Collaborating Domain-shared and Target-specific Feature Clustering for Cross-domain 3D Action Recognition

Qinying Liu[®] and Zilei Wang[®]

University of Science and Technology of China, Hefei, China lydyc@mail.ustc.edu.cn, zlwang@ustc.edu.cn

Abstract. In this work, we consider the problem of cross-domain 3D action recognition in the open-set setting, which has been rarely explored before. Specifically, there is a source domain and a target domain that contain the skeleton sequences with different styles and categories, and our purpose is to cluster the target data by utilizing the labeled source data and unlabeled target data. For such a challenging task, this paper presents a novel approach dubbed CoDT to collaboratively cluster the domain-shared features and target-specific features. CoDT consists of two parallel branches. One branch aims to learn domain-shared features with supervised learning in the source domain, while the other is to learn target-specific features using contrastive learning in the target domain. To cluster the features, we propose an online clustering algorithm that enables simultaneous promotion of robust pseudo label generation and feature clustering. Furthermore, to leverage the complementarity of domain-shared features and target-specific features, we propose a novel collaborative clustering strategy to enforce pair-wise relationship consistency between the two branches. We conduct extensive experiments on multiple cross-domain 3D action recognition datasets, and the results demonstrate the effectiveness of our method. Code is at CoDT.

Keywords: skeleton-based action recognition, cross-domain, open-set.

1 Introduction

Recent advances in 3D depth cameras and pose estimation algorithms have made it possible to estimate 3D skeletons quickly and accurately. In contrast to RGB images, the skeleton data only contain the coordinates of human keypoints, providing high-abstract and environment-free information. Thus, 3D action recognition (*a.k.a.*, skeleton-based action recognition) is attracting more and more attentions [74,50,80]. Nevertheless, existing methods mainly focus on the traditional supervised classification. In this learning paradigm, it is assumed that the labeled training (source) dataset and unlabeled test (target) dataset have the same distribution. In practical scenarios, it is not easy to hold such assumption, because labeling a dataset with the same distribution as the target data is a laborious task. In reality, it is preferred to utilize a related public annotated



Fig. 1. (a) Illustration of CD-SAR task and our proposed CoDT method. CD-SAR aims to group the target samples into semantic clusters by virtue of labeled source data and unlabeled target data. The core idea of CoDT is to collaborate the supervised learning of source data and contrastive learning of target data. (b) Illustration of different joint definitions in the NTU-60 [58] and Skeletics [25]. The figure is modified from [75].

skeleton or even image dataset as a source dataset. Unfortunately, there is typically discrepancy (a.k.a, domain gap [2]) between the source and target datasets due to various factors, including the devices (e.g., 3D sensors [83,18], imagebased pose estimation algorithms [5,37]), the camera setup (*e.g.*, viewpoints), the scenes (e.g., in-the-lab or in-the-wild), etc. These factors make the skeletons of different datasets distinct in styles (e.g., joint types, qualities, views) and action categories. For example, consider a very practical application that needs to recognize the action classes of the in-the-wild unlabeled 3D dataset Skeletics [25], we might seek help from the large-scale in-the-lab dataset NTU-60 [58] with highquality skeletons and annotated labels. However, the NTU-60 and Skeletics are captured by Microsoft Kinect V2 camera [18] and the pose estimation method VIBE [37] respectively, resulting in different styles of skeletons, e.g., different definitions of human joints (see Fig. 1(b)). Furthermore, NTU-60 mainly contains indoor actions, while Skeletics contains more unconstrained actions from internet videos, thereby leading to different action categories. In summary, the domain gap problem is very practical in skeleton-based action recognition but rarely studied in the literature. In this paper, we present the first systematic study to the domain gap problem in skeleton-based action recognition, where the source and target datasets have different styles and categories. Obviously, this is an open-set problem, and it is expected that the algorithm can automatically cluster the target samples into latent classes.

The labeled source dataset, although collected from a different domain from the target domain, is helpful to develop *domain-shared* representations. By training the models in the standard supervised manner, the representations are expected to be discriminative to different action categories. However, the models generally generalize poorly to the target domain due to the domain gap [75,24]. On the other hand, the target domain has many unlabeled samples, which can be used to learn more *target-specific* representations. This suggests the possibility of applying recently emerging contrastive learning [28,11,23] on target data. The contrastive learning optimizes the features by instance discrimination, *i.e.*, the features are enforced to be invariant for different transformations of an instance and distinct for different instances. By learning to attract or repel different instances, the features appear to automatically capture some extent of semantic similarity. Yet, they may not have enough discrimination to action categories.

Based on the above discussions, we consider that the domain-shared representations and target-specific representations are conceptually complementary. This motivates us to integrate the supervised learning on labeled source data and contrastive learning on unlabeled target data. Previous methods [56,89,33] commonly implement such integration through multi-task learning with a shared feature encoder. However, in our case, the two domains may differ considerably in styles and even joint types, and thus we argue that the domain-shared and target-specific features should be learned via different models in order to sufficiently exploit their respective merits. With the purpose of clustering target samples, a natural problem arises: is it possible to collaborate the two models on feature clustering? To achieve this goal, we need to address two key issues. The first one is how to cluster the features for a single model, which is expected to be optimized in an end-to-end manner rather than computed by some offline clustering algorithms (e.q., k-means) whose usefulness is proven to be limited [6]. The second one is how to collaborate both models to jointly optimize feature clustering. It is quite challenging since the learned clusters from two models cannot be matched exactly due to the lack of labels.

In this paper, we propose a novel Collaborating Domain-shared and Targetspecific features clustering (CoDT) network for the task of cross-domain skeletonbased action recognition (CD-SAR). Fig. 1 illustrates the overview of the task and our method. Specifically, to address the first issue, we propose an online clustering algorithm to generate robust pseudo labels to guide feature clustering. It is built upon the teacher-student framework [64], where the clustering is optimized via the pseudo labels generated by the teacher model. Under this framework, a straightforward way to determine pseudo label is to select the cluster with which the teacher is most confident. However, due to the discrepancy between target sample clustering and source-based supervised learning (or instance-based contrastive learning), there is a risk of obtaining a trivial solution that groups most samples into only a few clusters, as observed in Sec. 4.3. To make the clusters balanced, we propose to generate uniformly distributed pseudo labels. It is non-trivial to achieve it online as we have to take into consideration the global distribution of pseudo labels. Hence, we transform it to an optimal transport problem that can be solved by linear programming [14].

As for the second issue, we propose a collaborative clustering strategy that exchanges pseudo labels across models for collaborative training (co-training). In the traditional co-training methods [4,57,26], the categories of two models are pre-defined and consistent, and thus the pseudo labels produced by one model can be directly used to train another model. In our case, however, the semantics of the learned clusters of two models are agnostic and variable during

training, making it difficult to determine the correlation between the clusters of two models. To this end, we propose to perform co-training on the pair-wise relationship that represents whether a pair of samples are from the same cluster (positive pair) or distinct clusters (negative pair). Specifically, we first construct the pair-wise binary pseudo labels by comparing the instance-wise pseudo labels of samples. Thereafter, the pair-wise labels are used to train the other model, where a novel contrastive loss is particularly adopted to enforces the model to produce consistent/inconsistent predictions for the positive/negative pairs.

