



HAL
open science

View-invariant Skeleton Action Representation Learning via Motion Retargeting

Di Yang, Yaohui Wang, Antitza Dantcheva, Lorenzo Garattoni, Gianpiero
Francesca, François Brémont

► **To cite this version:**

Di Yang, Yaohui Wang, Antitza Dantcheva, Lorenzo Garattoni, Gianpiero Francesca, et al.. View-invariant Skeleton Action Representation Learning via Motion Retargeting. International Journal of Computer Vision, 2024, 10.1007/s11263-023-01967-8 . hal-03906649

HAL Id: hal-03906649

<https://hal.science/hal-03906649>

Submitted on 19 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ViA: View-invariant Skeleton Action Representation Learning via Motion Retargeting

Di Yang^{1,2} Yaohui Wang^{1,2,4*} Antitza Dantcheva^{1,2} Lorenzo Garattoni³
Gianpiero Francesca³ François Brémond^{1,2}

¹Inria ²Université Côte d’Azur ³Toyota Motor Europe ⁴Shanghai AI Laboratory
{di.yang, yaohui.wang, antitza.dantcheva, francois.bremond}@inria.fr
{lorenzo.garattoni, gianpiero.francesca}@toyota-europe.com

Abstract

Current self-supervised approaches for skeleton action representation learning often focus on constrained scenarios, where videos and skeleton data are recorded in laboratory settings. When dealing with estimated skeleton data in *real-world videos*, such methods perform poorly due to the large variations across subjects and camera viewpoints. To address this issue, we introduce ViA, a novel View-Invariant Autoencoder for self-supervised skeleton action representation learning¹. ViA leverages motion retargeting between different human performers as a pretext task, in order to disentangle the latent action-specific ‘Motion’ features on top of the visual representation of a 2D or 3D skeleton sequence. Such ‘Motion’ features are invariant to skeleton geometry and camera view and allow ViA to facilitate both, cross-subject and cross-view action classification tasks. We conduct a study focusing on transfer-learning for skeleton-based action recognition with self-supervised pre-training on real-world data (*e.g.*, Posetics). Our results showcase that skeleton representations learned from ViA are generic enough to improve upon state-of-the-art action classification accuracy, not only on 3D laboratory datasets such as NTU-RGB+D 60 and NTU-RGB+D 120, but also on real-world datasets where only 2D data are accurately estimated, *e.g.*, Toyota Smarthome, UAV-Human and Penn Action.

Introduction

Human action recognition is a crucial task in real-world video understanding. Recent works have made promising progress by adopting spatio-temporal Convolutional Neural Networks (CNNs) (Ji et al. 2013; Carreira and Zisserman 2017; Hara, Kataoka, and Satoh 2017; Feichtenhofer et al. 2019; Feichtenhofer 2020; Ryoo et al. 2020; Li et al. 2021a; Wang et al. 2021) or Transformer (Arnab et al. 2021) to effectively extract features from RGB videos and optical flows (Karen and Andrew 2014; Feichtenhofer, Pinz, and Zisserman 2016). However, such methods are generally poor at recognizing actions performed by subjects or recorded from viewpoints that are not present in the training data. One of the open challenges is to design methods that are invariant to subject and view. As an alternative, skeleton-based action recognition methods have been demonstrated robust

to changes in viewpoints and subjects (Vemulapalli, Arate, and Chellappa 2014; Caetano, Brémond, and Schwartz 2019; Yan, Xiong, and Lin 2018; Li et al. 2019; Shi et al. 2019; Song et al. 2020; Chen et al. 2021; Duan et al. 2022) when the models were trained (and evaluated) on laboratory indoor simulation datasets (Shahrudy et al. 2016; Liu et al. 2020; Wang et al. 2014), where high-quality 3D skeleton data are accessible. Unfortunately, 3D skeleton-based models often fail to generalize to real-world videos (Carreira and Zisserman 2017; Zhang, Zhu, and Derpanis 2013; Das et al. 2019; Li et al. 2021c) containing realistic and diverse conditions, whereas acquiring high-quality 3D skeleton labeled data is extremely expensive. On the contrary, many studies (Weinzaepfel and Rogez 2021; Yang et al. 2021c; Duan et al. 2022) have shown that 2D estimated skeleton, although sensitive to view and subject variations, is more accurate and more effective for action recognition compared to estimated 3D counterparts in many real-world scenarios (Zhang, Zhu, and Derpanis 2013; Das et al. 2019; Li et al. 2021c; Carreira and Zisserman 2017). Based on this observation, we hypothesize that action recognition, particularly based on 2D skeletons, could be improved by embedding a view-invariant representation of skeleton sequences.

In this work we propose ViA, a novel framework to train an embedding model that represents a view-invariant 2D action without the need to explicitly reconstruct 3D skeletons or camera parameters. Given a pair of videos, we denote the temporal static information of a skeleton sequence, *i.e.*, ‘viewpoints’, ‘body size’, etc., as ‘Character’, while the temporal dynamic information, *i.e.*, the specifics of the ‘action’ performed by the subject, as ‘Motion’. ViA is able to make subjects perform each other’s action, while maintaining viewpoint and body size invariance. As the learned ‘Motion’ representation is subject and view agnostic, it can be effectively applied for cross-subject and cross-view action recognition. Specifically, as shown in Fig. 1, ViA incorporates (i) an encoder that operates the skeleton sequences extracted from source and driving videos, along with (ii) a Latent Motion Disentanglement module (LMD) to learn the ‘Motion’ and ‘Character’ components by orthogonal decomposition from the visual skeleton representations. (iii) A lightweight skeleton sequence decoder is employed to generate the novel pair of skeleton sequences from the representations where their ‘Motion’ features are swapped with

^{*}Work done while the author was at Inria

¹Project website: <https://walker-a1y.github.io/ViA-project>

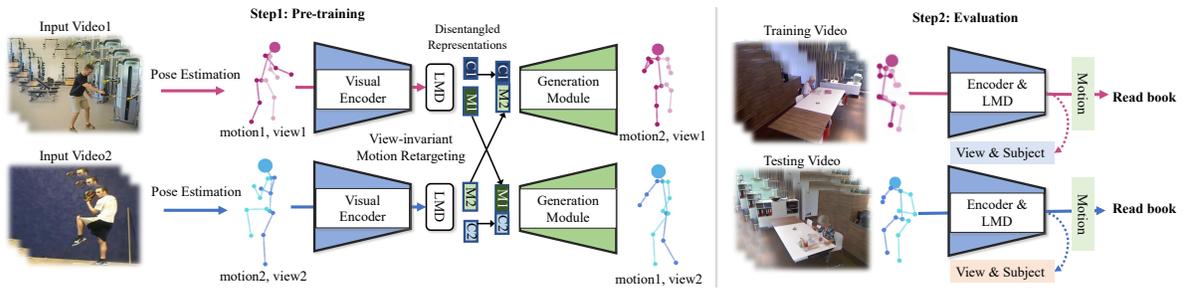


Figure 1: **General pipeline of ViA.** Our framework consists of two steps. Firstly, we apply view-invariant motion retargeting as a pretext task for pre-training a view- and subject-invariant autoencoder for given skeleton sequences. Secondly, the pre-trained visual encoder is evaluated by transferring to downstream cross-view and cross-subject action classification tasks. LMD: Latent Motion Disentanglement module. M: ‘Motion’ features, C: ‘Character’ features.

each other. In addition, towards independence from action annotations, we propose to train the autoencoder of ViA in a self-supervised manner by adopting contrastive loss and a cycle of reconstruction loss.

