55]. RNN-based methods treat skeleton data as vector sequences [11,39,44,58], while CNN approaches transform the skeleton data into image-like formats amenable to convolutions [17, 18, 21, 23, 24, 31].

Yan et al. [57] introduced a graph structure called Spatial-Temporal Graph Convolutional Network (ST-GCN), in which spatial edges connect body joints and temporal edges connect each joint across time steps. Many recent methods adopt the ST-GCN architecture [8, 22, 27, 32, 37,40,41,45] and enhance it with new features such as adaptive adjacency, self-attention, and ensembling of multiple input streams.

Self-supervised learning. A common approach to representation learning is to train a network on a large, labeled dataset. Self-supervised learning aims at learning transferable representations without using labels by extracting supervisory signals from the data itself. Early methods relied on pretext tasks related to high-level image understanding, including reconstruction tasks such as denoising [53], inpainting [36] or colorizing [20, 59, 60], or transformation prediction tasks such as jigsaw puzzles [10, 35] or image rotations [12]. Discriminative methods based on contrastive learning [4–6, 14, 15, 34, 50] have recently achieved new state-of-the-art results on linear evaluation and semisupervised learning tasks. Moreover, when transferring the learned representation to downstream tasks such as object detection, they also demonstrated performance on par with or surpassing supervised pre-training. These methods learn by constructing positive and negative sample pairs and training the network to generate representations that can be used to distinguish the positive pairs from the negative ones. A large number of negative samples is usually needed to achieve the best performance. For example, SimCLR [4,5] uses large batch sizes, and MoCo [6, 15] keeps a large queue of past representations as negative samples.

Several recent works showed that useful representations can be learnt without the need for explicit negative samples [3, 7, 13]. BYOL [13], in particular, trains an *online* network to predict the representation generated by a *target* network given two augmented views of the same input. Our work adapts BYOL for skeleton-based action recognition. While BYOL uses two slightly different augmentation sets, we use two very distinct conservative and aggressive augmentation pipelines, and we leverage recordings of the same action captured simultaneously from different camera angles to encourage learning view-invariant representations.

Self-supervision for Skeleton-based Action Recognition.

Self-supervised methods for learning action representations from skeleton sequences focused initially on solving pretext tasks. Zheng et al. [61] used a recurrent encoder-decoder GAN to learn representations by reconstructing masked sequences. Predict & Cluster [46] trained a sequential autoencoder to predict skeleton motion while keeping the decoder's weights fixed to force the encoder to learn discriminative latent representations. EnGAN [19] used a variational auto-encoder to learn a pose embedding manifold, and an additional encoder-decoder to learn action representations as trajectories on the pose manifold. Later works adopted the contrastive learning paradigm. MS²L [29] integrated contrastive learning in a multi-task learning framework, and AS-CAL [38] used different augmentations of skeleton sequences to produce positive and negative pairs. Thoker et al. [48] performed representation learning with two different network architectures operating on graphbased and sequence-based modalities, in a cross-contrastive manner. Recently, SkeletonCLR [26] learned skeleton sequence representations with a momentum contrast framework, while CrosSCLR [26] trained three networks using different modalities (joints, bones, motion) and mined similar sequences between modalities. In a concurrent work, AimCLR [49] builds on SkeletonCLR, adding an energybased attention-guided drop module and nearest neighbours mining. They also use normal and extreme augmentation pipelines, and propose to minimize the distributional divergence between the normally augmented and the extremely augmented views.

3. Method

In this section, we first review BYOL, a state-of-theart self-supervised representation learning algorithm upon which our method is based. We then describe our data augmentation and multi-viewpoint sampling strategies. Finally, we describe the steps for knowledge distillation.

3.1. Bootstrap Your Own Latent (BYOL)

As shown in Fig. 1, BYOL consists of two neural networks. The *target* network, defined by weights ξ , comprises an encoder f_{ξ} and a projector g_{ξ} . The *online* network is defined by parameters θ and consists of an encoder f_{θ} , a projector g_{θ} and a predictor q_{θ} . BYOL performs representation learning by maximizing the similarity between the embeddings of two augmentations of the same action sequence generated by these networks.

More specifically, given an action sequence s, two random augmentations $\mathbf{v_1} = t_1(\mathbf{s})$ and $\mathbf{v_2} = t_2(\mathbf{s})$ are generated using $t_1 \sim \mathcal{T}_1, t_2 \sim \mathcal{T}_2$, where \mathcal{T}_1 and \mathcal{T}_2 denote two sets of transformations. The view $\mathbf{v_1}$ is fed to the online network, where the encoder produces the representation $\mathbf{y_1} = f_{\theta}(\mathbf{v_1})$. This is then projected to an embedding $\mathbf{z_1} = g_{\theta}(\mathbf{y_1})$, which is fed to the predictor, yielding $q_{\theta}(\mathbf{z_1})$. Likewise, the view $\mathbf{v_2}$ is fed to the target network, which produces the representation $\mathbf{y_2} = f_{\xi}(\mathbf{v_2})$ and projection $\mathbf{z_2} = g_{\xi}(\mathbf{y_2})$. Given $\mathbf{z_1}$, the prediction network q_{θ} is trained to predict the target $\mathbf{z_2}$ by minimizing the normalized ℓ_2 distance