Our contributions are summarized as: 1) We provide a benchmark for CD-SAR. To solve this task, we propose a novel two-branch framework dubbed CoDT to collaborate domain-shared and target-specific features. 2) We propose an online clustering algorithm that can alternate the robust pseudo label generation and balanced feature clustering. 3) We propose a collaborative clustering algorithm, which enables co-training of two models to enforce their consistency in terms of pair-wise relationship. 4) We evaluate our method upon different cross-domain tasks, and the effectiveness of our method is well shown.

2 Related Work

Unsupervised Representation Learning and Visual Clustering In the field of unsupervised learning, there are two main research topics: representation learning and image clustering. The former focuses on training the feature encoder by self-supervised learning. To achieve this, existing methods either design numerous pre-designed pretext tasks [16,55,81,52] or perform contrastive learning [11,28,23]. Despite these efforts, these approaches are mainly used for pretraining. As an alternative, image clustering methods simultaneously optimize clustering and representation learning. Previous methods train the model using the pseudo labels derived from the most confident samples [10,70,54,12], or through cluster re-assignments [6,7]. Recently, [1,8] are proposed to apply a balanced label assignment. In this work, we take advantage of the idea in [1,8] but incorporate it into the student-teacher network.

Supervised and Unsupervised 3D Action Recognition To tackle skeletonbased action recognition, many RNN-based methods [37,79,61] and CNN-based methods [35,41] are carried out. Recently, GCN-based methods [74,40,50,59,80] have attracted increasing attention due to their outstanding performances. We adopt the widely-used ST-GCN [74] as the backbone. There are also many unsupervised methods [72,53,86,45,44] proposed to learn the skeleton representations. In particular, many methods [88,62,38,63,76] utilize the encoder-decoder structure to reconstruct skeletons from the encoded features. CrosSCLR [44] proposes to apply contrastive learning for skeleton representation learning. A recent work [75] studies the generalizability of the models by first pretraining a model on one dataset and then finetuning it on another dataset. Our task is different from the above works, as we do not require the ground-truth labels of target data to train or finetune the model.

5

Close-set and Open-set Transfer Learning In (close-set) unsupervised domain adaptation (UDA) [2,3,69], the source and target datasets are different in styles but have an identical label set. Many methods aim to learn domaininvariant features by adversarial learning [20,13] or explicitly reducing the distribution discrepancy [73]. For example, GVB [13] proposes a gradually vanishing bridge layer to facilitate the adversarial training. Yet, these methods are not suitable for the open-set problem [51]. Novel class discovery (NCD) [27] aims to transfer knowledge between datasets with different categories but almost the same style. UNO [19] trains the source set with ground-truth labels and the target set with pseudo labels generated by [8]. Another similar task to ours is crossdomain person re-identification (CD-ReID) [68,15]. Representative methods are based on clustering [17,82,87] or domain-invariant feature learning [49,31,46]. Among them, a series of methods [21,85,78] propose a collaborative learning scheme among multiple peer networks to alleviate the effects of label noise. However, the multiple networks are just initialized differently, making them essentially different from our method. Recently, cross-domain few-shot learning (CD-FSL) [65,24] has emerged, where the source and target datasets are drawn from different domains and classes. A few methods [56,77,32] relax the task to allow access to unlabeled target data during training. STARTUP [56] uses the model pretrained on the source dataset to produce soft pseudo labels for target samples, and then finetunes the model using target samples and their soft labels.

3 Method

3.1 Problem Formulation

Formally, we denote the labeled source domain as $\mathbb{D}^{l} = (\boldsymbol{x}_{n}^{s}, y_{n}^{s})|_{n=1}^{N^{s}}$, where \boldsymbol{x}_{n}^{s} and y_{n}^{s} are the *n*-th training sample and its associated action label, N^{s} is the number of source skeleton sequences. The N^{t} unlabeled target skeleton sequences are denoted as $\mathbb{D}^{t} = \{\boldsymbol{x}_{n}^{t}|_{n=1}^{N^{t}}\}$, which are not associated with any label. Our goal is to mine the latent classes of \mathbb{D}^{t} , which are disjoint with that of \mathbb{D}^{s} .

3.2 Overview

The framework of CoDT is depicted in Fig. 2. It is composed of two parallel branches denoted by \mathcal{B}_0 and \mathcal{B}_1 , where \mathcal{B}_0 processes the data of both domains for exploiting domain-shared knowledge and \mathcal{B}_1 only processes the target data for capturing target-specific characteristics. Each branch contains a base module (BM) dubbed BM- \mathcal{B}_0 and BM- \mathcal{B}_1 to optimize the feature encoders. The BM- \mathcal{B}_0 is to learn discriminative information from source data via supervised learning. To make the features more domain-invariant, we add a domain-shared decoder in BM- \mathcal{B}_0 for skeleton reconstruction. The BM- \mathcal{B}_1 is to learn semantic similarity of target data with contrastive learning. Then we aim to collaboratively train the two branches and encourage their agreement on the clustering of target data. Yet, there are two problems: 1) How to optimize the feature clustering? 2) How to achieve co-training? To address the problems, we propose the online clustering module (OCM) and collaborative clustering module (CCM), respectively.



Fig. 2. Illustration of our CoDT network. It consists of two branches (left and right), each of which contains a base module (BM), and an online clustering module (OCM). The two branches are connected by a collaborative clustering module (CCM).

3.3 Base Module

BM-B₀ There are two feature encoders dubbed F_0 and \hat{F}_0 . The \hat{F}_0 is updated by an exponential moving average (EMA) of F_0 and used by the subsequent OCM. The encoder F_0 first embeds the input skeleton sequences into features. The features of source samples are then fed into a source classifier C_0^s and optimized by cross-entropy loss with annotated labels

$$\mathcal{L}_{sup} = \frac{1}{n^s} \sum_{i=1}^{n^s} \text{CE}\left(C_0^s(F_0(\boldsymbol{x}_i^s)), y_i^s\right),$$
(1)

where n^s is the mini-batch size of source data, and CE is the short of Cross Entropy loss. The decoder D_0 reconstructs the input skeletons from the features for both source data and target data. The reconstruction loss is computed as

$$\mathcal{L}_{dec} = \frac{1}{n^s} \sum_{i=1}^{n^s} \text{MSE}\left(D_0(F_0(\boldsymbol{x}_i^s)), \boldsymbol{x}_i^s\right) + \frac{1}{n^t} \sum_{i=1}^{n^t} \text{MSE}\left(D_0(F_0(\boldsymbol{x}_i^t)), \boldsymbol{x}_i^t\right), \quad (2)$$

where n^t denotes the mini-batch size of target data, and MSE denotes Mean Square Error loss. The reconstruction enforces the representations to retain generic and meaningful human structure, which can strengthen the domain invariance of representations [22]. Note that, based on the principle of 'domainsharing', when the source data and target data have different joints, we only keep their shared joints in \mathcal{B}_0 .

