

Learning Clothing and Pose Invariant 3D Shape Representation for Long-Term Person Re-Identification

Feng Liu, Minchul Kim, ZiAng Gu, Anil Jain, Xiaoming Liu
Department of Computer Science and Engineering
Michigan State University, East Lansing MI 48824
{liufeng6, kimminc2, guziang, jain, liuxm}@msu.edu

Abstract

Long-Term Person Re-Identification (LT-ReID) has become increasingly crucial in computer vision and biometrics. In this work, we aim to extend LT-ReID beyond pedestrian recognition to include a wider range of real-world human activities while still accounting for cloth-changing scenarios over large time gaps. This setting poses additional challenges due to the geometric misalignment and appearance ambiguity caused by the diversity of human pose and clothing. To address these challenges, we propose a new approach **3DInvarReID** for (i) disentangling identity from non-identity components (pose, clothing shape, and texture) of 3D clothed humans, and (ii) reconstructing accurate 3D clothed body shapes and learning discriminative features of naked body shapes for person ReID in a joint manner. To better evaluate our study of LT-ReID, we collect a real-world dataset called **CCDA**, which contains a wide variety of human activities and clothing changes. Experimentally, we show the superior performance of our approach for person ReID. Code is available at <http://cvlab.cse.msu.edu/project-reid3dinvar.html>.

1. Introduction

Person Re-Identification (ReID) aims to recognize and match a specific pedestrian in various locations and at different times [1, 26, 29, 57, 62]. This is a crucial task for various applications, including crime prevention, forensic identification and security monitoring [13, 63].

Most existing works [11, 12, 24, 25] in this field concentrate on the short-term scenarios, assuming that pedestrians' clothing remains unchanged. However, in this paper, we focus on a more challenging yet practical scenario of Long-Term Person Re-Identification (LT-ReID), where the objective is to recognize individuals over long time periods while taking into account **variations in clothing** and **diverse human activities**. For the first time, we extend person re-

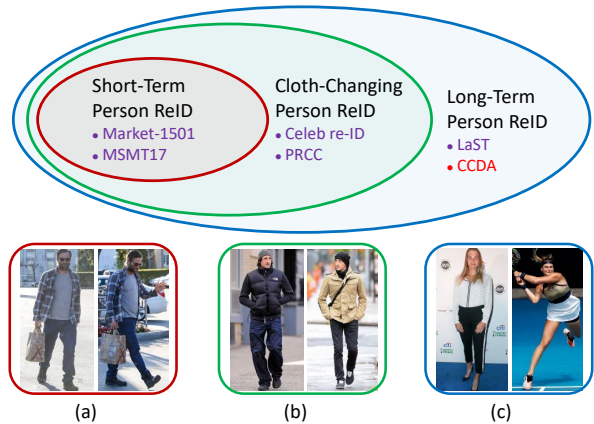


Figure 1. Illustration of the differences between various person re-identification (ReID) settings. Both (a) *conventional/short-term* and (b) *cloth-changing* person ReID benchmarks often restrict subjects to walking or standing, limiting their applications in real-world scenarios. This paper expands on the long-term person re-identification (LT-ReID) setting by tackling a wider range of human activities, increasing its practicality.

identification beyond pedestrian recognition to encompass a wider range of human activities, such as identifying students playing tennis or soldiers crawling in the field (see Fig. 1). This setting poses new challenges due to the geometric misalignment and appearance ambiguity caused by the diversity of human poses and their clothing.

Recently, various approaches [14, 18, 20, 23, 27, 46, 56, 59] have been proposed to investigate LT-ReID under clothing changes. They extract clothing-irrelevant features for robust person ReID by custom-designed architectures [19, 20], training process [27], loss functions [14], and data augmentation [64]. However, these methods only attempt to mine texture-insensitive body-structural cues in 2D space while ignoring the prior knowledge that the human body is a 3D non-rigid object. A new line of research introduces 3D priors for LT-ReID by either lifting 2D images to a 3D space [67] or including 3D body reconstruction as an auxil-

iary task [5]. However, without modeling the 3D clothing, the clothing-sensitive features can not be properly disentangled in either method. Moreover, none of the methods above handle body images with diverse activities.