Towards assessing the performance of our framework, we first pre-train ViA on the large-scale real-world Posetics dataset with a rich variety of subjects and viewpoints and we evaluate the quality of the learned action representation by fine-tuning and linear evaluation protocols on unseen 2D real-world action recognition datasets (*e.g.*, Toyota Smarthome, UVA-Human and Penn Action). As ViA is not limited to 2D skeletons, we additionally validate the effectiveness of ViA on laboratory 3D datasets (*e.g.*, NTU-RGB+D 60 and 120). Experimental analyses confirm that through motion retargeting, ViA outperforms state-of-the-art methods (Sun et al. 2020; Li et al. 2021b; Yang et al. 2021b; Su, Lin, and Wu 2021) on self-supervised action representation learning and the learned video representations can notably transfer to videos with cross-view and cross-subject challenges.

In summary, the contributions of this paper are the following. (i) We introduce ViA, a novel joint generative and discriminative framework. ViA leverages motion retargeting as a pretext task to learn view- and subject-invariant skeleton-based action representations. (ii) We introduce a novel Latent Motion Disentanglement mechanism to decompose and regroup the view-agnostic ‘Motion’ features for skeleton sequences by orthogonal decomposition on visual representations in the latent space. (iii) We set a new state-of-the-art for self-supervised skeleton action recognition on the real-world Posetics dataset. (iv) We conduct a study and show that pre-training ViA on Posetics and transferring it onto an unseen target dataset represents a generic and effective methodology for view- and subject-invariant action classification.

Related Work

Self-supervised Skeleton Action Representation. Self-supervised skeleton representation learning involves learning spatio-temporal features from numerous unlabeled data by means of pretext tasks, *e.g.*, motion consistency learning (Su, Lin, and Wu 2021) and skeleton colorization (Yang et al. 2021d) for specific skeleton sequence processing. These methods highly rely on the quality of pretext tasks. Recently, contrastive learning and its variants (Wu et al.

2018; Hjelm et al. 2019; Bachman, Hjelm, and Buchwalter 2019; Tian, Krishnan, and Isola 2020; He et al. 2020; Chen et al. 2020; Jiao et al. 2020; Yang et al. 2021b) have established themselves as an important direction for self-supervised representation learning, due to their promising performances. Inspired by (Tian, Krishnan, and Isola 2020), CroSCLR (Li et al. 2021b) trains action models (Yan, Xiong, and Lin 2018) in three modalities including joint, bone, and motion by encouraging cross-view consistency. Recent techniques (Su, Lin, and Wu 2021; Yang et al. 2021d; Li et al. 2021b) have only shown promising results on laboratory datasets based on high-quality 3D skeleton sequences that are more robust to view variations. They thus struggle to deal with a large diversity of subjects and viewpoints when generalizing onto real-world action recognition tasks especially on 2D skeleton datasets (Das et al. 2019; Li et al. 2021c; Zhang, Zhu, and Derpanis 2013).

View-invariant Skeleton Representation. To explore the view-invariant representation ability of human skeletons, previous methods (Li et al. 2018; Kundu et al. 2019; Nie, Liu, and Liu 2020; Nie and Liu 2021) aim at disentangling the view-dependent and pose-dependent features by two independent encoders on top of a single 3D skeleton using probabilistic embedding for view-invariant action recognition. To further address inherent ambiguities in 2D skeleton due to 3D-to-2D projection for action recognition, recent methods (Sun et al. 2020; Zhao et al. 2021; Sardari, Ommer, and Mirmehdi 2021) perform the disentanglement learning on specific sensors (*e.g.*, motion capture system) capturing multi-view 2D skeletons. However, The aforementioned methods all process the skeleton sequence frame by frame, which are challenged in capturing the temporal features of the sequence and they are often not available when applying for common 2D datasets (Das et al. 2019; Li et al. 2021c; Zhang, Zhu, and Derpanis 2013) where collecting data in multi-view is expensive and challenging.

In our work, ViA applies generative task for the disentanglement and does not need multi-view data and 3D reconstruction. Moreover, unlike previous works only disentangling static aspects ‘view’ and ‘pose’ for a single frame, the disentanglement of ViA is designed for ‘Character’ (including ‘view’ and ‘pose’) and ‘Motion’ coded in a sequence. The important temporal dimension is considered to better generalize to action understanding tasks. By disentangling ‘Motion’ features using orthogonal decomposition in the la-

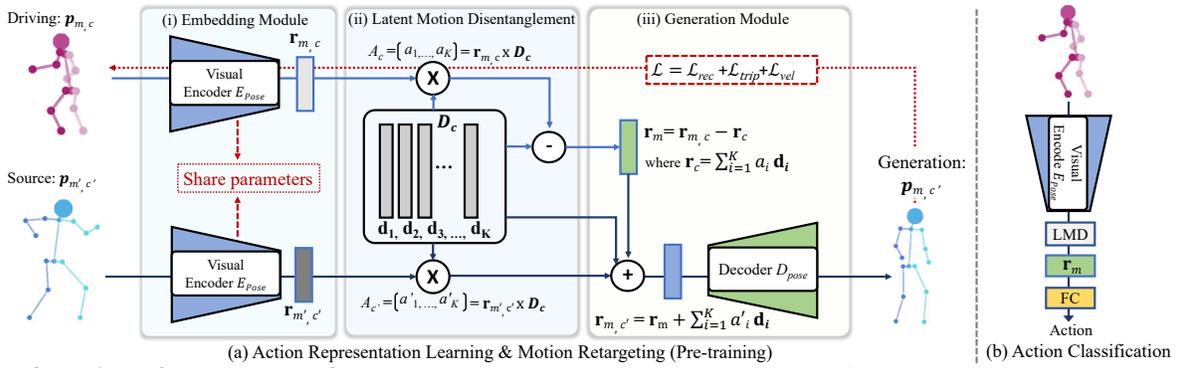


Figure 2: **Overview of the proposed framework.** ViA is an autoencoder consisting of two networks, a visual encoder E_{Pose} and a decoder D_{Pose} . In the latent space, we apply Latent Motion Disentanglement (LMD) by learning a character dictionary D_c , which is an orthogonal basis where each vector represents a basic visual transformation of the ‘Character’ aspects. (a) In the **pre-training** stage, ViA takes two skeleton sequences $\mathbf{p}_{m,c}$ and $\mathbf{p}_{m',c'}$ as driving and source sequence respectively. Firstly, the two samples are encoded into latent codes $\mathbf{r}_{m,c}$ and $\mathbf{r}_{m',c'}$. Then, their projections A_c (or $A_{c'}$) along D_c can be computed and the linear combination of A_c (or $A_{c'}$) with D_c is considered as the ‘Character’ features. The ‘Motion’ features of the driving sequence \mathbf{r}_m can be disentangled and regrouped with the ‘Character’ features of $\mathbf{p}_{m',c'}$ to become the target code $\mathbf{r}_{m,c'}$. Finally, the target skeleton sequence $\mathbf{p}_{m,c'}$ is generated from $\mathbf{r}_{m,c'}$ by D_{Pose} . (b) **In the action classification** stage, we transfer the pre-trained visual Encoder E_{Pose} stacked with LMD onto downstream tasks for fine-tuning or feature extraction.

tent space, ViA eliminates the requirement of explicit regularization terms (Sun et al. 2020; Zhao et al. 2021) that encourage disentanglement and smoothness of the learned representation.

Skeleton Generation by Motion Retargeting. Motion retargeting is one of the important applications of video generation (Tulyakov et al. 2018; Wang et al. 2020; WANG et al. 2020; Yu et al. 2022; Skorokhodov, Tulyakov, and Elhoseiny 2022). Recent methods have used both spatially structured representations (e.g., 2D skeleton sequences (Villegas et al. 2018; Aberman et al. 2019; Chan et al. 2019)) and non-structure representations (e.g., latent codes (Wang et al. 2022)) to transfer motion between humans. Current skeleton generation method (Aberman et al. 2019) shows that transferring motion across characters can enforce view and motion disentanglement. In this work, we propose to leverage such a generative task to improve action recognition downstream tasks. As opposed to finding directions corresponding to individual visual transformations, we seek to learn a set of directions that cooperatively allows for high-level visual transformations that can be beneficial in skeleton generation. Hence, deviating from existing methods (Villegas et al. 2018; Aberman et al. 2019; Chan et al. 2019) using separate networks to learn disentangled features, ViA integrates both, latent motion code, as well as view features in a single encoder, which highly reduces the model complexity and simplifies training. In addition, we employ cycle-consistency (Zhu et al. 2017; Ma et al. 2022) to generalize to real-world videos using non-paired skeleton sequences as the only training data.