BM-\mathcal{B}_1 To learn the target-specific representations from the target set, we employ the popular contrastive learning based on instance discrimination [28],

where the features of the different augments of a sample are pulled together and the features of different samples are pushed apart.

Specifically, similar to BM- \mathcal{B}_0 , BM- \mathcal{B}_1 contains a feature encoder F_1 and an EMA encoder \hat{F}_1 . For a target sample, we first transform it to two augments $(\boldsymbol{x}_i^t, \hat{\boldsymbol{x}}_i^t)$ by data augmentation, and then pass them into F_1 and \hat{F}_1 , respectively. Here we denote the outputs as $\boldsymbol{z}_i^t = F_1(\boldsymbol{x}_i^t)$, $\hat{\boldsymbol{z}}_i^t = \hat{F}_1(\hat{\boldsymbol{x}}_i^t)$. Besides, to enlarge the number of negative samples, following [28,42], we maintain a memory bank $\boldsymbol{M}^t = \{\hat{\boldsymbol{z}}_m^t\}|_m^{N^t}$ to store the features from \hat{F}_1 of all target samples. In each iteration, the \boldsymbol{M}^t is updated by the $\hat{\boldsymbol{z}}_i^t$ in current mini-batch, and then we compute the contrastive loss as

$$\mathcal{L}_{cont} = \frac{1}{n^t} \sum_{i=1}^{n^t} -\log \frac{e^{\rho \cdot \cos(\boldsymbol{z}_i^t, \hat{\boldsymbol{z}}_i^t)}}{\sum_{m=1}^{N^t} e^{\rho \cdot \cos(\boldsymbol{z}_i^t, \hat{\boldsymbol{z}}_m^t)}},\tag{3}$$

where $\cos(\cdot)$ denotes the cosine similarity function and ρ is the temperature.

3.4 Online Clustering Module

Pipeline of OCM In both \mathcal{B}_0 and \mathcal{B}_1 , we employ an OCM to guide the feature clustering via pseudo labels. The OCM is based on the student-teacher framework [39,64,71], where a teacher model (e.g., \hat{F}_0 , \hat{F}_1) is updated by EMA of a student model (e.g., F_0 , F_1). To learn the optimal cluster assignment, we add a target classifier to both student and teacher. The classifiers are denoted by C_0^t and \hat{C}_0^t in \mathcal{B}_0 , and C_1^t and \hat{C}_1^t in \mathcal{B}_1 . Following image clustering [70,34,66], we set the numbers of categories of the classifiers to the number of ground-truth categories for the purpose of evaluation¹. Given a target sample \boldsymbol{x}_i^t , inspired by [60], a strong augmentation operator A and a weak augmentation operator a are applied to transform \boldsymbol{x}_i^t into $A(\boldsymbol{x}_i^t)$ and $a(\boldsymbol{x}_i^t)$, respectively. The teacher generates the pseudo label based on its prediction on the $a(\boldsymbol{x}_i^t)$. Then the student is trained to predict the pseudo label by feeding $A(\boldsymbol{x}_i^t)$. These practices can enhance the invariance of representations to varying degrees of transformations.

Pseudo Label Generation Before elaborating on the details, we first review the objective of a classification task. Specifically, given a dataset with N instances $\{\boldsymbol{x}_n\}|_{n=1}^N$ drawn from K classes, the task is to maximize the mutual information between labels and input data through minimizing the Kullback–Leibler divergence $D(\boldsymbol{Q}||\boldsymbol{P})$ between the model's predictions \boldsymbol{P} and labels \boldsymbol{Q} :

$$D(\mathbf{Q}||\mathbf{P}) = \frac{1}{N} \sum_{n=1}^{N} \sum_{y=1}^{K} Q_{yn} \log \frac{Q_{yn}}{P_{yn}} = \frac{1}{N} \sum_{n=1}^{N} \sum_{y=1}^{K} \left(\underbrace{Q_{yn} \log Q_{yn}}_{-H(\mathbf{Q})} \underbrace{-Q_{yn} \log P_{yn}}_{E(\mathbf{Q}||\mathbf{P})} \right),$$
(4)

¹ In fact, as proven in [66], even if the exact number of ground-truth categories is unknown, we can overcluster to a larger amount of clusters.

where Q_{yn} is the (y, n) element of $\mathbf{Q} \in \mathbb{R}^{K \times N}$ denoting the label probability of \mathbf{x}_n being assigned to the label $y \in \{1, \ldots, K\}$. P_{yn} is the element of $\mathbf{P} \in \mathbb{R}^{K \times N}$ that denotes the probability predicted by the model. $D(\mathbf{Q}||\mathbf{P})$ can be split into a cross-entropy term $E(\mathbf{Q}||\mathbf{P})$ and an entropy term $H(\mathbf{Q})$. Unlike the supervised task where \mathbf{Q} is deterministic, labels are unavailable in our case, and we have to minimize $D(\mathbf{Q}||\mathbf{P})$ w.r.t. both the \mathbf{Q} and \mathbf{P} . Further, we extend it to the student-teacher framework. Formally, it is to alternate the following steps

$$\begin{cases} \hat{\boldsymbol{Q}} \leftarrow \arg\min_{\hat{\boldsymbol{Q}}} D(\hat{\boldsymbol{Q}} || \hat{\boldsymbol{P}}), \qquad (5) \end{cases}$$

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \epsilon \frac{\partial \left(D(\hat{\boldsymbol{Q}} || \boldsymbol{P}) + \mathcal{L} \right)}{\partial \boldsymbol{\theta}}, \tag{6}$$

$$\left(\hat{\boldsymbol{\theta}} \leftarrow \alpha \hat{\boldsymbol{\theta}} + (1 - \alpha)\boldsymbol{\theta},\right)$$
(7)

where \boldsymbol{P} and $\boldsymbol{\theta}$ are the predictions and parameters of student. $\hat{\boldsymbol{P}}$ and $\hat{\boldsymbol{\theta}}$ correspond to teacher. $\hat{\boldsymbol{Q}}$ represents the pseudo labels. First, $\hat{\boldsymbol{Q}}$ is calculated by Eq. (5). Then, we fix $\hat{\boldsymbol{Q}}$ and optimize the predictions \boldsymbol{P} of student by minimizing $D(\hat{\boldsymbol{Q}}||\boldsymbol{P})$ along with other auxiliary loss \mathcal{L} (*e.g.*, \mathcal{L}_{sup} , \mathcal{L}_{cont}) on a mini-batch via Eq. (6), where ϵ is the learning rate of gradient descent. Finally, the teacher is updated by an EMA of the student in Eq. (7), where α is the decay rate.