Given the numerous variations in body images, including body pose, clothing, and view angles, we posit that the most reliable identity cue for LT-ReID is the 3D naked (unclothed) body shape, if it can be accurately and discriminatively estimated from a 2D body image. Obviously this is extremely challenging due to confounding factors and the lack of supervision, such as paired images and 3D naked body scans. However, taking inspiration from advancements in 3D feature learning for face recognition [32, 41], we propose a new algorithm, **3DInvarReID**, to disentangle identity (naked body) from non-identity components (pose, clothing shape and texture) of 3D clothed humans. This innovative approach not only reconstructs accurate 3D clothed body shapes that faithfully represent the input 2D images, but it also simultaneously learns discriminative naked shape features that effectively enhance LT-ReID.

An effective representation of the 3D shape and texture of the human body is a key component of such a learning-based process. To this end, we propose a joint two-layer neural implicit function to represent 3D humans, where identity, clothing shape, and texture components are disentangled into latent representations. Based on the composite model, we jointly learn a model fitting module to disentangle identity from non-identity components (body pose, clothing shape and texture) from 2D images. Modeling texture, along with a differentiable renderer enables us to compare the rendered image with the input image in a self-supervised manner. This allows the learning process to be supervised by both image reconstruction loss and identification loss, using a set of 2D images with identity labels *only*. Comprehensive experiments demonstrate the superiority of our method in diverse ReID benchmarks. Additionally, to advance the research in the field of LT-ReID, we collect a Cloth-Changing and Diverse Activities (CCDA) dataset (see Fig. 1). The CCDA dataset is specifically designed to evaluate the ReID of the person undergoing both human activities and changes in clothing.

In summary, the contributions of this work include:

- ◊ We propose a novel LT-ReID method, **3DInvarReID**, to learn clothing/pose invariant 3D shape representation.
- ◊ We devise a novel joint two-layer implicit model that fully models a textured 3D clothed human. Our approach includes a robust and discriminative fitting process that disentangles identity and non-identity features in reconstructing two-layer 3D body shapes from real-world images.
- ◊ We achieve superior performance in both LT-ReID accuracy and 3D body shape reconstruction.

Method	End-to-End trainable	Model texture	Discriminative	Model type
<i>3D modeling methods</i>				
SCANimate [44]	-	✗	✗	universal
SMPLicit [10]	-	✗	✗	universal
Neural-GIF [49]	-	✗	✗	individual
SNARF [7]	-	✗	✗	individual
gDNA [6]	-	✗	✗	universal
<i>3D fitting methods</i>				
PiFu [42]	✗	✓	✗	-
PiFuHD [43]	✗	✗	✗	-
Arch [21]	✗	✓	✗	-
Arch++ [15]	✗	✓	✗	-
ICON [55]	✗	✓	✗	-
ClothWild [37]	✗	✗	✗	-
PHORHUM [2]	✓	✓	✗	-
3DInvarReID	✓	✓	✓	universal

Table 1. Overview of the 3D clothed human modeling (top) and fitting (bottom) methods. Our method is the only one that models clothing texture and learns discriminative information compared to 3D modeling methods. Compared to 3D fitting methods, our end-to-end trainable pipeline enables disentangling identity-sensitive shape features from whole-body images.

2. Prior Work

Person Re-identification. Person ReID aims to match a person across images captured by a distributed camera system. The majority of prior methods [11, 12, 24, 25, 28, 52, 58, 60] assume a short-term application scenario without clothing changes by the person. This limitation has generated a growing interest in long-term cloth-changing person ReID [14, 23, 27, 56]. Datasets such as Real28 [50], VC-Clothes [50], PRCC [56], LTCC [46], COCAS [59] and Celebrities-reID [18, 20] are collected to facilitate this research. These datasets, however, either ignore or only minimally consider human activities, assuming that subjects are pedestrians with a restricted set of activities, limiting their applicability in real-world scenarios. As a result, there is a noticeable discrepancy between published approaches and the real-world LT-ReID problem. In contrast to the focus on the clothing-change person ReID, our research takes a step further by addressing a more challenging and practical issue of person ReID that involves diverse human activities, which are not limited to walking.