Proposed Approach

In this section, we introduce the full architecture and the training strategy of the ViA framework.

Overview. ViA is an autoencoder consisting of a visual encoder and a decoder for skeleton (i.e., pose) sequences

(see Fig. 2). We aim at training the visual encoder that can embed the input skeleton sequence to a view- and subject-invariant representation, then transferring the visual encoder for skeleton-based action classification tasks. To disentangle ‘Motion’ from ‘Character’, we apply motion retargeting as the pretext task (i.e., transferring the motion of driving skeleton sequence to the source skeleton sequence maintaining the source skeletons invariant in viewpoint and body size). We note that input skeletons can be in 3D or 2D only. The framework is learned with a cycle of reconstruction (the autoencoder is reused taking generated skeleton sequences as the input to recover the original source and driving skeleton sequences) and contrastive losses as self-supervision. We present the architecture details in the following sections.

Motion Retargeting Architecture

Skeleton Sequence Modeling. The input skeleton sequence with ‘Character’ c and ‘Motion’ m is modeled by a spatio-temporal matrix, noted as $\mathbf{p}_{m,c} \in \mathbb{R}^{T \times V \times C_{in}}$. T , V , and C_{in} respectively represent the length of the video, the number of body joints in each frame, and the input channels ($C_{in} = 2$ for 2D data, or $C_{in} = 3$ if we use 3D skeletons). For each frame, the body joint coordinates are arranged in a matrix, with consistent order between different frames. The matrices of different frames are then stacked along the temporal dimension to obtain the matrix $\mathbf{p}_{m,c}$.

Skeleton Sequence Embedding As shown in Fig. 2 a.(i), ViA adopts a visual encoder E_{Pose} to extract features of the input driving skeleton sequence $\mathbf{p}_{m,c}$ and represents the features as $\mathbf{r}_{m,c} \in \mathbb{R}^{T' \times C_{out}} = E_{\text{Pose}}(\mathbf{p}_{m,c})$. T' is the size of temporal dimension after convolutions and C_{out} is the output channel size. In order to encourage the motion retargeting, we select a source skeleton sequence $\mathbf{p}_{m',c'}$ as a contrastive sample. For the selection of this contrastive sample which is expected to have different motion with driving sequence, we first cluster the skeleton sequences in the

Stages	E_{Pose}	D_{Pose}
Input	2D sequence [$T \times V, 2$]	Rep. [$T', 256$]
1	$\text{Conv} \begin{pmatrix} 1 \times 1, & 64 \\ 9 \times 1, & 64 \end{pmatrix} \times 4$	Upsample(2) $\text{Conv}(7, 128)$
2	$\text{Conv} \begin{pmatrix} 1 \times 1, & 128 \\ 9 \times 1, & 128 \end{pmatrix} \times 3$	Upsample(2) $\text{Conv}(7, 64)$
3	$\text{Conv} \begin{pmatrix} 1 \times 1, & 256 \\ 9 \times 1, & 256 \end{pmatrix} \times 3$	Upsample(2) $\text{Conv}(7, 2V)$
4	S-GAP ($2 \times V, 256$)	-
Rep.	[$T', 256$]	-
5	T-GAP ($T', 256$)	-
6	FC, Softmax	-
Output	Action Class	2D sequence [$T, 2V$]

Table 1: Main building blocks of the Encoder and Decoder networks. The dimensions of the kernels are denoted by $t \times s, c$ (2D kernels) and t, c (1D kernels) for temporal, spatial, channel sizes. S/T-GAP, FC denotes temporal/spatial global average pooling, and fully-connected layer respectively. Rep. indicates the learned representation.

training set using K-Means, and then randomly sample a sequence excluding the cluster where $\mathbf{p}_{m,c}$ belongs. The visual representation of the contrastive sample is obtained as $\mathbf{r}_{m',c'} = E_{Pose}(\mathbf{p}_{m',c'})$.

Our goal is to learn the view-invariant and generalizable skeleton action representation, which is generated by the visual encoder E_{Pose} . This encoder thus needs to have a strong capability to extract skeleton spatio-temporal features. To this end, we adopt the recent topology-free skeleton backbone network UNIK (Yang et al. 2021c) as the visual encoder. Specifically, the E_{Pose} is composed of 10 convolutional building blocks. Each building block contains a spatial network and a temporal convolutional network to extract both spatial and temporal multi-scale features from the skeleton sequence (see Tab. 1). For the spatial processing, we utilize 1×1 convolutions to expand the data channels and then multiply the features by uniformly initialized (He et al. 2015) and learnable dependency matrices (which replace the adjacency matrices used in GCN-based methods (Yan, Xiong, and Lin 2018; Shi et al. 2019; Liu et al. 2020; Chen et al. 2021)). For the temporal processing, we utilize 9×1 convolutions with strides. The size of the temporal dimension of embedded latent ‘Motion’ T' depends on the duration of the input sequence. For transfer-learning on downstream tasks, we attach the visual encoder to a temporal global average pooling layer and a fully-connected layer followed by a Softmax Layer. The output size of each fully-connected layer depends on the number of action classes. Then, we re-train the network with action labels.

Latent Motion Disentanglement. Latent Motion Disentanglement (LMD) is the key module of ViA to obtain the ‘Motion’ features on top of the visual representations (see Fig. 2 a.(ii)). Our insight is that the ‘Character’ and ‘Motion’ features are independent, these two components can be explicitly discovered in the latent code using the vector orthogonal decomposition method. As the ‘Character’ features are temporal static, we simply represent such features by a one-dimensional vector $\mathbf{r}_c \in \mathbb{R}^{C_{out}}$ for the skeleton

sequence $\mathbf{p}_{m,c}$. To model \mathbf{r}_c , we first pre-define a learnable orthogonal basis, noted as $\mathbf{D}_c = \{\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_K\}$ with $K \in [1, C_{out})$ where each vector indicates a basic visual transformation of the ‘Character’ aspects of the skeletons. Then, we consider the ‘Character’ features \mathbf{r}_c as a linear combination between this orthogonal basis \mathbf{D}_c , and associated magnitudes $A_c = \{a_1, a_2, \dots, a_K\}$. For $\mathbf{p}_{m',c'}$, we can obtain its ‘Character’ component $\mathbf{r}_{c'}$ in the same way:

$$\mathbf{r}_c = \sum_{i=1}^K a_i \mathbf{d}_i, \quad \mathbf{r}_{c'} = \sum_{i=1}^K a'_i \mathbf{d}_i. \quad (1)$$

Subsequently, the ‘Motion’ features $\mathbf{r}_m \in \mathbb{R}^{T' \times C_{out}}$ can be obtained by the vector decomposition: $\mathbf{r}_m = \mathbf{r}_{m,c} - \mathbf{r}_c$. As $\mathbf{r}_{m,c}$ has the temporal dimension of size T' , for each feature in the temporal dimension, we repeat the decomposition process with the same ‘Character’ component. According to the Gram-Schmidt algorithm, the set of magnitudes A_c and A'_c can be computed as the projections of $\mathbf{r}_{m,c}$ onto \mathbf{D}_c (see Eq. 2) to satisfy the orthogonality of \mathbf{r}_m and \mathbf{r}_c .

$$a_i = \frac{\langle \mathbf{r}_{m,c}, \mathbf{d}_i \rangle}{\|\mathbf{d}_i\|^2}, \quad a'_i = \frac{\langle \mathbf{r}_{m',c'}, \mathbf{d}_i \rangle}{\|\mathbf{d}_i\|^2}. \quad (2)$$

With such trained LMD, in the inference stage, we can generate skeleton sequences in multiple viewpoints with a single input skeleton sequence by only changing the magnitudes A_c along the orthogonal basis in the latent space without the need for paired data.