$$\mathcal{L}_{\theta,\xi}(\mathbf{v_1}, \mathbf{v_2}) = \left\| \overline{q_{\theta}(\mathbf{z_1})} - \overline{\mathbf{z_2}} \right\|_2^2, \qquad (1)$$

where $\overline{\mathbf{u}} = \mathbf{u}/\|\mathbf{u}\|_2$. The loss is symmetrized by also inputting \mathbf{v}_2 to the online network and \mathbf{v}_1 to the target network and calculating $\mathcal{L}_{\theta,\xi}(\mathbf{v}_2, \mathbf{v}_1)$. This gives the symmetric loss

$$\mathcal{L}_{\theta,\xi}^{\text{BYOL}} = \mathcal{L}_{\theta,\xi}(\mathbf{v_1}, \mathbf{v_2}) + \mathcal{L}_{\theta,\xi}(\mathbf{v_2}, \mathbf{v_1}).$$
(2)

At each iteration, we optimize the loss in Eq. (2) only with regard to the online weights θ , while the target network's weights ξ are kept fixed. The target network is updated from past iterations of the online network, using an exponential moving average (EMA) on the weights θ after each training step. i.e.

$$\xi \leftarrow \lambda \xi + (1 - \lambda)\theta, \ \lambda \in [0, 1].$$
(3)

Momentum BatchNorm. BYOL performs best when calculating stable batch statistics over large mini-batches across GPUs and compute nodes. This synchronized batch normalization (SyncBN) relies on cross-GPU communication and leads to slower training speeds. As the target network does not propagate gradients, Cai *et al.* [1] and Li, Liu and Sun [28] showed that the need for SyncBN can be eliminated by employing momentum BN layers in the target network, which keep a moving average of the batch statistics. After each training iteration, the mean μ and variance σ of every momentum BN layer in the target network is updated using an exponential moving average of the current batch statistics of the online network, i.e.

$$\mu \leftarrow \alpha \mu + (1 - \alpha)\mu_b, \quad \sigma \leftarrow \alpha \sigma + (1 - \alpha)\sigma_b, \quad (4)$$

where $\alpha \in [0, 1]$ is the momentum coefficient, μ_b and σ_b are the current batch statistics. We find that using momentum BN instead of SyncBN reduces training times by 30%.

3.2. Data Augmentation

As the encoder learns to map augmented views of the same skeleton sequence close to each other in latent space, it is trained to ignore the variances induced by the augmentations. The choice of data augmentation is thus critical for learning representations that are semantically relevant for action recognition. We use the following augmentation strategies: **Shear.** Shearing [38] is a linear transformation that slants the shape of an object in a given direction. The shear transformation matrix is defined as

$$\boldsymbol{S} = \begin{bmatrix} 1 & s_X^Y & s_X^Z \\ s_Y^X & 1 & s_Y^Z \\ s_Z^X & s_Z^Y & 1 \end{bmatrix},$$
(5)

where $s_X^Y, s_X^Z, s_Y^X, s_Y^Z, s_Z^X, s_Z^Y$ are the shear factors defining the displacement direction. We apply the transformation by multiplying the skeleton sequence with S on the channel dimension.

Left/Right Drop. This transformation randomly zeroes out the left or right limbs of the skeleton during the entire sequence so the network learns to extract as much information as possible from partial observations of the body.

Resampling. Leveraging the fact that the same action can be performed at various speeds, we speed up or slow down the input sequence by resampling it along the temporal dimension according to a random rate.

Filtering. Skeleton sequences captured by Kinect cameras are quite noisy, which is apparent in the NTU-60 and NTU-120 datasets. We train the network to be robust to noise by randomly smoothing the sequences with a low-pass filter.

Temporal shift. In real-world applications, action recognition algorithms are typically applied on overlapping temporal windows, which means that the action of interest might have already started, or might start or end anytime during a window. We therefore shift the sequence cyclically along the time dimension with a random offset.

3.3. Asymmetric augmentation pipelines

While strong augmentation has been shown to be beneficial for self-supervised learning [13], some of the augmentations above heavily distort the augmented samples, which may become very different from the original data. By contrast, overly aggressive augmentations empirically lead to worse results during supervised fine-tuning, therefore only conservative augmentations are used in that setting. To enforce consistency between the representations of the data seen during self-supervised pre-training and in the fine-tuning stage, we use two very different transformation distributions during self-supervised learning.