3D Clothed Human Modeling and Fitting. In early attempts [22, 34, 40], a clothed person was modeled as displacements over naked body meshes, obtained by SMPL [33]. However, the fixed mesh topology and bounded resolution approach limit geometric expressivity. Recently, neural implicit representations have been explored to model 3D body shapes due to their topological flexibility and resolution independence [6, 7, 10, 38, 44, 45, 49, 53]. However, as shown in Tab. 1, modeling texture in 3D remains a challenge. While these approaches provide rich geometric detail, insufficient attention has been paid




Figure 2. Overview of the proposed joint learning framework for long-term person re-identification and 3D clothed body shape reconstruction. During the inference of ReID, the identity shape feature \mathbf{z}_{id} is utilized for matching.

to the discriminativeness of the resulting body shapes. In contrast, we build a *discriminative* and *textured* 3D clothed human, serving the purpose of LT-ReID.

These 3D clothed models can be naturally applied to monocular 3D reconstruction (2D-to-3D fitting). Generally, the input image is encoded as a latent vector, from which the generative model reconstructs the 3D shape [10, 30, 31, 37]. Alternatively, two-step pipelines [15, 21, 42, 43, 55] firstly recover 2.5D sketches (*e.g.*, surface normal), and then infer a full 3D shape. A common limitation of these works is that they require 3D body scans for training, as they are trained on synthetic datasets derived from 3D scans and their rendered images. Furthermore, existing methods do not explicitly consider the discriminative ability of reconstructed 3D clothed body shapes. LVD [9] and SHAPY [8] introduce new discriminative fitting pipelines to reconstruct naked body shapes from images. However, without modeling 3D clothing, 2D image cues can not be fully exploited.

We propose a novel joint two-layer shape and texture representation of a 3D clothed human model, consisting of both shape and texture. Together with a model fitting module, our representation allows semi-supervised training from images without 3D labels. More importantly, guided by the completed 3D model and the discriminative 2D-to-3D fitting module, our approach disentangles identity-features from identity-irrelevant features in 3D space for LT-ReID. Tab. 1 compares our method with prior works.

3. Proposed Method

3.1. Problem Formulation

A 3D clothed human model is described by three disentangled latent variables: identity shape, clothing shape and clothing texture. As shown in Fig. 2, these latent repre-

sentations can be sequentially decoded into *canonical* 3D shape and texture, respectively by three decoders. To enable self-supervised training on real images, we estimate these latent codes along with the body pose and camera projection parameters. In this work, we use an off-the-shelf method [36] as our PoseNet to predict pose and camera projection, while our image encoder focuses on identity disentanglement learning, *i.e.*, the fitting module. These networks disentangle identity and non-identity components of 3D shapes and reconstruct the input body images via a differentiable render.

Formally, given a training set of T images $\{\mathbf{I}_i\}_{i=1}^T$ and the corresponding identity labels $\{l_i\}_{i=1}^T$, the image encoder $\mathcal{E}(\mathbf{I}) : \mathbf{I} \rightarrow \mathbf{z}_{id}, \mathbf{z}_{cloth}, \mathbf{z}_{tex}$ predicts the identity shape code of naked body $\mathbf{z}_{id} \in \mathbb{R}^{L_{id}}$, clothed shape code $\mathbf{z}_{cloth} \in \mathbb{R}^{L_{cloth}}$ and texture code $\mathbf{z}_{tex} \in \mathbb{R}^{L_{tex}}$. Functions \mathcal{F} , \mathcal{C} and \mathcal{T} decode the latent codes to identity shape, clothing shape and texture components, respectively. Additionally, PoseNet \mathcal{P} predicts the camera projection parameters Ω and SMPL body pose θ : $(\Omega, \theta) = \mathcal{P}(\mathbf{I})$.

Mathematically, the learning objective is defined as:

$$\arg \min_{\mathcal{E}, \mathcal{F}, \mathcal{C}, \mathcal{T}} \sum_{i=1}^T \left(\left\| \hat{\mathbf{I}}_i - \mathbf{I}_i \right\|_1 + \mathcal{L}_{cla}(\mathbf{z}_{id}, l_i) \right), \quad (1)$$

where \mathcal{L}_{cla} is the classification loss. $\hat{\mathbf{I}}$ is the rendered image. This objective enables us to jointly learn accurate 3D clothed shape and discriminative shape for the naked body.

3.2. Joint Two-Layer Implicit Model

We jointly model 3D naked body shape, clothed shape and texture in a *canonical* space by implicit representations.