Skeleton Sequence Generation. To generate skeleton sequences, the output \mathbf{r}_m and $\mathbf{r}_{c'}$ are composed and fed into a decoder D_{Pose} (Fig. 2 a.(iii)). Unlike E_{Pose} , the Pose Sequence Decoder D_{Pose} is lightweight and designed by multiple 1D temporal convolutions with temporal global max pooling and upsampling to respectively encode and decode the skeletons as in (Aberman et al. 2019). The new retargeted skeleton sequence with ‘Motion’ m , and ‘Character’ c' , noted as $\mathbf{p}_{m,c'}$ is generated from the recombined features, $\mathbf{r}_m + \mathbf{r}_{c'}$. Similarly, $\mathbf{p}_{m',c}$ can also be generated by swapping the source and driving sequences. The skeleton sequence generation can be formulated as $\mathbf{p}_{m,c'} = D_{Pose}(\mathbf{r}_m, \mathbf{r}_{c'})$ and $\mathbf{p}_{m',c} = D_{Pose}(\mathbf{r}_{m'}, \mathbf{r}_c)$.

Self-supervised Training

In this section, we introduce the self-supervised training loss function \mathcal{L} , which consists of three components, namely reconstruction loss \mathcal{L}_{rec} , contrastive triplet loss \mathcal{L}_{trip} , and temporal velocity loss \mathcal{L}_{vel} :

$$\mathcal{L} = \mathcal{L}_{rec} + \mathcal{L}_{trip} + \mathcal{L}_{vel}. \quad (3)$$

Reconstruction Loss. The reconstruction loss aims at guiding the network towards a high generation quality at the global sequence level. It consists of two components: $\mathcal{L}_{rec} = \mathcal{L}_{self} + \mathcal{L}_{cycle}$. Specifically, at every training iteration, the decoder network D_{Pose} is firstly required to reconstruct each of the original input samples $\mathbf{p}_{m,c}$ using its representations \mathbf{r}_m and \mathbf{r}_c . This component of the loss is denoted as \mathcal{L}_{self} and formulated as a standard autoencoder reconstruction loss (see Eq. 4). Moreover, at each iteration,

the decoder is also encouraged to re-compose new combinations. To this end, we can explicitly apply the cycle reconstruction loss \mathcal{L}_{cycle} (see Eq. 4, 5) through the cycle generation. Specifically, the Embedding and LMD modules are used again to disentangle and re-combine the features for the previous generated skeleton sequence $\mathbf{p}_{m',c}$ and $\mathbf{p}_{m,c}$. Then, the generation module is also used again to reconstruct the original sequence $\mathbf{p}_{m,c}$.

$$\begin{aligned}\mathcal{L}_{self} &= \mathbb{E}[\|D_{\text{Pose}}(\mathbf{r}_m, \mathbf{r}_c) - \mathbf{p}_{m,c}\|^2], \\ \mathcal{L}_{cycle} &= \mathbb{E}[\|D_{\text{Pose}}(\hat{\mathbf{r}}_m, \hat{\mathbf{r}}_c) - \mathbf{p}_{m,c}\|^2],\end{aligned}\quad (4)$$

where

$$\begin{aligned}\hat{\mathbf{r}}_c &= \sum_{i=1}^K \frac{\langle E_{\text{Pose}}(\mathbf{p}_{m',c}) \cdot \mathbf{d}_i \rangle}{\|\mathbf{d}_i\|^2} \mathbf{d}_i, & \hat{\mathbf{r}}_{m'} &= E_{\text{Pose}}(\mathbf{p}_{m',c}) - \hat{\mathbf{r}}_c, \\ \hat{\mathbf{r}}_{c'} &= \sum_{i=1}^K \frac{\langle E_{\text{Pose}}(\mathbf{p}_{m,c'}) \cdot \mathbf{d}_i \rangle}{\|\mathbf{d}_i\|^2} \mathbf{d}_i, & \hat{\mathbf{r}}_m &= E_{\text{Pose}}(\mathbf{p}_{m,c'}) - \hat{\mathbf{r}}_{c'}.\end{aligned}\quad (5)$$

Contrastive Triplet Loss. At the representation level, we adopt a triplet loss \mathcal{L}_{trip} for the sampled skeletons (both driving skeleton sequence and its contrastive sample) as $\mathcal{L}_{trip} = \mathcal{L}_{trip_M} + \mathcal{L}_{trip_C}$. This loss aims to enhance the mutual information of \mathbf{r}_m and $\hat{\mathbf{r}}_m$, which are the representations of the same ‘Motion’ performed by different characters, to produce view-invariant ‘Motion’ representations. Specifically, we encourage the similarity between \mathbf{r}_m and $\hat{\mathbf{r}}_m$, while discouraging the similarity between $\hat{\mathbf{r}}_m$ and the other ‘Motion’ performed in its contrastive sample, $\mathbf{r}_{m'}$. Similarly, we define the ‘Character’ triplet loss in the same way. The loss components are described as follows (the triplet margin $\alpha = 1.0$):

$$\begin{aligned}\mathcal{L}_{trip_M} &= \mathbb{E}[\|\hat{\mathbf{r}}_m - \mathbf{r}_m\| - \|\hat{\mathbf{r}}_m - \mathbf{r}_{m'}\| + \alpha]_+, \\ \mathcal{L}_{trip_C} &= \mathbb{E}[\|\hat{\mathbf{r}}_c - \mathbf{r}_c\| - \|\hat{\mathbf{r}}_c - \mathbf{r}_{c'}\| + \alpha]_+.\end{aligned}\quad (6)$$

Velocity Loss. As described in (Aberman et al. 2019), the use of reconstruction loss only for sequence-level generation produces large errors for end joints such as hands and feet, which gives rise to the foot-skating phenomenon. We argue that the reconstruction loss constrains the network to generate the original skeletons with minimum global errors for all joints, however it misses the important temporal consistencies of each individual joint. We thus explicitly adopt a temporal consistency restriction loss for all V body joints (noted as an ensemble \mathcal{J}), which constrains the velocity—*i.e.*, joints shifting along the temporal dimension—of the skeleton sequence (see Eq. 7).

$$\mathcal{L}_{vel} = \lambda \mathbb{E}[\sum_{n \in \mathcal{J}} \|\mathcal{V}_n(D_{\text{Pose}}(\hat{\mathbf{r}}_m, \hat{\mathbf{r}}_c)) - \mathcal{V}_n(\mathbf{p}_{m,c})\|^2],\quad (7)$$

where \mathcal{V}_n denotes the velocity of the n -th joint, which can be calculated by the distance between this skeleton joint at frame τ -th and at frame $\tau + 1$ -th. λ indicates the weighting factor of the velocity loss.

Transfer-learning for Action Classification

ViA aims at pre-training a generic and view-invariant visual encoder. The model properties are verified by transfer-

learning of E_{Pose} (stacked with LMD) for action recognition tasks (see Fig. 2 (b)). In practice, we attach the visual encoder, where the pre-trained weights are used as initialization, to a temporal global average pooling layer and a fully-connected layer followed by a Softmax Layer. The output size of each fully-connected layer depends on the number of action classes. Then, we re-train the network with action labels on the target datasets. Following common evaluation protocols used in previous unsupervised action representation frameworks (Li et al. 2021b; Yang et al. 2021d; Su, Lin, and Wu 2021; Yang et al. 2021b), we conduct both *linear evaluation* by training only the fully-connected layer with the backbone frozen, and *fine-tuning evaluation* by further refining the whole network on downstream tasks.