In the *aggressive* distribution \mathcal{T}_1 , filtering, random resampling, shear, temporal shift and left/right drop are available. Filtering and left/right drop are randomly activated, and each side of the body has equal probability of being dropped. In the *conservative* distribution \mathcal{T}_2 , only random resampling is enabled. The augmentation pipelines are illustrated in Fig. 2.

Our experiments show that networks pre-trained with the asymmetric pipelines achieve better performance on down-stream tasks.



Figure 2. Asymmetric Augmentation Pipelines.

3.4. Multi-Viewpoint Sampling

It is common for skeleton sequence datasets to contain recordings of the same action sequences captured simultaneously from different angles by different cameras. Although the 3D skeletal representation provided by RGB-D cameras such as Kinect is in essence view-invariant, it is still sensitive to errors of the body tracker, due *e.g.* to selfocclusions [39]. Unlike action labels, viewpoint information is free to obtain as it is provided by the capturing system. We propose to leverage this information to learn representations that are invariant to changes of viewpoint and to the errors of the pose algorithm.

Let S be the dataset. Each skeleton sequence $\mathbf{s}_{i,c} \in S$ represents a viewpoint of the action sequence \mathbf{s}_i captured by camera $c \in \{1..C\}$. Under multi-viewpoint sampling, instead of iterating over the skeleton sequences $\mathbf{s}_{i,c}$, we iterate over the action sequences $\mathbf{s}_i = \{\mathbf{s}_{i,1}, ..., \mathbf{s}_{i,C}\}$. When \mathbf{s}_i contains more than one viewpoint, we draw two random viewpoints \mathbf{s}_{i,c_1} , \mathbf{s}_{i,c_2} from \mathbf{s}_i , with replacement (i.e., the two viewpoints may be identical), to construct a sequence pair. The two viewpoints are passed to the transformations $t_1 \sim \mathcal{T}_1, t_2 \sim \mathcal{T}_2$ to form $\mathbf{v}_1 = t_1(\mathbf{s}_{i,c_1})$ and $\mathbf{v}_2 = t_2(\mathbf{s}_{i,c_2})$, respectively. We find that multi-viewpoint sampling significantly improves the accuracy of networks trained with our method.

3.5. Knowledge Distillation

For self-supervised learning, it is beneficial to use larger, wider networks. This, however, increases the cost for inference. Inspired by [5], we use knowledge distillation [16] to transfer the knowledge from a large model to a smaller model to address this problem. Given a fixed *teacher network* (fine-tuned on a fraction of the labels), we train a randomly-initialized *student network* to match the output predictions of the teacher by minimizing the loss function

$$L = H_{CE}\left(\sigma\left(\frac{z_t(x)}{\tau}\right), \sigma\left(\frac{z_s(x)}{\tau}\right)\right), \qquad (6)$$

where $H_{CE}(.)$ is the cross-entropy loss, $\sigma(.)$ is the softmax function, $z_s(x)$ and $z_t(x)$ are the output logits of the student and teacher networks respectively, and τ is a temperature parameter controlling the softness of the teacher's output. Note that this step does not require any labeled data.

4. Experiments

We follow standard practice and evaluate our representations under the linear evaluation and semi-supervised protocols. We also demonstrate the performance of our method on a transfer learning task. We perform an ablation study to illustrate the effectiveness of each component in our method, and conduct comprehensive comparisons with existing state-of-the-art methods.

4.1. Datasets

NTU-60. NTU-60 [39] is a large-scale action-recognition dataset containing over 56K samples of 60 different actions performed by 40 subjects, recorded simultaneously from different viewpoints by Kinect v2 sensors. We follow the splits recommended by the authors: in Cross-Subject (CS), the training and test sets each contains data from 20 different subjects, and in Cross-View (CV) the training set contains action sequences captured by cameras 2 and 3 while the test set contains the same sequences captured by camera 1.

NTU-120. NTU-120 [30] is an extension of NTU-60 and contains 60 additional action classes. It contains 114K samples of 120 actions performed by 106 different subjects. The dataset contains 32 setups denoting specific locations and backgrounds. We follow the standard splits defined by the authors: for Cross-Subject (CSub), the training and test sets contain data from 53 subjects each, and for Cross-Setup (CSet) the training set contains data from all samples with even setup IDs, and the test set those with odd setup IDs.

PKU-MMD II. PKU-MMD II [9] is a large-scale multimodal dataset for continuous 3D human action detection, consisting of 2000 short sequences containing 7K instances of 41 action classes performed by 13 subjects and recorded from three viewpoints by Kinect v2 cameras. We follow the Cross-Subject split provided by the authors.

4.2. Implementation details

Encoder network and data representation. We use the original ST-GCN [57], both as the encoder in our framework and as a strong supervised baseline for our experiments. In experiments with encoders of different widths, we multiply the number of channels in all layers. For example, a $2\times$ network has twice the number of channels in each layer compared with the standard ST-GCN, whereas a $0.5\times$ network has half as many.