Let's first take a look at Eq. (5). Since we only consider the one-hot pseudo labels, *i.e.*, $\hat{Q}_{yn} \in \{0, 1\}$, then $H(\hat{Q}) \equiv 0$ and $D(\hat{Q}||\hat{P}) \equiv E(\hat{Q}||\hat{P})$. Thus we can obtain the solution to Eq. (5) by taking the index that has the maximum value in the prediction of teacher, *i.e.*,

$$\hat{Q}_{yn}^* = \delta \big(y - \arg\max_k \hat{P}_{kn} \big), \tag{8}$$

where $\delta(\cdot)$ is the Dirac delta function with $\delta(\cdot) = 0$ except $\delta(0) = 1$. The Eq. (8) is similar in spirit to the semi-supervised learning method [60]. Obviously, if $\mathcal{L} = 0$, there is a shortcut to minimize both $D(\mathbf{Q}||\mathbf{P})$ and $D(\mathbf{Q}||\mathbf{P})$ by assigning all samples to a single label, which is known as clustering degeneration. This phenomenon can be avoided in [60], since the supervised training on a few labeled samples can regularize the model via \mathcal{L} . However, in our case, the \mathcal{L} is either the supervised loss of source dataset with a distribution different from that of target dataset, or the contrastive loss for instance discrimination rather than clustering. In other words, there are discrepancies between the objectives of \mathcal{L} and the ultimate goal (*i.e.*, target sample clustering). It renders existing \mathcal{L} unable to prevent clustering degeneration, as proven in Sec. 4.3. To avoid clustering degeneration, we propose to constrain the distribution of clusters to be as uniform as possible, *i.e.*, all clusters contain the same number of samples. Achieving it within Eq. (6) in an end-to-end fashion is difficult as the constraint involves the global distribution of the entire dataset. Alternatively, we propose to balance the pseudo labels generated by Eq. (5) without gradient descent. Formally, we restrict \hat{Q} to be an element of transportation polytope [14]

$$\mathcal{U} = \{ \hat{\boldsymbol{Q}} \in \mathbb{R}^{K \times N}_{+} | \hat{\boldsymbol{Q}}^{\top} \boldsymbol{1}_{K} = \boldsymbol{1}_{N}, \hat{\boldsymbol{Q}} \boldsymbol{1}_{N} = \frac{N}{K} \boldsymbol{1}_{K} \},$$
(9)

where $\mathbf{1}_K$ denotes the vector of ones with dimension K. The Eq. (9) indicates that each class has equal number (*i.e.*, $\frac{N}{K}$) of samples. Inspired by [1,8], solving Eq. (5) subject to Eq. (9) can transformed to an optimal transport problem for mapping the N data points to the K centers, whose solution is

$$\hat{\boldsymbol{Q}}^* = N \operatorname{diag}(\mathbf{u}) (\hat{\boldsymbol{P}}/N)^{\xi} \operatorname{diag}(\mathbf{v}), \qquad (10)$$

where ξ is a pre-defined scalar, and $\mathbf{u} \in \mathbb{R}^{K}$, $\mathbf{v} \in \mathbb{R}^{N}$ are two vectors computed by the fast version of Sinkhorn-Knopp algorithm [14]. Please refer to Supplementary for more details.

Since it is costly to calculate \hat{P} from scratch each iteration, we maintain the same memory bank as that used in contrastive learning, and the extra cost of computing predictions from the bank is affordable.

Objective of OCM Thus, for a target instance \boldsymbol{x}_i^t , we can obtain its pseudo labels denoted by $\hat{y}_{0,i}^t, \hat{y}_{0,i}^t \in \{1, \ldots, K\}$ in \mathcal{B}_0 and \mathcal{B}_1 , respectively. Then we use the pseudo labels to train the corresponding student models by minimizing following loss function

$$\mathcal{L}_{ocm} = \frac{1}{n^t} \sum_{i=1}^{n^t} \left(\operatorname{CE}(\hat{C}_0^t(\hat{F}_0(A(\boldsymbol{x}_i^t))), \hat{y}_{0,i}^t) + \operatorname{CE}(\hat{C}_1^t(\hat{F}_1(A(\boldsymbol{x}_i^t))), \hat{y}_{1,i}^t) \right), \quad (11)$$

3.5 Collaborative Clustering Module

To make use of complementarity between \mathcal{B}_0 and \mathcal{B}_1 , we propose the collaborative clustering module (CCM) to allow co-training between two branches. Since the two branches have different parameters and regularized by different losses (*e.g.*, \mathcal{L}_{sup} for \mathcal{B}_0 , \mathcal{L}_{cont} for \mathcal{B}_1), the classes of the target classifiers of \mathcal{B}_0 and \mathcal{B}_1 are difficult to be matched exactly. Hence, the pseudo labels from one model cannot be used directly to train the other model. To overcome this challenge, we present a conceptually simple but practically effective co-training method. Instead of matching clusters between models, we propose to exchange the pair-wise relationships of samples across branches, which are explicitly matched between two branches and meanwhile can determine the performance of clustering.

Specifically, for two target instances $\boldsymbol{x}_{i}^{t}, \boldsymbol{x}_{j}^{t}$, we first define the binary pairwise pseudo label as $\mathcal{G}_{0,ij}^{t} = \delta(\hat{y}_{0,i}^{t} - \hat{y}_{0,j}^{t})$ in \mathcal{B}_{0} and $\mathcal{G}_{1,ij}^{t} = \delta(\hat{y}_{1,i}^{t} - \hat{y}_{1,j}^{t})$ in \mathcal{B}_{1} , indicating that whether $\boldsymbol{x}_{i}^{t}, \boldsymbol{x}_{j}^{t}$ are from the same cluster (positive pair) or from different clusters (negative pair) in each branch. Besides, we define the pairwise similarity of a pair of samples as the inner product of their predictions, *i.e.*, $\mathcal{P}_{0,ij}^{t} = \boldsymbol{p}_{0,i}^{t} \boldsymbol{p}_{0,j}^{t}$ and $\mathcal{P}_{1,ij}^{t} = \boldsymbol{p}_{1,i}^{t} \boldsymbol{p}_{1,j}^{t}$, where $\boldsymbol{p}_{0,i}^{t} = \operatorname{softmax}(\hat{C}_{0}^{t}(\hat{F}_{0}(\boldsymbol{A}(\boldsymbol{x}_{i}^{t})))) \in \mathbb{R}^{K}$ (likewise for $\boldsymbol{p}_{0,j}^{t}, \boldsymbol{p}_{1,i}^{t}, \boldsymbol{p}_{1,j}^{t}$). For co-training, we use $\mathcal{G}_{0,ij}^{t}$ and $\mathcal{G}_{1,ij}^{t}$ as the supervisions to optimize $\mathcal{P}_{1,ij}^{t}$ and $\mathcal{P}_{0,ij}^{t}$, respectively, aiming to make the similarities of positive/negative pairs increase/diminish. To achieve it, we design the

objective function similar in spirit to the supervised contrastive loss [36]:

$$\mathcal{L}_{ccm} = -\frac{1}{n^t} \sum_{i=1}^{n^t} \Big(\frac{\sum_{j=1}^{n^t} \mathcal{G}_{0,ij}^t \log \overline{\mathcal{P}_{i,ij}^t}}{\sum_{j=1}^{n^t} \mathcal{G}_{0,ij}^t} + \frac{\sum_{j=1}^{n^t} \mathcal{G}_{1,ij}^t \log \overline{\mathcal{P}_{0,ij}^t}}{\sum_{j=1}^{n^t} \mathcal{G}_{1,ij}^t} \Big), \quad (12)$$

where $\overline{\mathcal{P}^{t}}_{0,ij}$ and $\overline{\mathcal{P}^{t}}_{1,ij}$ are defined as $\overline{\mathcal{P}^{t}}_{0,ij} = \frac{\mathcal{P}^{t}_{0,ij}}{\sum_{j=1}^{n^{t}} \mathcal{P}^{t}_{0,ij}}$, $\overline{\mathcal{P}^{t}}_{1,ij} = \frac{\mathcal{P}^{t}_{1,ij}}{\sum_{j=1}^{n^{t}} \mathcal{P}^{t}_{1,ij}}$. Note that, to maximize the similarity (*i.e.*, the inner product of the predictions) of a positive pair, both predictions need to be one-hot and assigned to the same cluster [66]. This property enforces the two branches to be consistent in the pair-wise relationships of cluster assignments, yielding consistent clustering performance on the two branches.