Experiments and Analysis

We conduct extensive experiments to evaluate ViA. Firstly, we compare ViA against the state-of-the-art self-supervised models on the large-scale pre-training dataset **Posetics**, by linear evaluation. Secondly, we study the generalizability of ViA to quantify the performance improvement obtained by transfer-learning on the target 2D datasets (*i.e.*, **Toyota Smarthome**, **UAV-Human** and **Penn Action**) as well as 3D datasets (*i.e.*, **NTU-RGB+D 60** and **NTU-RGB+D 120**) after pre-training on Posetics. Thirdly, we evaluate the quality of the motion (*i.e.*, action) generated by the retargeting module on the synthetic dataset **Mixamo**. Finally, we provide an exhaustive ablation study of ViA.

Datasets and Evaluation Protocols

Posetics (Yang et al. 2021c) is created on top of Kinetics-400 (Carreira and Zisserman 2017) videos. It contains 142,000 real-world video clips of 320 action classes with the corresponding 2D and 3D skeletons. We use the Posetics dataset to pre-train our action representation learning framework with skeleton data and we study the transfer-learning on skeleton-based action classification. We use Top-1 and Top-5 accuracy as evaluation metrics (Yang et al. 2021c).

Toyota Smarthome (Das et al. 2019) (Smarthome) is a real-world dataset for daily living action classification and contains 16,115 videos of 31 action classes. It provides RGB videos, 2D and 3D skeleton data (Yang et al. 2021a). As the provided 2D data is more robust for action recognition even for cross-view evaluation (Yang et al. 2021a,c), unless otherwise stated, we use 2D data for the experiments. For the evaluation, we report mean per-class accuracy following the cross-subject (CS) and cross-view (CV1 and CV2) evaluation protocols.

UAV-Human (Li et al. 2021c) contains 22,476 video sequences collected by a flying UAV including 2D skeleton data estimated by (Fang et al. 2017). In this work, we use only 2D skeleton data and we follow Cross-subject (CS1 and CS2) evaluation protocols.

Penn Action (Zhang, Zhu, and Derpanis 2013) contains 2,326 video sequences of 15 different actions. In this work, we use 2D skeletons obtained by LCRNet++ (Rogez, Weinzaepfel, and Schmid 2019) for experiments and we report Top-1 accuracy following the standard train-test split.

NTU-RGB+D 60 (Shahroudy et al. 2016) consists of 56,880 sequences of high-quality 3D skeletons, captured by the Microsoft Kinect v2 sensor. We only use sequences of 3D skeletons in this work and we follow the cross-subject (CS) and cross-view (CV) evaluation protocol.

NTU-RGB+D 120 (Liu et al. 2020) extends the number of action classes and videos of NTU-RGB+D 60 to 120 classes 114,480 videos. We use 3D skeleton sequences and we follow the cross-subject (CS) and cross-set (CSet) evaluation protocols.

Mixamo (Inc. 2018) is a 3D animation collection, which contains approximately 2,400 unique motion sequences, including elementary actions, and various dancing moves. Each of these motions may be applied to 71 distinct characters, which share a human skeleton topology, but may differ in their body size and proportions. We use such a synthetic dataset, which includes the cross-reconstruction ground truth (*i.e.*, the same motion pattern performed by different characters and in different viewpoints obtained by rotated 3D and projected 2D skeletons) for pre-training and evaluating the motion retargeting module in ViA.

Evaluation on Self-supervised Action Classification

Our objective is to improve action recognition performance on 2D skeleton datasets by learning an action representation on a sufficiently large dataset. Hence, in this section, we evaluate ViA on self-supervised action classification (*i.e.*, linear evaluation) on the large-scale Posetics dataset and then compare ViA with state-of-the-art approaches.

Comparison with State-of-the-art (SoTA). For fair comparison, we re-implement recent state-of-the-art skeleton-based action representation learning approaches (Sun et al. 2020; Li et al. 2021b; Yang et al. 2021b) on the Posetics dataset using 2D skeleton data. Results are depicted in Tab. 2 (top): ViA is more effective when compared to 3D-based methods (Li et al. 2021b) applied onto 2D real-world datasets. Intuitively, we think that the variation of subject body sizes and of the viewpoints might weaken the robustness of the SoTA embedding networks. In contrast, ViA encourages similarity of the representation for actions performed by different subjects under different viewpoints. This shows that our model is more effective and robust to real-world videos. Compared to previous view-invariant embedding approaches (Sun et al. 2020) based on single frame, our method considering temporal features is more robust for action recognition. In Tab. 2 (bottom) we compare fine-tuning results of ViA to other supervised methods (Kim and Reiter 2017; Yan, Xiong, and Lin 2018; Shi et al. 2019; Song et al. 2020; Liu et al. 2020; Yang et al. 2021c) that are trained without representation learning (*i.e.*, training from scratch). Compared to the UNIK backbone model used in our work (Yang et al. 2021c), the pre-training provides minor improvement, as the training data (*i.e.*, Posetics) is sufficiently large. However, when transferring ViA onto smaller benchmark datasets, the impact of representation learning is significant (see Sec. Evaluation on Transfer-learning).

Methods	Posetics	
	Top-1(%)	Top-5(%)
Linear (Baseline)	8.2	21.4
Pr-ViPE (Sun et al. 2020)	17.2	35.3
OR-VPE (Yang et al. 2021b)	14.6	31.2
3s-CrosSCLR (Li et al. 2021b)	18.8	38.1
ViA (Ours)	20.7	40.1
TCNs (Kim and Reiter 2017)	34.0	57.2
ST-GCN (Yan, Xiong, and Lin 2018)	43.3	67.3
2s-AGCN (Shi et al. 2019)	47.0	70.8
Res-GCN (Song et al. 2020)	46.7	70.6
MS-G3D Net (Liu et al. 2020)	47.1	70.0
UNIK (Yang et al. 2021c)	47.6	71.3
ViA (Ours ft.)	48.0	72.6

Table 2: Comparison of Top-1 and Top-5 classification accuracy with state-of-the-art **unsupervised methods (top)** on Posetics. **Fully-supervised results (bottom)** with fine-tuning (reported as ft.) are also reported for reference.

Methods	NTU-60		NTU-120	
	CS(%)	CV(%)	CS(%)	CSet(%)
SeBiRe (Nie, Liu, and Liu 2020)	-	79.7	-	-
CrosSLR (Li et al. 2021b)	77.8	83.4	67.9	66.7
Colorization (Yang et al. 2021d)	75.2	83.1	-	-
ViA (Ours)	78.1	85.8	69.2	66.9
W/o pre-training	86.5	94.6	80.1	84.5
Self-supervised pre-training	89.6	96.4	85.0	86.5

Table 3: Comparison with previous self-supervised state-of-the-art by **linear evaluation (top)** on NTU-RGB+D 60 and NTU-RGB+D 120. Transfer learning results by **fine-tuning (bottom)** are also reported for reference.

Evaluation on Transfer-learning

In this section, we study the transfer ability of ViA by both linear evaluation and fine-tuning evaluation with self-supervised training on Posetics. We transfer the model onto three 2D skeleton action classification benchmarks *i.e.*, Toyota Smarhome, UAV-Human and Penn Action with no additional pre-training. As Smarhome and UAV-Human mainly focus on the cross-subject and cross-view challenges, the results measure the view- and subject-invariance of the 2D action representation of ViA models. We also report the results with supervised pre-training for reference (*i.e.*, we add a classifier at the end of the visual encoder and adopt cross entropy loss using action labels during training).

Linear Evaluation. Tab. 4 (top) shows the linear results on the three datasets. This evaluates the effectiveness of transfer-learning with fewer parameters (only the classifier is trained) compared to classification from random initialization. The results suggest that the weights of the model can be well pre-trained without action labels, providing a strong transfer ability especially on smaller benchmarks (*e.g.*, +31.9% Smarhome on CV2 and +70.4% on Penn Action) and the pre-trained visual encoder is generic enough to extract meaningful action features from skeleton sequences.