In addition to the 3D joint coordinates, we use the bone representation, calculated as the vectors between the two vertices of each edge of the skeleton graph [41]. Several recent action recognition methods use multiple modalities, including joint coordinates and bone representation, either by ensembling different models or by employing multi-stream architectures [41,45]. Here, we simply concatenate the joint and bone representations along the channel dimension. This only affects the input layer of the network and does not require any further modification to the ST-GCN architecture.

Data processing. All skeleton sequences are cropped to 150 frames, and sequences shorter than 150 frames are repeated. The skeletons are centered, and rotated so that the body has a canonical orientation in the first frame.

Data augmentation. In the *aggressive* augmentation pipeline, we resample the skeleton sequences by a rate uniformly sampled in the interval [0.7, 1.3]. Filtering is applied randomly with a probability of 0.5, and is performed using a Savitzky-Golay filter with order 2 and window 15. The shear factors are sampled randomly from [-1, 1]. The sequence is cyclically shifted along the temporal dimension with an offset uniformly sampled in the interval [0, 150]. The left/right drop augmentation is applied with a probability of 0.5, and each side of the body has equal probability of being dropped.

In the *conservative* augmentation pipeline, only random resampling is performed.

Supervised baseline training. We train the fully supervised model for 80 epochs with a mini-batch size of 8 on a single GPU. We use SGD with weight decay 0.0003 and momentum 0.9. The initial learning rate is 0.0125, and is multiplied by 0.1 at epochs 56 and 70.

Self-supervised pre-training. During self-supervised pretraining, we train the model for 1600 epochs with a minibatch size of 512 on 4 V100 GPUs, which takes approximately 16 hours. For the much smaller PKU-MMD II dataset, we train for 4800 epochs. For wider $(2\times)$ networks, pre-training takes 23 hours on 8 V100 GPUs. We use SGD with weight decay 10^{-4} and momentum 0.9. The learning rate starts at 10^{-6} and is linearly increased to 0.2 in the first

	NTU	J-60	NTU	-120	PKU-MMD II
Method	CS	CV	CSub	CSet	CS
ST-GCN (sup.)	88.5	94.3	83.0	85.1	61.9
LongT GAN [61]	39.1	48.1	-	-	26.5
$MS^{2}L$ [29]	52.6	-	-	-	27.6
PCRP [56]	53.9	63.5	-	-	-
AS-CAL [38]	58.5	64.8	48.6	49.2	-
Thoker <i>et al</i> . [48]	76.3	85.2	67.1	67.9	36.0
3s-CrosSCLR [26]	77.8	83.4	67.9	66.7	-
3s-AimCLR [49]	78.9	83.8	68.2	68.8	-
Ours	86.8	91.2	77.1	79.2	55.25

Table 1. Linear evaluation protocol. Performance comparison (Top-1 accuracy, %) with the fully-supervised baseline and other self-supervised methods on NTU-60, NTU-120 and PKU-MMD.

10 epochs of training, after which it is decreased to 0 following a cosine decay schedule [33]. We initially set the BYOL momentum λ to 0.99 and increase it to 1 following a cosine schedule during training.

4.3. Linear Classification Protocol

We first evaluate the quality of the representations learned with our method under the linear classification protocol, training a linear classifier on top of the frozen pretrained online encoder f_{θ} . We train for 100 epochs with a batch size of 256 on 4 V100 GPUs, using SGD with momentum 0.9 and without weight decay. The initial learning rate is 30, and is multiplied by 0.1 at epochs 60 and 80. We only use conservative augmentations in this step.

Comparison with previous results We compare our method with the fully-supervised baseline and with other self-supervised methods by linear evaluation on the NTU-60, NTU-120 and PKU-MMD II datasets. As shown in Tab. 1, our method outperforms all other self-supervised methods and considerably reduces the gap to fully-supervised training.

Ablation Studies

We evaluate the usefulness of the components of our method by removing them in turn. Note that removing both multi-viewpoint sampling and the asymmetric augmentation pipelines corresponds to applying the classical BYOL method to skeleton sequence data.

The results are shown in Tab. 2. They demonstrate the performance of the aggressive augmentation pipeline, and show that all data augmentations contribute to the method's performance, with shearing being the most effective transformation. The results also demonstrate the performance gains obtained with multi-viewpoint sampling and with the



Figure 3. **t-SNE visualization of the features** learned under self-supervised training with a) Baseline BYOL with aggressive augmentations, b) Our method without multi-viewpoint learning, c) Our method with only aggressive augmentations and d) Our method. Each point represents a skeleton sequence. We show 10 random action classes of the NTU-60 dataset, indicated by colors.

use of two asymmetric augmentation pipelines. In the crossview setting, the shear transformation and multi-viewpoint sampling both provide the most significant performance increase, reinforcing our claim that these features force the network to learn invariances to change of viewpoint. Finally, for both benchmarks, the combination of all our contributions provides the best results.