Discussion Here we discuss the vital differences between our proposed cotraining method and some related methods. [8] swaps instance-wise labels of two views of the same image, forcing the (single) model to produce consistent predictions for different views. In contrast, CoDT exchanges pair-wise labels across models for co-training to utilize the complementarity of domain-shared and target-specific features. Recently, [84] proposes a mutual knowledge distillation algorithm across two different branches by comparing the feature similarity distribution between each instance and a queue of features. Differently, CoDT conducts co-training on the predictions of classifiers, leading to a direct optimization on the cluster assignments.

Apart from the idea of co-training, the objective function \mathcal{L}_{ccm} , to our knowledge, is also different from existing methods. To name a few, the differences between [36] and ours include: 1) [36] applies the loss on features, while our \mathcal{L}_{ccm} is applied on predictions; 2) \mathcal{L}_{ccm} has a simpler formula without any hyperparameter (*e.g.*, temperature). Recently, [43] proposes to adopt the instance-based contrastive loss on the predictions, showing promising performance in close-set domain adaption. Yet, [43] handles the close-set tasks with known classes, and its loss is used for instance discrimination, which is evidently different from ours. \mathcal{L}_{ccm} is also different from [66,27,42] where a pair of samples is positive only when the similarity between their features meets a heuristic condition (*e.g.*, is larger than a threshold). Note that, none of the above competitors involve co-training.

3.6 Training and Test

In training phase, we first pretrain $BM-\mathcal{B}_0$ and $BM-\mathcal{B}_1$ with the loss

$$\mathcal{L}_{base} = \lambda_{sup} \mathcal{L}_{sup} + \lambda_{dec} \mathcal{L}_{dec} + \lambda_{cont} \mathcal{L}_{cont}, \tag{13}$$

where λ_* denotes the loss weight. After obtaining a good initialization, OCM and CCM are successively included for finetuning. The overall loss is

$$\mathcal{L}_{all} = \mathcal{L}_{base} + \lambda_{ocm} \mathcal{L}_{ocm} + \lambda_{ccm} \mathcal{L}_{ccm}.$$
 (14)

In test phase, following [44,32], we use the student models for testing. Specifically, if there is a target classifier in the model, we use its predictions as cluster assignments, otherwise, we use the spherical k-means [29] to cluster the features.

4 Experiment

4.1 Datasets and Metrics

NTU-60 [58] is a 3D action recognition dataset with 60 classes. The skeletons are shot by Microsoft Kinect V2 [18]. NTU-120 [48] is an extended version of NTU-60, containing 120 classes. Here we only take the classes of NTU-120 that are not overlapped with those of NTU-60, denoted by NTU-60+. PKUMMD [47] is a dataset for temporal action detection, where the trimmed action instances have been used for action recognition [45]. The skeletons are also collected by Kinect V2. There are 51 classes different from that of NTU-60+. Different from the above in-the-lab datasets, *Skeletics* [25] is a carefully curated dataset sourced from real-world videos [9]. The 3D poses are estimated by the pose estimation method VIBE [37]. Three cross-domain tasks are chosen to imitate various situations: NTU-60 \rightarrow Skeletics (xview), Skeletics \rightarrow PKUMMD (xview), and $NTU-60+ \rightarrow PKUMMD$ (xsub). Let's take $NTU-60+ \rightarrow PKUMMD$ (xsub) as an example. The NTU-60+ and PKUMMD are the source and target datasets, respectively, and 'xsub' is the evaluation protocol. The former two tasks are more challenging than the last one because the former two tasks simulate the transfer learning between an in-the-lab dataset shot by 3D sensors and a realworld dataset estimated by pose estimation algorithm, while in the last task, the two datasets are shot by similar device, and are therefore more similar in style. In the rest of the paper, we abbreviate the three tasks as $N \to S, S \to P$, and $N \rightarrow P$. For evaluation, we use three widely-used clustering performance metrics [30], including Accuracy (ACC), Normalised Mutual Information (NMI), Adjusted Rand Index (ARI). Due to the space constraint, we refer readers to Supplementary for more details about datasets, metrics, implementation details.

4.2 Comparison with Different Baselines

To verify the effectiveness of CoDT, we first compare CoDT with related baselines in Table 1. Here the results of both branches are reported. For fairness, we use the pretrained weights of our base modules to initialize the weights of other baselines. The 'BM- \mathcal{B}_0 w/o D_0 ' means that the decoder D_0 is not used. Compared to 'BM- \mathcal{B}_0 ', we can see that the decoder improves the performances in most cases. We further evaluate the performance when we combine 'BM- \mathcal{B}_0 ' and 'BM- \mathcal{B}_1 ' via multi-task learning, denoted by 'BM- \mathcal{B}_0 + BM- \mathcal{B}_1 '. We can see that it outperforms both 'BM- \mathcal{B}_0 ' and 'BM- \mathcal{B}_1 ', demonstrating the complementarity of 'BM- \mathcal{B}_0 ' and 'BM- \mathcal{B}_1 '. Thereafter, we combine BM- \mathcal{B}_0 with other advanced methods of unsupervised skeleton/image learning [44,1]. We can see that the performances are improved. However, their performance is still far behind that of CoDT, demonstrating that multi-task framework is suboptimal for combining 'BM- \mathcal{B}_0 ' and 'BM- \mathcal{B}_1 '. What's more, we re-implement some representative methods of other cross-domain tasks, including UDA [13], NCD [19], CD-ReID [21] and CD-FSL [56]. These methods commonly share the same feature extractor among different domains. It can be seen that our CoDT significantly