Fine-tuning. Tab. 4 (middle) shows the fine-tuning results, when the whole network is re-trained. These results suggest that pre-training can improve upon previous SoTA (Yang et al. 2021c) with no pre-training. The self-supervised pre-trained model also performs competitively

Methods	Pre-training	Toyota Smarthome				UAV-Human			Penn Action	
		#Params	CS(%)	CV1(%)	CV2(%)	#Params	CS1(%)	CS2(%)	#Params	Top-1(%)
Random init.	Scratch	7.97K	24.6	17.2	20.7	39.85K	3.8	4.1	3.85K	29.8
Supervised	Posetics w/ labels	7.97K	51.9	35.4	52.2	39.85K	32.9	56.1	3.85K	97.3
Self-supervised	Posetics w/o labels	7.97K	49.5	33.6	52.6	39.85K	29.5	46.7	3.85K	90.2
Random init.	Scratch	3.45M	63.1	22.9	61.2	3.45M	39.2	67.3	3.45M	94.0
Supervised	Posetics w/ labels	3.45M	64.5	36.1	65.2	3.45M	42.6	69.5	3.45M	98.0
Self-supervised	Posetics w/o labels	3.45M	64.0	35.6	65.4	3.45M	41.3	68.5	3.45M	97.7
Previous SoTA	-	(Yang et al. 2021c)			(Cheng et al. 2020)			(Sun et al. 2020)		
	-	-	63.1	22.9	61.2	-	38.0	67.0	-	97.5
ViA (Ours)	-	-	64.5	36.1	65.4	-	42.6	69.5	-	98.0

Table 4: Transfer learning results by **linear evaluation (top)** and **fine-tuning (middle)** on Smarthome, UAV-Human and Penn Action with self-supervised pre-training on Posetics compared to Baseline (random initialization). Results with supervised pre-training and **previous state-of-the-art (bottom)** are also reported.

Methods	Pre-training	Training data	Toyota Smarthome			UAV-Human		Penn Action	
			CS(%)	CV1(%)	CV2(%)	CS1(%)	CS2(%)	Top-1 Accuracy(%)	
Random init. (Yang et al. 2021c)	Scratch	5%	22.9	5.6	33.7	10.9	10.4	32.4	
Self-supervised	Posetics w/o labels	5%	38.6	16.8	42.6	21.7	33.3	65.8	
Random init. (Yang et al. 2021c)	Scratch	10%	33.8	8.5	39.5	17.8	25.6	39.8	
Self-supervised	Posetics w/o labels	10%	45.3	22.7	46.6	31.0	43.7	85.2	

Table 5: Transfer learning results by **fine-tuning** on all benchmarks of Smarthome, UAV-Human and Penn Action with randomly selected **5% (top)** and **10% (bottom)** of labeled training data.

compared to supervised pre-trained models. From these results we conclude that collecting a large-scale video dataset, without the need of action annotation, can be beneficial to downstream tasks, especially when using our proposed view-invariant ViA for the 2D action classification task (*e.g.*, +12.7% on Smarthome CV1 and +4.2% on CV2). Furthermore, we compare our fine-tuning results to other SoTA skeleton-based supervised approaches (Yang et al. 2021c; Cheng et al. 2020; Sun et al. 2020). The results in Tab. 4 (bottom) show that ViA outperforms all previous approaches on all the three real-world datasets.

Training with Fewer Labels. In some real-world applications, labeled data may be lacking, which makes it challenging to train models with good performance. To evaluate ViA in such cases, we pre-train with Posetics and then fine-tune the visual encoder with 5% and 10% of the labeled data. As shown in Tab. 5, without pre-training, the accuracy of the baseline (Yang et al. 2021c) significantly decreases with the amount of training data. In contrast, ViA still achieves good performance on all three datasets.

3D Skeleton Action Classification. As ViA can be simply extended to take 3D skeleton sequence as input, we further analyze the transfer ability of ViA onto 3D skeleton action recognition tasks. We firstly compare (Li et al. 2021b), (Yang et al. 2021b) and ViA on Posetics using officially provided 3D data (we get 17.1%, 12.9% and 19.3%, respectively for linear evaluation). These results are lower than related results in Tab. 2 using 2D data. We argue that, although 3D skeletons are more robust to the view variation, 2D skeletons extracted from images or videos tend to be more accurate than extracted 3D skeletons. In contrast, laboratory datasets (*e.g.*, NTU-RGB+D 60 & 120) provide 3D skeleton data obtained by RGBD sensors that have a higher quality than the one provided by 2D data. To study the impact of action representation learning, we also trans-

Methods	Sup.	Unsup.
NKN (Villegas et al. 2018)	1.51	-
MotionRetargeting2D (Aberman et al. 2019)	0.96	2.56
ViA w/o \mathbf{D}_c (Ours)	2.42	-
ViA w/ \mathbf{D}_c (Ours)		
size $K = 2$	1.16	-
size $K = 64$	0.89	-
size $K = 32$	0.86	2.47

Table 6: Quantitative comparisons of Mean Square Error (MSE) show that our framework outperforms other SoTA motion retargeting methods on Mixamo.

fer the ViA pre-trained on Posetics (3D skeletons) without action labels onto NTU-RGB+D-60 and NTU-RGB+D-120 by fine-tuning. The action recognition performance can still be improved (*e.g.*, +4.9% on NTU-RGB+D-120 CS, see Tab. 3 (bottom)). To compare with other recent self-supervised methods (Nie, Liu, and Liu 2020; Li et al. 2021b; Yang et al. 2021d), we follow the same pre-training setting and linear evaluation protocol and report state-of-the-art accuracy (see Tab. 3 (top)).

Evaluation on Motion Retargeting

As motion retargeting is our pretext task, we here evaluate the proposed LMD mechanism of ViA by motion retargeting performance. We randomly split training and test sets on the Mixamo dataset and follow the same setting and protocol described in (Aberman et al. 2019). As previous SoTA (Villegas et al. 2018; Aberman et al. 2019) are supervised approaches, for fair comparison, we also train ViA using cross reconstruction loss on (Inc. 2018) in a supervised manner. To validate the impact of proposed \mathbf{D}_c , we learn \mathbf{r}_c by decomposing $\mathbf{r}_{m,c}$ on a pre-defined and fixed subspace without the learnable \mathbf{D}_c . From the evaluation results reported in Tab. 6, we observe that in the absence of \mathbf{D}_c , model fails to generate high-quality skeletons, which proves the effectiveness

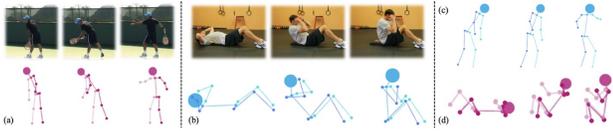


Figure 3: Qualitative results on **Motion Retargeting**. (a) and (b) are the input pair of videos and corresponding 2D skeleton sequences. (c) is the generated 2D skeleton sequence that represents the character of (b) performing the motion in (a) while maintaining the viewpoint and body size invariance. (d) is the generated 2D skeleton sequence that represents the character of (a) performing the motion in (b).

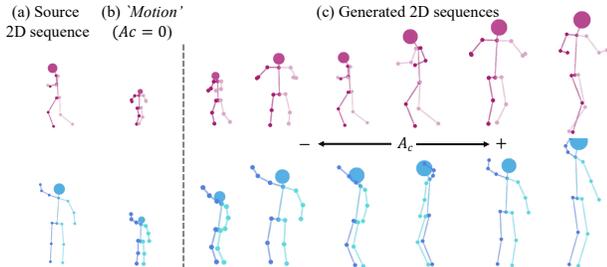


Figure 4: Qualitative results on **2D Motion Generation**. Given a source skeleton sequence, we can generate multiple sequences by latent space manipulation on disentangled ‘Motion’ and ‘Character’ magnitudes (A_c).

of D_c . Then we conduct an ablation analysis on the size of D_c . The results suggest that motion retargeting by ViA (w/ D_c in including 32 directions) performs the best and achieve SoTA accuracy. We also report unsupervised results by cycle consistency learning.