Visualization of Learned Representations We present in Fig. 3 a t-SNE [51] visualization of the representations learned under self-supervised training with the BYOL baseline and with our proposed method, with and without multiviewpoint learning and asymmetric augmentation pipelines. Our proposed multi-viewpoint learning and asymmetric augmentation pipelines lead to a better separation of the action classes.

4.4. Semi-Supervised Evaluation

We now evaluate our learned representation on semisupervised learning tasks. We sample 1%, 5% and 10% of the NTU-60 training set in a class-balanced way as labeled samples and fine-tune the encoder $f_{\theta}(.)$ with a linear layer on top. We report top-1 accuracy on the entire test set. We present results averaged over five randomly sampled subsets of the labeled data.

Method	MVS	Asymm. Aug.	Augmentations	NTU-60 (CS)	NTU-60 (CV)
Baseline BYOL			Only conservative augmentation	51.1	55.1
			Only aggressive augmentation	79.6	81.0
Ours	\checkmark	\checkmark	No shear	78.0	81.1
	\checkmark	\checkmark	No low-pass filtering	84.4	90.2
	\checkmark	\checkmark	No temporal shift	84.7	89.4
	\checkmark	\checkmark	No resampling	85.6	90.6
	\checkmark	\checkmark	No left/right drop	85.6	89.7
		\checkmark	All augmentations	81.8	82.3
	\checkmark		Only aggressive augmentation	83.8	88.7
	\checkmark	\checkmark	All augmentations	86.8	91.2

Table 2. **Ablation study** on the components of our self-supervised training method. We remove multi-viewpoint sampling (MVS), the asymmetric augmentation pipelines (Asymm. Aug.) and each augmentation in turn. Linear evaluation (Top-1 accuracy, %) on NTU-60.

	NTU-60 (CS)		NT	NTU-60 (CV)		
Label fraction	1%	5%	10%	1%	5%	10%
ST-GCN (sup.)	19.3	59.1	71.7	20.6	61.6	75.9
SESAR-KT [25]	48.1	55.0	58.2	-	-	-
$MS^{2}L$ [29]	33.1	-	65.2	-	-	-
ASSL [43]	-	57.3	64.3	-	63.6	69.8
3s-CrosSCLR [26]	51.1	-	74.4	50.0	-	77.8
Thoker <i>et al</i> . [48]	35.7	59.6	65.9	38.1	65.7	72.5
3s-AimCLR [49]	54.8	-	78.2	54.3	-	81.6
Ours 1×, distilled	79.4	83.6	84.6	81.7	87.5	89.5
Ours $2 \times$	79.3	84.5	86.0	81.5	87.5	89.8

Table 3. **Semi-supervised learning on NTU-60** with different label fractions for fine-tuning. The labels are sampled in a class-balanced way. For our method, we present results averaged over five random seeds for sampling the labeled data.

When fine-tuning the network on 1%, 5% and 10% of the labeled samples, we train for 100, 20 and 10 epochs respectively, with a mini-batch size of 8 on a single Nvidia V100 GPU using SGD with momentum 0.9 and without weight decay. The initial learning rate is 10, and is decayed to 0 using a cosine schedule. We keep the pre-trained encoder frozen during the first third of the iterations. Afterwards, the encoder's learning rate follows the same schedule as the fully connected layer's but is multiplied by a factor of 10^{-3} . Only conservative augmentations are used.

We then distill the fine-tuned model to another network using the whole training set *without labels*. We train the student network for 100 epochs with a batch size of 256, using SGD with weight decay 10^{-4} , momentum 0.9, and distillation temperature 1.0. The initial learning rate is 0.2, and is decayed to 0 with a cosine schedule.

	NTU-120 (CSub)			NTU-120 (CSet)		
Label fraction	1%	5%	10%	1%	5%	10%
ST-GCN (sup.)	15.2	51.2	63.0	14.0	49.9	61.9
Ours $1 \times$, distilled Ours $2 \times$	66.8 65.9	74.2 73.9	77.0 76.7	67.1 66.6	76.2 76.1	78.0 78.6

Table 4. Semi-supervised learning on NTU-120 with different label fractions for fine-tuning.

Comparison with previous results In Tabs. 3 and 4 we compare our method with the fully-supervised baseline and other semi-supervised methods. The results further highlight the quality of the action representations learned by our approach: fine-tuning the model on 1% (~400) labeled sequences of NTU-60 (CS) yields a top-1 accuracy of 79.4%, outperforming both the fully-supervised baseline and other semi-supervised methods by a large margin.

Effect of distillation To evaluate the effect of knowledge distillation, we compare semi-supervised and distilled ST-GCN models of varied widths on the NTU-60 (CS) dataset. The semi-supervised models are pre-trained on the whole unlabeled training set and fine-tuned on different label fractions of the training set, while the distilled models are trained using fine-tuned $2\times$ models as teachers. The results are shown in Tab. 5. We observe that knowledge distillation is effective and consistently provides generalization gains, especially at lower label fractions, and for lower-capacity students.