Methoda	$N \rightarrow S$			$S \rightarrow P$			$N+ \rightarrow P$		
Methous	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI
$BM-\mathcal{B}_0^{\dagger}$	17.9	22.5	6.1	40.0	58.9	30.2	54.8	73.7	44.3
BM- $\mathcal{B}_0 w/o D_0^{\dagger}$	16.6	20.5	6.1	38.5	55.6	27.6	53.3	73.4	43.4
$BM-\mathcal{B}_1^{\dagger}$	18.7	23.3	6.4	43.3	58.4	32.1	42.7	57.4	32.3
$BM-\mathcal{B}_0 + BM-\mathcal{B}_1^{\dagger}$	19.4	22.8	6.4	47.7	62.1	36.9	58.3	74.0	47.7
$BM-\mathcal{B}_0 + CrossCLR [44]^{\dagger}$	20.5	24.1	7.0	49.1	63.8	36.6	60.4	75.2	47.9
$BM-\mathcal{B}_0 + Asano. et.al. [1]$	21.2	26.1	8.4	51.4	65.4	38.4	62.3	74.4	49.4
GVB [13]†	19.3	22.2	6.0	37.8	56.9	28.4	59.5	75.6	50.0
STARTUP [56]†	19.0	22.1	5.9	48.5	63.1	38.1	59.1	72.5	47.2
MMT [21]	20.8	25.5	7.6	52.4	67.3	41.2	65.4	76.2	55.6
UNO [19]	22.5	26.4	9.1	54.1	70.0	43.2	66.8	76.9	56.7
$CoDT-\mathcal{B}_0$	25.0	28.0	10.7	59.5	74.1	50.0	68.2	78.5	58.8
$CoDT-\mathcal{B}_1$	25.4	${\bf 28.7}$	11.3	59.4	73.9	50.2	67.8	78.1	58.1

Table 1. Comparison between different methods. '†' indicates that the spherical k-means is used for clustering in these methods.

outperforms all of them. CoDT achieves salient performances since it can 1) disentangle the learning of domain-shared and target-specific features via a twobranch framework to fully exploit their respective characteristics; 2) coordinate the clustering of domain-shared and target-specific features to fully utilize their complementarity. Especially, these two properties are more advantageous when the domain gaps are larger, since the domain-shared and target-specific features are more different and complementary. It makes our CoDT extremely superior in the former two tasks with large domain gaps.

4.3 Ablation Study

Effectiveness of Proposed Components We conduct ablation studies in Table 2 to investigate the effectiveness of the proposed components, *i.e.*, OCM and CCM. After introducing OCM, the performances on all branches and tasks are substantially improved when compared to the base modules. This is because the base modules are only trained to learn good representations and an offline clustering criterion (*e.g.*, k-means) is still needed for clustering, whereas OCM is capable of optimizing both feature learning and clustering simultaneously. By taking account of CCM, \mathcal{B}_0 and \mathcal{B}_1 perform very closely. It is because CCM can enforce the consistency between \mathcal{B}_0 and \mathcal{B}_1 in terms of pair-wise relationships on cluster assignments. More importantly, after using CCM, the performances are greatly improved, verifying that CCM can effectively leverage the complementarity of domain-shared features and target-specific features. In Fig. 3, we show the evolution of the learned representations of target samples on 'N+ \rightarrow P'. It is shown that while the clusters overlap in the beginning, they become more and more separated as the OCM and CCM are successively involved.

Table 2. Effect of different components.

Mathada	ĸ	$N \rightarrow S$			$S \rightarrow P$			$N+ \rightarrow P$		
methous	\mathcal{D}_*	ACC	NMI	ARI	ACC	NMI	ARI	ACC	NMI	ARI
BM	\mathcal{B}_0	17.9	22.5	6.1	40.0	58.9	30.2	54.8	73.7	44.3
DW	\mathcal{B}_1	18.7	23.3	6.4	43.3	58.4	32.1	42.7 5	57.4	32.3
DMLOCM	\mathcal{B}_0	22.0	26.8	9.1	55.3	70.3	45.1	63.7	80.1	53.9
DM+OOM	\mathcal{B}_1	21.8	26.4	8.7	54.7	67.5	44.3	52.6	66.1	41.9
PM + OCM + CCM	\mathcal{B}_0	25.0	28.0	10.7	59.5	74.1	50.0	68.1	78.5	58.8
DM+00M+00M	\mathcal{B}_1	25.4	28.7	11.3	59.4	73.9	50.2	67.8 7	78.1	58.1



Fig. 3. t-SNE visualization of features learned by different methods. Different colors represent different clusters.

Analysis of OCM In Table 3, we give the detailed analysis of OCM, where the models with and without CCM are evaluated. Without loss of generalization, we take the task of 'S \rightarrow P' as an example. Apart from 'ACC', we also report the metric of 'Uniformity' which is measured by the entropy of clustering results. The more uniform the distribution of clusters, the greater its value. It takes the maximum value of 3.93 when all clusters are equal in size. The value corresponding to the ground-truth labels is 3.88, indicating that the ground-truth distribution is not exactly uniform. The \mathcal{L} in Table 3 denotes the loss regularization in Eq. (6), including \mathcal{L}_{sup} , \mathcal{L}_{dec} , and \mathcal{L}_{con} . The \mathcal{U} represents the uniform constraint of Eq. (9). Note that, the pseudo labels are generated by Eq. (10) when \mathcal{U} is used, and by Eq. (8) otherwise. From Table 3, we have following observations: 1) When neither \mathcal{L} nor \mathcal{U} is used, the 'Uniformity' is very small, indicating that the models encounter catastrophic clustering degeneration. 2) When \mathcal{L} is used, the degeneration is alleviated and the performances increase, demonstrating that \mathcal{L} is helpful to generate meaningful groups for the target data. Besides, we observe that in the row of 'BM+OCM', the 'Uniformity' of \mathcal{B}_1 is larger than that of \mathcal{B}_0 . The reason may be that the features learned by the contrastive learning have the property of uniformity in instance-level [67]. Nevertheless, the clusters are still not balanced enough, indicating that using \mathcal{L} alone is not enough. 3) When \mathcal{U} is further imposed, the clusters become more balanced, resulting in better performances. These results well prove the effectiveness of our design.

Table 3. Analysis of OCM, where the numbers before and after the '/' denote the results of \mathcal{B}_0 and \mathcal{B}_1 , respectively.

Methods	\mathcal{L}	\mathcal{U}	ACC(%)	Uniformity	-	'/' denote the	e ACC of <i>l</i>	\mathcal{B}_0 and \mathcal{B}_1 .
	X	X	9.0/6.9	0.70/1.09	-			
BM+OCM	1	X	31.2/48.4	2.76/3.68	-	Methods	$N \rightarrow S$	$\mathrm{N}+ \to \mathrm{P}$
	1	1	55.7/54.7	3.91/3.93	-	$\operatorname{CCM-}FP$	22.3/24.2	63.9/55.3
BM+OCM+CCM	X	X	9.1/8.9	0.80/0.80	-	CCM-PF	22.8/24.7	64.7/55.9
	1	X	48.4/48.7	3.48/3.47	-	CCM (ours)	25.0/25.4	68.1/67.8
	1	1	59.5/59.4	3.91/3.93	-			

 Table 4. Comparison between

our CCM with its variants, where

the numbers before and after the

Analysis of CCM The CCM can be divided into two stages. The first is to construct pair-wise pseudo labels by comparing the pseudo labels of sample pairs in each branch, and the second is to optimize the pair-wise similarities of predictions using the pair-wise labels of the other branch. To verify the necessity of our design at each stage, we device two variants of CCM dubbed CCM-PF and CCM-FP, which are extended from the related methods discussed in Sec. 3.5. The CCM-PF modifies the first stage of CCM, where the pair-wise pseudo labels are constructed based on the pair-wise similarities of features (rather than the pair-wise comparisons of pseudo labels), following [27]. The CCM-PF changes the second stage of CCM, where the pair-wise similarities of *features* (rather than predictions) are optimized, following [36]. Please refer to the Supplementary for details about the variants. Their performances are shown in Table 4. We can see that neither CCM-*PF* nor CCM-*FP* can ensure performance consistency between two branches, showing that in these two variants, the cluster assignments are not fully inherited across branches. Due to such reason, their performances are also inferior to ours, demonstrating the effectiveness of our design.