Qualitative Evaluation. To demonstrate that the ‘View/Subject’ and ‘Motion’ are well disentangled by the proposed framework, we visualize an example of motion retargeting inference of two Penn Action’s videos (see Fig. 3). Then we visualize the representations of all the skeletons on Mixamo with t-SNE (see Fig. 5 with both supervised and unsupervised motion retargeting). Qualitative results validate that the ‘View/Subject’ and ‘Motion’ parts of 2D skeleton sequences have been effectively disentangled. To further understand the learned ‘Motion’ features, we generate skeletons for each single input sequence with only r_m (see Fig. 4 (b)), and then with r_m combined by different r_c obtained by a linearly grown A_c (see Fig. 4 (c)). We observe that r_m represents the motion in a ‘canonical view’, regardless of original views of the input skeleton sequences. As such a ‘canonical view’ can be considered as a normalized form of the given skeleton sequence, learning transformations between generative sequence and source sequence using D_c and A_c is considerably more efficient than directly generation once the ‘canonical view’ is fixed.

Ablation Study

To understand the contribution of each loss function in ViA, we conduct ablation experiments on Smarthome CV2 with fine-tuning protocol. To perform more studies on the characteristics of view-invariant representations, we additionally

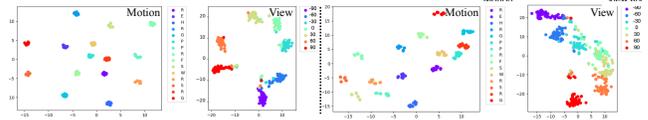


Figure 5: Skeleton representations (marked by different colors with ‘Motion’ and ‘View’) on Mixamo with ViA by supervised (left) and unsupervised (right) motion retargeting.

Methods	\mathcal{L}_{self}	\mathcal{L}_{cycle}	\mathcal{L}_{trip}	\mathcal{L}_{vel}	Smarthome	Mixamo
					CV2(%)	CV(%)
L0: Baseline					61.7	71.7
L1: +Self	✓				62.9	76.5
L2: +Cycle	✓	✓			63.8	82.5
L3: +Trip	✓	✓	✓		65.0	85.8
L4: +Vel (Full)	✓	✓	✓	✓	65.4	87.2

Table 7: Ablation study of ViA on Smarthome CV2 and Mixamo CV with transfer learning (fine-tuning).

set a Cross-view (CV) action recognition protocol on the Mixamo dataset (*i.e.*, Top-1 classification accuracy) using two different 2D skeleton projections generated by random 3D rotations of the same action for 2D cross-view evaluation. We start from a baseline model that has been previously pre-trained on the synthetic dataset (*i.e.*, Mixamo) using motion retargeting annotations for cross-character reconstruction. Already this visual encoder has a strong capability to embed the 2D skeleton sequence into a view-invariant representation. The results in Tab. 7 (see L0) suggest that the generalizability is hindered by the lack of action diversity in the synthetic training dataset if directly transferring the baseline visual encoder for action classification. Therefore, from the full results in Tab. 7, we infer that additional self-supervised training on Posetics can improve the real-world generalizability. Specifically, a **self reconstruction loss (L1)** can help the visual encoder learn the global characteristics of the real-world data and thus facilitate the classification. The **cycle reconstruction loss (L2)** and the **triplet loss (L3)** aim at maximizing the embedding similarity between the same action performed from two viewpoints, while minimizing the embedding similarity between different actions. These losses are instrumental in extracting a more generic representation for the downstream action classification task. Finally, The **velocity loss (L4)** contributes to minor boosts.

Conclusion

In this work, we presented ViA, a generic framework aimed at learning view-invariant skeleton action representation via Latent Motion Disentanglement. We showed that self-supervised motion retargeting with contrastive learning can be an effective pretext task to learn view-invariant action representation for real-world 2D skeleton sequences. Experimental analysis confirmed that a visual encoder extracting such representation on large-scale datasets such as Posetics significantly boosts the performance when transferred onto downstream target datasets for cross-subject and cross-view action classification tasks. Future work will extend our method to RGB data in order to capture the information on objects used to perform actions, while maintaining view- and subject-invariance.

References

- Aberman, K.; Wu, R.; Lischinski, D.; Chen, B.; and Cohen-Or, D. 2019. Learning character-agnostic motion for motion retargeting in 2D. *ACM TOG*.
- Arnab, A.; Deghani, M.; Heigold, G.; Sun, C.; Lučić, M.; and Schmid, C. 2021. ViViT: A Video Vision Transformer. *ICCV*.
- Bachman, P.; Hjelm, R. D.; and Buchwalter, W. 2019. Learning Representations by Maximizing Mutual Information Across Views. In *NeurIPS*.
- Caetano, C.; Brémond, F.; and Schwartz, W. 2019. Skeleton Image Representation for 3D Action Recognition Based on Tree Structure and Reference Joints. *SIBGRAPI*.
- Carreira, J.; and Zisserman, A. 2017. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *CVPR*.
- Chan, C.; Ginosar, S.; Zhou, T.; and Efros, A. A. 2019. Everybody Dance Now. In *ICCV*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *ICML*.
- Chen, Y.; Zhang, Z.; Yuan, C.; Li, B.; Deng, Y.; and Hu, W. 2021. Channel-Wise Topology Refinement Graph Convolution for Skeleton-Based Action Recognition. In *ICCV*.
- Cheng, K.; Zhang, Y.; He, X.; Chen, W.; Cheng, J.; and Lu, H. 2020. Skeleton-Based Action Recognition With Shift Graph Convolutional Network. In *CVPR*.
- Das, S.; Dai, R.; Koperski, M.; Minciullo, L.; Garattoni, L.; Bremond, F.; and Francesca, G. 2019. Toyota Smarthome: Real-World Activities of Daily Living. In *ICCV*.
- Duan, H.; Zhao, Y.; Chen, K.; Shao, D.; Lin, D.; and Dai, B. 2022. Revisiting Skeleton-based Action Recognition. In *CVPR*.
- Fang, H.-S.; Xie, S.; Tai, Y.-W.; and Lu, C. 2017. RMPE: Regional Multi-person Pose Estimation. In *ICCV*.
- Feichtenhofer, C. 2020. X3D: Expanding Architectures for Efficient Video Recognition. In *CVPR*.
- Feichtenhofer, C.; Fan, H.; Malik, J.; and He, K. 2019. Slow-Fast Networks for Video Recognition. In *ICCV*.
- Feichtenhofer, C.; Pinz, A.; and Zisserman, A. 2016. Convolutional Two-Stream Network Fusion for Video Action Recognition. In *CVPR*.
- Hara, K.; Kataoka, H.; and Satoh, Y. 2017. Learning Spatio-Temporal Features with 3D Residual Networks for Action Recognition. In *ICCV*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *CVPR*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *ICCV*.
- Hjelm, R. D.; Fedorov, A.; Lavoie-Marchildon, S.; Grewal, K.; Bachman, P.; Trischler, A.; and Bengio, Y. 2019. Learning deep representations by mutual information estimation and maximization. In *ICLR*.
- Inc., A. S. 2018. Mixamo. <https://www.mixamo.com>. <https://www.mixamo.com>. Accessed: 2018-12-27.
- Ji, S.; Xu, W.; Yang, M.; and Yu, K. 2013. 3D Convolutional Neural Networks for Human Action Recognition. *IEEE TPAMI*.
- Jiao, Y.; Xiong, Y.; Zhang, J.; Zhang, Y.; Zhang, T.; and Zhu, Y. 2020. Sub-graph Contrast for Scalable Self-Supervised Graph Representation Learning. In *ICDM*.
- Karen, S.; and Andrew, Z. 2014. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*.
- Kim, T. S.; and Reiter, A. 2017. Interpretable 3D Human Action Analysis with Temporal Convolutional Networks. In *CVPRW*.
- Kundu, J. N.; Gor, M.; Uppala, P. K.; and Radhakrishnan, V. B. 2019. Unsupervised Feature Learning of Human Actions As Trajectories in Pose Embedding Manifold. In *WACV*.
- Li, J.; Wong, Y.; Zhao, Q.; and Kankanhalli, M. S. 2018. Unsupervised Learning of View-invariant Action Representations. In *NeurIPS*.
- Li, K.; Li, X.; Wang, Y.; Wang, J.; and Qiao, Y. 2021a. CT-Net: Channel Tensorization Network for Video Classification. In *ICLR*.
- Li, L.; Wang, M.; Ni, B.; Wang, H.; Yang, J.; and Zhang, W. 2021b. 3D Human Action Representation Learning via Cross-View Consistency Pursuit. In *CVPR*.
- Li, M.; Chen, S.; Chen, X.; Zhang, Y.; Wang, Y.; and Tian, Q. 2019. Actional-Structural Graph Convolutional Networks for Skeleton-Based Action Recognition. In *CVPR*.
- Li, T.; Liu, J.; Zhang, W.; Ni, Y.; Wang, W.; and Li, Z. 2021c. UAV-Human: A Large Benchmark for Human Behavior Understanding With Unmanned Aerial Vehicles. In *CVPR*.
- Liu, J.; Shahroudy, A.; Perez, M.; Wang, G.; Duan, L. Y.; and Kot, A. C. 2020. NTU RGB+D 120: A Large-Scale Benchmark for 3D Human Activity Understanding. *IEEE TPAMI*.
- Liu, Z.; Zhang, H.; Chen, Z.; Wang, Z.; and Ouyang, W. 2020. Disentangling and Unifying Graph Convolutions for Skeleton-Based Action Recognition. In *CVPR*.
- Ma, X.; Rahmani, H.; Fan, Z.; Yang, B.; Chen, J.; and Liu, J. 2022. Remote: Reinforced motion transformation network for semi-supervised 2d pose estimation in videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Nie, Q.; and Liu, Y. 2021. View transfer on human skeleton pose: Automatically disentangle the view-variant and view-invariant information for pose representation learning. *IJCV*.
- Nie, Q.; Liu, Z.; and Liu, Y. 2020. Unsupervised Human 3D Pose Representation with Viewpoint and Pose Disentanglement. In *ECCV*.
- Rogez, G.; Weinzaepfel, P.; and Schmid, C. 2019. LCR-Net++: Multi-person 2D and 3D Pose Detection in Natural Images. *IEEE TPAMI*.
- Ryoo, M.; Piergiovanni, A.; Kangaspunta, J.; and Angelova, A. 2020. AssembleNet++: Assembling Modality Representations via Attention Connections. *ECCV*.