4.5. Transfer Learning

We evaluate the transfer learning performance of our method by pre-training the encoder $f_{\theta}(.)$ on NTU-60 and fine-tuning it with a linear layer on top on the PKU-MMD

]	Label fractio	n
Model	KD	1%	5%	10%
Ours $0.25 \times$	\checkmark	$53.5 \pm 1.5 \\ 76.8 \pm 0.7$	$\begin{array}{c} 68.9 \pm 0.4 \\ 80.4 \pm 1.4 \end{array}$	$71.2 \pm 0.3 \\ 81.3 \pm 0.9$
Ours $0.5 \times$	\checkmark	$\begin{array}{c} 66.9 \ \pm 0.8 \\ 78.3 \ \pm 0.3 \end{array}$	$\begin{array}{c} 78.4 \ \pm 0.6 \\ 81.7 \ \pm 0.7 \end{array}$	$\begin{array}{c} 80.4 \pm 0.3 \\ 82.8 \pm 0.4 \end{array}$
Ours 1×	\checkmark	$\begin{array}{c} 75.8 \pm 1.0 \\ 79.4 \pm 0.6 \end{array}$	$\begin{array}{c} 82.8 \pm 0.6 \\ 83.6 \pm 0.7 \end{array}$	$\begin{array}{c} 84.5 \pm 0.4 \\ 84.6 \pm 0.3 \end{array}$
Ours $2 \times$		79.3 ± 0.6	84.5 ± 0.4	86.0 ± 0.1

Table 5. Effect of distillation. Top-1 accuracy (%) for semisupervised and distilled models of varied sizes on different label fractions of the NTU-60 (CS) dataset. Knowledge distillation (KD) is performed with the whole *unlabeled* training set, using the $2 \times$ models fine-tuned on each label fraction as teachers. We present the mean and standard deviation of the results obtained over five random seeds for sampling the labeled data.

Method	PKU-MMD II
ST-GCN (supervised)	61.9
LongT GAN [61]	44.8
$MS^{2}L$ [29]	45.8
Thoker <i>et al</i> . [48]	45.9
Ours	68.3

Table 6. **Transfer learning from NTU-60 to PKU-MMD II.** Performance comparison (Top-1 accuracy, %) when fine-tuning on PKU-MMD II a model pre-trained on NTU-60.

II dataset. The results are shown in Tab. 6. When finetuned, our pre-trained model significantly outperforms both previous work and the supervised baseline.

5. Discussion

5.1. Limitations

Not all datasets provide multiple camera views of the same action sequence, making it impossible to take advantage of our proposed multi-viewpoint sampling. However, as seen in Sec. 4.3, our method still provides competitive results without this contribution.

The benchmarks used by this and previous work are based on datasets in which the classes are balanced and the class distributions of the training and test sets are similar. The results may not be representative of the performance obtained in real applications, where unlabeled datasets are often imbalanced, with long-tailed label distributions. Investigating the robustness of self-supervised skeleton-based action representation learning to dataset imbalance requires further exploration, which we leave to future work.

5.2. Potential Negative Societal Impact

In contrast to fully-supervised learning, self-supervised representation learning requires training with large batches for a large number of epochs on power-intensive computing hardware, and has therefore a higher environmental impact. However, as shown in Sec. 4.4, knowledge distillation allows transferring the learned representations to smaller, more energy-efficient models with little accuracy loss, which may help offset this impact.

The motivations behind our work have been to decrease the cost and increase the ease of adoption of action recognition technologies. Camera-based applications becoming pervasive in public spaces raises obvious privacy concerns. Persons moving in monitored areas might not be aware of the presence of cameras, and may not have the possibility to express (or deny) consent. In a long-term perspective, if the technology generalizes the omnipresence of cameras might affect individuals' sense of freedom, and could generate anxiety. Alleviating such concerns will require a considerable effort of transparency and pedagogy to inform about the presence and purpose of capturing devices, proactive security and data protection, and clear accountability.

Performing self-supervised representation learning on large unlabeled datasets has the potential to reduce human bias introduced when curating and annotating data. However, some groups that are underrepresented in society will be likewise underrepresented in the unlabeled data. Ensuring that models used to make decisions do not impact those groups unfairly will require sizable effort.

Finally, the applications we envisage are overwhelmingly positive, in domains such as healthcare, sports, entertainment or livestock farming. However, we are conscious that the same technology may be used for other purposes.

6. Conclusion

We presented a simple framework for self-supervised representation learning for skeleton-based action recognition based on BYOL. It includes a data augmentation strategy for skeleton data based on two distinct transformation pipelines, and a multi-viewpoint learning method that makes better use of action sequences that are captured simultaneously by different cameras. Experiments on three large scale action recognition datasets show that our method consistently outperforms the current state of the art on linear evaluation, semi-supervised and transfer learning tasks.