5 Conclusion and Future Work

In this paper, we propose a novel method dubbed CoDT for CD-SAR. The main idea of CoDT is to leverage the complementarity of domain-shared and target-specific features. To this end, we introduce the OCM to obtain robust pseudo labels to guide feature clustering, and the CCM to collaborate the two kinds of features. The experimental results show that our method significantly outperforms other cross-domain training or unsupervised learning methods. In principle, our method can be adopted to other cross-domain tasks in an open-set setting (*e.g.*, CD-FSL [56], CD-ReID [68]). We leave this as our future work.

Acknowledge This work is supported by the National Natural Science Foundation of China under Grant No.62176246 and No.61836008.

References

- 1. Asano, Y.M., Rupprecht, C., Vedaldi, A.: Self-labelling via simultaneous clustering and representation learning. In: ICLR (2020)
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W.: A theory of learning from different domains. Machine learning 79(1), 151–175 (2010)
- 3. Ben-David, S., Blitzer, J., Crammer, K., Pereira, F., et al.: Analysis of representations for domain adaptation. In: NIPS (2007)
- Blum, A., Mitchell, T.: Combining labeled and unlabeled data with co-training. In: COLT. pp. 92–100 (1998)
- 5. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: CVPR (2017)
- Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: ECCV (2018)
- Caron, M., Bojanowski, P., Mairal, J., Joulin, A.: Unsupervised pre-training of image features on non-curated data. In: ICCV (2019)
- 8. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. In: NIPS (2020)
- 9. Carreira, J., Noland, E., Hillier, C., Zisserman, A.: A short note on the kinetics-700 human action dataset. arXiv preprint arXiv:1907.06987 (2019)
- Chang, J., Wang, L., Meng, G., Xiang, S., Pan, C.: Deep adaptive image clustering. In: ICCV (2017)
- 11. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML (2020)
- Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297 (2020)
- Cui, S., Wang, S., Zhuo, J., Su, C., Huang, Q., Tian, Q.: Gradually vanishing bridge for adversarial domain adaptation. In: CVPR (2020)
- 14. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. In: NIPS (2013)
- Deng, W., Zheng, L., Ye, Q., Kang, G., Yang, Y., Jiao, J.: Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person reidentification. In: CVPR (2018)
- Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: ICCV (2015)
- Fan, H., Zheng, L., Yan, C., Yang, Y.: Unsupervised person re-identification: Clustering and fine-tuning. ACM Transactions on Multimedia Computing, Communications, and Applications 14(4), 1–18 (2018)
- Fankhauser, P., Bloesch, M., Rodriguez, D., Kaestner, R., Hutter, M., Siegwart, R.: Kinect v2 for mobile robot navigation: Evaluation and modeling. In: ICAR (2015)
- 19. Fini, E., Sangineto, E., Lathuilière, S., Zhong, Z., Nabi, M., Ricci, E.: A unified objective for novel class discovery. In: ICCV (2021)
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. The journal of machine learning research 17(1), 2096–2030 (2016)
- 21. Ge, Y., Chen, D., Li, H.: Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. In: ICLR (2019)
- Ghifary, M., Kleijn, W.B., Zhang, M., Balduzzi, D.: Domain generalization for object recognition with multi-task autoencoders. In: ICCV (2015)

- 16 Q. Liu, Z. Wang
- Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Pires, B., Guo, Z., Azar, M., et al.: Bootstrap your own latent: A new approach to self-supervised learning. In: NIPS (2020)
- Guo, Y., Codella, N.C., Karlinsky, L., Codella, J.V., Smith, J.R., Saenko, K., Rosing, T., Feris, R.: A broader study of cross-domain few-shot learning. In: ECCV (2020)
- Gupta, P., Thatipelli, A., Aggarwal, A., Maheshwari, S., Trivedi, N., Das, S., Sarvadevabhatla, R.K.: Quo vadis, skeleton action recognition? International Journal of Computer Vision 129(7), 2097–2112 (2021)
- 26. Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I.W., Sugiyama, M.: Co-teaching: Robust training of deep neural networks with extremely noisy labels. In: NIPS (2018)
- Han, K., Rebuffi, S.A., Ehrhardt, S., Vedaldi, A., Zisserman, A.: Automatically discovering and learning new visual categories with ranking statistics. In: ICLR (2019)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020)
- Hornik, K., Feinerer, I., Kober, M., Buchta, C.: Spherical k-means clustering. Journal of statistical software 50, 1–22 (2012)
- Huang, J., Gong, S., Zhu, X.: Deep semantic clustering by partition confidence maximisation. In: CVPR (2020)
- Huang, Y., Peng, P., Jin, Y., Xing, J., Lang, C., Feng, S.: Domain adaptive attention model for unsupervised cross-domain person re-identification. In: AAAI (2019)
- Islam, A., Chen, C.F., Panda, R., Karlinsky, L., Feris, R., Radke, R.J.: Dynamic distillation network for cross-domain few-shot recognition with unlabeled data. In: NIPS (2021)
- 33. Islam, A., Chen, C.F., Panda, R., Karlinsky, L., Radke, R., Feris, R.: A broad study on the transferability of visual representations with contrastive learning. In: ICCV (2021)
- 34. Ji, X., Henriques, J.F., Vedaldi, A.: Invariant information clustering for unsupervised image classification and segmentation. In: ICCV (2019)
- Ke, Q., Bennamoun, M., An, S., Sohel, F., Boussaid, F.: A new representation of skeleton sequences for 3d action recognition. In: CVPR (2017)
- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. In: NIPS (2020)
- 37. Kocabas, M., Athanasiou, N., Black, M.J.: Vibe: Video inference for human body pose and shape estimation. In: CVPR (2020)
- Kundu, J.N., Gor, M., Uppala, P.K., Radhakrishnan, V.B.: Unsupervised feature learning of human actions as trajectories in pose embedding manifold. In: WACV (2019)
- Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: ICLR (2016)
- 40. Li, B., Li, X., Zhang, Z., Wu, F.: Spatio-temporal graph routing for skeleton-based action recognition. In: AAAI (2019)
- Li, C., Zhong, Q., Xie, D., Pu, S.: Skeleton-based action recognition with convolutional neural networks. In: 2017 IEEE International Conference on Multimedia & Expo Workshops. pp. 597–600. IEEE (2017)
- 42. Li, J., Li, G., Shi, Y., Yu, Y.: Cross-domain adaptive clustering for semi-supervised domain adaptation. In: CVPR (2021)