- Sardari, F.; Ommer, B.; and Mirmehdi, M. 2021. Unsupervised View-Invariant Human Posture Representation. In *BMVC*.
- Shahroudy, A.; Liu, J.; Ng, T.-T.; and Wang, G. 2016. NTU RGB+D: A Large Scale Dataset for 3D Human Activity Analysis. *CVPR*.
- Shi, L.; Zhang, Y.; Cheng, J.; and Lu, H. 2019. Two-Stream Adaptive Graph Convolutional Networks for Skeleton-Based Action Recognition. In *CVPR*.
- Skorokhodov, I.; Tulyakov, S.; and Elhoseiny, M. 2022. StyleGAN-V: A Continuous Video Generator with the Price, Image Quality and Perks of StyleGAN2. In *CVPR*.
- Song, Y.-F.; Zhang, Z.; Shan, C.; and Wang, L. 2020. Stronger, Faster and More Explainable: A Graph Convolutional Baseline for Skeleton-Based Action Recognition. In *ACM MM*.
- Su, Y.; Lin, G.; and Wu, Q. 2021. Self-Supervised 3D Skeleton Action Representation Learning With Motion Consistency and Continuity. In *ICCV*.
- Sun, J. J.; Zhao, J.; Chen, L.-C.; Schroff, F.; Adam, H.; and Liu, T. 2020. View-Invariant Probabilistic Embedding for Human Pose. In *ECCV*.
- Tian, Y.; Krishnan, D.; and Isola, P. 2020. Contrastive Multiview Coding. In *ECCV*.
- Tulyakov, S.; Liu, M.-Y.; Yang, X.; and Kautz, J. 2018. MoCoGAN: Decomposing motion and content for video generation. In *CVPR*.
- Vemulapalli, R.; Arrate, F.; and Chellappa, R. 2014. Human Action Recognition by Representing 3D Skeletons as Points in a Lie Group. *CVPR*.
- Villegas, R.; Yang, J.; Ceylan, D.; and Lee, H. 2018. Neural Kinematic Networks for Unsupervised Motion Retargeting. In *CVPR*.
- Wang, J.; Nie, X.; Xia, Y.; Wu, Y.; and Zhu, S.-C. 2014. Cross-view Action Modeling, Learning and Recognition. In *CVPR*.
- Wang, L.; Tong, Z.; Ji, B.; and Wu, G. 2021. TDN: Temporal Difference Networks for Efficient Action Recognition. In *CVPR*.
- Wang, Y.; Bilinski, P.; Bremond, F.; and Dantcheva, A. 2020. G3AN: Disentangling Appearance and Motion for Video Generation. In *CVPR*.
- WANG, Y.; Bilinski, P.; Bremond, F.; and Dantcheva, A. 2020. ImaGINator: Conditional Spatio-Temporal GAN for Video Generation. In *WACV*.
- Wang, Y.; Yang, D.; Bremond, F.; and Dantcheva, A. 2022. Latent Image Animator: Learning to Animate Images via Latent Space Navigation. In *ICLR*.
- Weinzaepfel, P.; and Rogez, G. 2021. Mimetics: Towards Understanding Human Actions Out of Context. *IJCV*.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised Feature Learning via Non-parametric Instance Discrimination. In *CVPR*.
- Yan, S.; Xiong, Y.; and Lin, D. 2018. Spatial Temporal Graph Convolutional Networks for Skeleton-Based Action Recognition. *AAAI*.
- Yang, D.; Dai, R.; Wang, Y.; Mallick, R.; Minciullo, L.; Francesca, G.; and Bremond, F. 2021a. Selective Spatio-Temporal Aggregation Based Pose Refinement System: Towards Understanding Human Activities in Real-World Videos. In *WACV*.
- Yang, D.; Wang, Y.; Dantcheva, A.; Garattoni, L.; Francesca, G.; and Bremond, F. 2021b. Self-Supervised Video Pose Representation Learning for Occlusion-Robust Action Recognition. In *FG*.
- Yang, D.; Wang, Y.; Dantcheva, A.; Garattoni, L.; Francesca, G.; and Bremond, F. 2021c. UNIK: A Unified Framework for Real-world Skeleton-based Action Recognition. In *BMVC*.
- Yang, S.; Liu, J.; Lu, S.; Er, M. H.; and Kot, A. C. 2021d. Skeleton Cloud Colorization for Unsupervised 3D Action Representation Learning. In *ICCV*.
- Yu, S.; Tack, J.; Mo, S.; Kim, H.; Kim, J.; Ha, J.-W.; and Shin, J. 2022. Generating Videos with Dynamics-aware Implicit Generative Adversarial Networks. In *ICLR*.
- Zhang, W.; Zhu, M.; and Derpanis, K. G. 2013. From Actemes to Action: A Strongly-Supervised Representation for Detailed Action Understanding. In *ICCV*.
- Zhao, L.; Wang, Y.; Zhao, J.; Yuan, L.; Sun, J. J.; Schroff, F.; Adam, H.; Peng, X.; Metaxas, D.; and Liu, T. 2021. Learning View-Disentangled Human Pose Representation by Contrastive Cross-View Mutual Information Maximization. In *CVPR*.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *ICCV*.