Acknowledgment

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

References

- Zhaowei Cai, Avinash Ravichandran, Subhransu Maji, Charless Fowlkes, Zhuowen Tu, and Stefano Soatto. Exponential moving average normalization for self-supervised and semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 194–203, 2021. 3
- [2] Zhe Cao, Gines Hidalgo Martinez, Tomas Simon, Shih-En Wei, and Yaser A Sheikh. OpenPose: Realtime Multi-Person 2D pose estimation using part affinity fields. *IEEE Trans. Pattern Anal. Mach. Intell.*, July 2019. 1
- [3] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in Self-Supervised vision transformers. In Proceedings of the International Conference on Computer Vision (ICCV), 2021. 2
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. Feb. 2020. 1, 2
- [5] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big Self-Supervised models are strong Semi-Supervised learners. June 2020. 1, 2, 4
- [6] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. Mar. 2020. 1, 2
- [7] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. Nov. 2020. 2
- [8] Ke Cheng, Yifan Zhang, Xiangyu He, Weihan Chen, Jian Cheng, and Hanqing Lu. Skeleton-Based action recognition with shift graph convolutional network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 183–192, 2020. 1, 2
- [9] Liu Chunhui, Hu Yueyu, Li Yanghao, Song Sijie, and Liu Jiaying. Pku-mmd: A large scale benchmark for continuous multi-modal human action understanding. *arXiv preprint arXiv:1703.07475*, 2017. 5
- [10] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015. 2
- [11] Yong Du, Wei Wang, and Liang Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1110–1118, 2015. 2
- [12] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. Mar. 2018. 2
- [13] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to Self-Supervised learning. In H Larochelle, M Ranzato, R Hadsell, M F Balcan, and H Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21271–21284. Curran Associates, Inc., 2020. 1, 2, 4

- [14] R Hadsell, S Chopra, and Y LeCun. Dimensionality reduction by learning an invariant mapping. In 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06), volume 2, pages 1735–1742, June 2006. 2
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, June 2020. 1, 2
- [16] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In NIPS Deep Learning and Representation Learning Workshop, 2015. 4
- [17] Qiuhong Ke, Mohammed Bennamoun, Senjian An, Ferdous Sohel, and Farid Boussaid. A new representation of skeleton sequences for 3d action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3288–3297, 2017. 2
- [18] Tae Soo Kim and Austin Reiter. Interpretable 3D human action analysis with temporal convolutional networks, 2017.
- [19] Jogendra Nath Kundu, Maharshi Gor, Phani Krishna Uppala, and R Venkatesh Babu. Unsupervised feature learning of human actions as trajectories in pose embedding manifold. Dec. 2018. 1, 3
- [20] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Learning representations for automatic colorization. In *Computer Vision – ECCV 2016*, pages 577–593. Springer International Publishing, 2016. 2
- [21] Bo Li, Yuchao Dai, Xuelian Cheng, Huahui Chen, Yi Lin, and Mingyi He. Skeleton based action recognition using translation-scale invariant image mapping and multi-scale deep CNN. In 2017 IEEE International Conference on Multimedia Expo Workshops (ICMEW), pages 601–604, July 2017. 2
- [22] Bin Li, Xi Li, Zhongfei Zhang, and Fei Wu. Spatio-Temporal graph routing for Skeleton-Based action recognition. AAAI, 33(01):8561–8568, July 2019. 1, 2
- [23] Chao Li, Qiaoyong Zhong, Di Xie, and Shiliang Pu. Skeleton-based action recognition with convolutional neural networks, 2017. 2
- [24] Chao Li, Qiaoyong Zhong, Di Xie, and Shiliang Pu. Cooccurrence feature learning from skeleton data for action recognition and detection with hierarchical aggregation. Apr. 2018. 2
- [25] Jingyuan Li and Eli Shlizerman. Sparse Semi-Supervised action recognition with active learning. Dec. 2020. 7
- [26] Linguo Li, Minsi Wang, Bingbing Ni, Hang Wang, Jiancheng Yang, and Wenjun Zhang. 3D human action representation learning via Cross-View consistency pursuit. Apr. 2021. 1, 3, 6, 7
- [27] Maosen Li, Siheng Chen, Xu Chen, Ya Zhang, Yanfeng Wang, and Qi Tian. Actional-structural graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3595–3603, 2019. 1, 2