- Li, J., Zhang, Y., Wang, Z., Tu, K.: Semantic-aware representation learning via probability contrastive loss. arXiv preprint arXiv:2111.06021 (2021)
- 44. Li, L., Wang, M., Ni, B., Wang, H., Yang, J., Zhang, W.: 3d human action representation learning via cross-view consistency pursuit. In: CVPR (2021)
- Lin, L., Song, S., Yang, W., Liu, J.: Ms2l: Multi-task self-supervised learning for skeleton based action recognition. In: ACMMM (2020)
- 46. Lin, S., Li, H., Li, C.T., Kot, A.C.: Multi-task mid-level feature alignment network for unsupervised cross-dataset person re-identification. In: BMVC (2018)
- 47. Liu, C., Hu, Y., Li, Y., Song, S., Liu, J.: Pku-mmd: A large scale benchmark for continuous multi-modal human action understanding. arXiv preprint arXiv:1703.07475 (2017)
- 48. Liu, J., Shahroudy, A., Perez, M., Wang, G., Duan, L.Y., Kot, A.C.: Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding. IEEE transactions on pattern analysis and machine intelligence 42(10), 2684–2701 (2019)
- 49. Liu, X., Zhang, S.: Domain adaptive person re-identification via coupling optimization. In: ACMMM (2020)
- 50. Liu, Z., Zhang, H., Chen, Z., Wang, Z., Ouyang, W.: Disentangling and unifying graph convolutions for skeleton-based action recognition. In: CVPR (2020)
- 51. Mekhazni, D., Bhuiyan, A., Ekladious, G., Granger, E.: Unsupervised domain adaptation in the dissimilarity space for person re-identification. In: ECCV (2020)
- 52. Misra, I., Maaten, L.v.d.: Self-supervised learning of pretext-invariant representations. In: CVPR (2020)
- 53. Nie, Q., Liu, Z., Liu, Y.: Unsupervised 3d human pose representation with viewpoint and pose disentanglement. In: ECCV (2020)
- 54. Park, S., Han, S., Kim, S., Kim, D., Park, S., Hong, S., Cha, M.: Improving unsupervised image clustering with robust learning. In: CVPR (2021)
- 55. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: CVPR (2016)
- 56. Phoo, C.P., Hariharan, B.: Self-training for few-shot transfer across extreme task differences. In: ICLR (2020)
- 57. Qiao, S., Shen, W., Zhang, Z., Wang, B., Yuille, A.: Deep co-training for semisupervised image recognition. In: ECCV (2018)
- 58. Shahroudy, A., Liu, J., Ng, T.T., Wang, G.: Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In: CVPR (2016)
- 59. Shi, L., Zhang, Y., Cheng, J., Lu, H.: Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In: CVPR (2019)
- 60. Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In: NIPS (2020)
- Song, S., Lan, C., Xing, J., Zeng, W., Liu, J.: Spatio-temporal attention-based lstm networks for 3d action recognition and detection. IEEE Transactions on image processing 27(7), 3459–3471 (2018)
- 62. Su, K., Liu, X., Shlizerman, E.: Predict & cluster: Unsupervised skeleton based action recognition. In: CVPR (2020)
- 63. Su, Y., Lin, G., Wu, Q.: Self-supervised 3d skeleton action representation learning with motion consistency and continuity. In: ICCV (2021)
- 64. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: NIPS (2017)
- Tseng, H.Y., Lee, H.Y., Huang, J.B., Yang, M.H.: Cross-domain few-shot classification via learned feature-wise transformation. In: ICLR (2019)

- 18 Q. Liu, Z. Wang
- 66. Van Gansbeke, W., Vandenhende, S., Georgoulis, S., Proesmans, M., Van Gool, L.: Scan: Learning to classify images without labels. In: ECCV (2020)
- 67. Wang, T., Isola, P.: Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In: ICML (2020)
- Wei, L., Zhang, S., Gao, W., Tian, Q.: Person transfer gan to bridge domain gap for person re-identification. In: CVPR (2018)
- Wilson, G., Cook, D.J.: A survey of unsupervised deep domain adaptation. ACM Transactions on Intelligent Systems and Technology 11(5), 1–46 (2020)
- 70. Xie, J., Girshick, R., Farhadi, A.: Unsupervised deep embedding for clustering analysis. In: ICML (2016)
- 71. Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: CVPR (2020)
- Xu, S., Rao, H., Hu, X., Hu, B.: Prototypical contrast and reverse prediction: Unsupervised skeleton based action recognition. arXiv preprint arXiv:2011.07236 (2020)
- Yan, H., Ding, Y., Li, P., Wang, Q., Xu, Y., Zuo, W.: Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. In: CVPR (2017)
- 74. Yan, S., Xiong, Y., Lin, D.: Spatial temporal graph convolutional networks for skeleton-based action recognition. In: AAAI (2018)
- Yang, D., Wang, Y., Dantcheva, A., Garattoni, L., Francesca, G., Bremond, F.: Unik: A unified framework for real-world skeleton-based action recognition. In: BMVC (2021)
- Yang, S., Liu, J., Lu, S., Er, M.H., Kot, A.C.: Skeleton cloud colorization for unsupervised 3d action representation learning. In: ICCV (2021)
- Yao, F.: Cross-domain few-shot learning with unlabelled data. arXiv preprint arXiv:2101.07899 (2021)
- Zhai, Y., Ye, Q., Lu, S., Jia, M., Ji, R., Tian, Y.: Multiple expert brainstorming for domain adaptive person re-identification. In: ECCV (2020)
- Zhang, P., Lan, C., Xing, J., Zeng, W., Xue, J., Zheng, N.: View adaptive recurrent neural networks for high performance human action recognition from skeleton data. In: ICCV (2017)
- Zhang, P., Lan, C., Zeng, W., Xing, J., Xue, J., Zheng, N.: Semantics-guided neural networks for efficient skeleton-based human action recognition. In: CVPR (2020)
- 81. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: ECCV (2016)
- 82. Zhang, X., Cao, J., Shen, C., You, M.: Self-training with progressive augmentation for unsupervised cross-domain person re-identification. In: ICCV (2019)
- Zhang, Z.: Microsoft kinect sensor and its effect. IEEE multimedia 19(2), 4–10 (2012)
- 84. Zhao, B., Han, K.: Novel visual category discovery with dual ranking statistics and mutual knowledge distillation. In: NIPS (2021)
- Zhao, F., Liao, S., Xie, G.S., Zhao, J., Zhang, K., Shao, L.: Unsupervised domain adaptation with noise resistible mutual-training for person re-identification. In: ECCV (2020)
- Zhao, L., Wang, Y., Zhao, J., Yuan, L., Sun, J.J., Schroff, F., Adam, H., Peng, X., Metaxas, D., Liu, T.: Learning view-disentangled human pose representation by contrastive cross-view mutual information maximization. In: CVPR (2021)
- Zheng, K., Liu, W., He, L., Mei, T., Luo, J., Zha, Z.J.: Group-aware label transfer for domain adaptive person re-identification. In: CVPR (2021)

- Zheng, N., Wen, J., Liu, R., Long, L., Dai, J., Gong, Z.: Unsupervised representation learning with long-term dynamics for skeleton based action recognition. In: AAAI (2018)
- 89. Zhong, Z., Zheng, L., Luo, Z., Li, S., Yang, Y.: Invariance matters: Exemplar memory for domain adaptive person re-identification. In: CVPR (2019)