- [28] Zeming Li, Songtao Liu, and Jian Sun. Momentum² teacher: Momentum teacher with momentum statistics for self-supervised learning. Jan. 2021. 3
- [29] Lilang Lin, Sijie Song, Wenhan Yang, and Jiaying Liu. MS2L: Multi-Task Self-Supervised learning for skeleton based action recognition. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2490–2498. Association for Computing Machinery, New York, NY, USA, Oct. 2020. 1, 3, 6, 7, 8
- [30] Jun Liu, Amir Shahroudy, Mauricio Lisboa Perez, Gang Wang, Ling-Yu Duan, and Alex Kot Chichung. NTU RGB+D 120: A Large-Scale benchmark for 3D human activity understanding. *IEEE Trans. Pattern Anal. Mach. Intell.*, May 2019. 5
- [31] Mengyuan Liu, Hong Liu, and Chen Chen. Enhanced skeleton visualization for view invariant human action recognition. *Pattern Recognit.*, 68:346–362, Aug. 2017. 2
- [32] Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. Disentangling and unifying graph convolutions for skeleton-based action recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 143–152, 2020. 1, 2
- [33] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. Aug. 2016. 6
- [34] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, June 2020. 1, 2
- [35] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles, 2016. 2
- [36] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016. 2
- [37] Wei Peng, Xiaopeng Hong, Haoyu Chen, and Guoying Zhao. Learning graph convolutional network for Skeleton-Based human action recognition by neural searching. AAAI, 34(03):2669–2676, Apr. 2020. 1, 2
- [38] Haocong Rao, Shihao Xu, Xiping Hu, Jun Cheng, and Bin Hu. Augmented skeleton based contrastive action learning with momentum LSTM for unsupervised action recognition. *Inf. Sci.*, 569:90–109, 2021. 1, 3, 4, 6
- [39] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. NTU RGB+D: A large scale dataset for 3D human activity analysis. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, June 2016. 2, 4, 5
- [40] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Skeleton-based action recognition with directed graph neural networks. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 7912– 7921, 2019. 1, 2
- [41] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Twostream adaptive graph convolutional networks for skeletonbased action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12026–12035, 2019. 1, 2, 5

- [42] Jamie Shotton, Andrew Fitzgibbon, Mat Cook, Toby Sharp, Mark Finocchio, Richard Moore, Alex Kipman, and Andrew Blake. Real-time human pose recognition in parts from single depth images, 2011. 1
- [43] Chenyang Si, Xuecheng Nie, Wei Wang, Liang Wang, Tieniu Tan, and Jiashi Feng. Adversarial self-supervised learning for semi-supervised 3D action recognition. In *Computer Vision – ECCV 2020*, pages 35–51. Springer International Publishing, 2020. 7
- [44] Sijie Song, Cuiling Lan, Junliang Xing, Wenjun Zeng, and Jiaying Liu. An End-to-End Spatio-Temporal attention model for human action recognition from skeleton data. *AAAI*, 31(1), Feb. 2017. 2
- [45] Yi-Fan Song, Zhang Zhang, Caifeng Shan, and Liang Wang. Stronger, faster and more explainable: A graph convolutional baseline for skeleton-based action recognition, 2020. 2, 5
- [46] 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. 1, 3
- [47] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5693– 5703, 2019. 1
- [48] Fida Mohammad Thoker, Hazel Doughty, and Cees G M Snoek. Skeleton-Contrastive 3D action representation learning. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1655–1663. Association for Computing Machinery, New York, NY, USA, Oct. 2021. 3, 6, 7, 8
- [49] Guo Tianyu, Liu Hong, Chen Zhan, Liu Mengyuan, Wang Tao, and Ding Runwei. Contrastive learning from extremely augmented skeleton sequences for self-supervised action recognition. In AAAI, 2022. 3, 6, 7
- [50] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. July 2018.1, 2
- [51] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. 6
- [52] Raviteja Vemulapalli, Felipe Arrate, and Rama Chellappa. 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, pages 588–595, 2014. 2
- [53] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the* 25th international conference on Machine learning, ICML '08, pages 1096–1103, New York, NY, USA, July 2008. Association for Computing Machinery. 2
- [54] Jiang Wang, Zicheng Liu, Ying Wu, and Junsong Yuan. Learning actionlet ensemble for 3D human action recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(5):914– 927, May 2014. 2
- [55] Lu Xia, Chia-Chih Chen, and J K Aggarwal. View invariant human action recognition using histograms of 3D joints. In

2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, pages 20–27, June 2012. 2

- [56] Shihao Xu, Haocong Rao, Xiping Hu, and Bin Hu. Prototypical contrast and reverse prediction: Unsupervised skeleton based action recognition. *ArXiv*, 2020. 6
- [57] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018. 1, 2, 5
- [58] Pengfei Zhang, Cuiling Lan, Junliang Xing, Wenjun Zeng, Jianru Xue, and Nanning Zheng. View adaptive recurrent neural networks for high performance human action recognition from skeleton data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2117–2126, 2017. 2
- [59] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *Computer Vision – ECCV 2016*, pages 649–666. Springer International Publishing, 2016. 2
- [60] Richard Zhang, Phillip Isola, and Alexei A Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1058–1067, 2017. 2
- [61] 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. *AAAI*, 32(1), Apr. 2018. 1, 3, 6, 8