Discriminative Body Shape Component. We represent the 3D naked body shape as the $\tau = 0.5$ level set of the




Figure 3. Joint two-layer implicit model. The naked body shape model \mathcal{F} and clothed body shape model \mathcal{C} take identity shape code \mathbf{z}_{id} , clothing shape code \mathbf{z}_{cloth} and a spatial point \mathbf{x} , and produces two occupancy values o_1 and o_2 . The texture model \mathcal{T} takes \mathbf{z}_{tex} and \mathbf{z}_{cloth} to estimate RGB color at \mathbf{x} . A 3D generator \mathcal{G} uses \mathbf{z}_{id} to produce a 3D feature volume, enabling hierarchical point-wise feature representation.

occupancy function [35]:

$$\mathbf{S}_{id}(\mathbf{z}_{id}) = \{\mathbf{x} | \mathcal{F}(\mathbf{z}_{id}, \mathbf{x}) = \tau\}, \quad (2)$$

where \mathcal{F} predicts the occupancy value, o_1 for any point \mathbf{x} in the canonical space. Specifically,

$$\mathcal{F} : \mathbb{R}^{L_{id}} \times \mathbb{R}^3 \rightarrow (o_1, \mathbf{f}), \quad (3)$$

where $\mathbf{f} \in \mathbb{R}^{L_f}$ is a point-wise feature and will be utilized to predict clothing details.

Following [6], we make use of function \mathcal{F} , implemented via a Multi-Layer Perceptron (MLP), coupled with a 3D CNN-based generator \mathcal{G} to model 3D naked bodies. As shown in Fig. 3(a), the generator produces a 3D feature volume using \mathbf{z}_{id} as input. We then use trilinear interpolation to query continuous 3D points and feed the feature at \mathbf{x} to the MLP.

Clothed Shape Component. Similarly, we also represent the clothed body shape as the $\tau = 0.5$ level set function:

$$\mathbf{S}_{cloth}(\mathbf{z}_{cloth}) = \{\mathbf{x} | \mathcal{C}(\mathbf{z}_{cloth}, \mathbf{f}, \mathbf{x}) = \tau\}, \quad (4)$$

where \mathcal{C} is implemented as a MLP:

$$\mathcal{C} : \mathbb{R}^{L_{cloth}} \times \mathbb{R}^{L_f} \times \mathbb{R}^3 \rightarrow o_2. \quad (5)$$

\mathcal{C} outputs the occupancy value o_2 to represent the clothed shape information.

Texture Component. We define a texture field as a mapping function \mathcal{T} from a point \mathbf{x} in the canonical space, texture latent \mathbf{z}_{tex} and \mathbf{z}_{cloth} to a RGB value $\mathbf{c} \in \mathbb{R}^3$:

$$\mathcal{T} : \mathbb{R}^{L_{tex}} \times \mathbb{R}^{L_{cloth}} \times \mathbb{R}^3 \rightarrow \mathbf{c}. \quad (6)$$

The design of our joint two-layer implicit model is inspired by the approach in [6]. However, as shown in Fig. 3, our model has two novel traits: 1) *Instead of simply decomposing the 3D clothed human into coarse and fine models, we apply a two-layer implicit model to represent the naked body and clothing shapes.* 2) *We additionally model texture to form a complete 3D human model.*




Figure 4. Neural blend skinning network. This module deforms the pose space to canonical space. Given a deformed point \mathbf{x}' , we compute its corresponding position $\hat{\mathbf{x}}$ in canonical space by iteratively finding the root of Eqn. 9.

3.3. Neural Linear Blend Skinning Network

Our joint two-layer implicit model is built within the canonical space. However, the 3D clothed human data, usually captured in various poses, introduces misalignments between this canonical space and the deformed counterpart. To predict the occupancy values o_1 , o_2 and the texture \mathbf{c} for a given observed point \mathbf{x}' within the deformed space, it is essential to first determine its canonical correspondence point $\hat{\mathbf{x}}$. Once \mathbf{x} is identified, we can then compute o_1 , o_2 and \mathbf{c} using Eqns. 3, 5 and 6. The objective of this step is to find the canonical correspondence $\hat{\mathbf{x}}$ of any query point \mathbf{x}' . To achieve this goal, similar to [6, 7], we learn a linear blend skinning (LBS) [33] using neural networks in an unsupervised manner.

Regressing Blend Weight. We follow [6, 7] which define the skinning field in canonical space conditioned on our identity latent code \mathbf{z}_{id} :

$$\begin{aligned} \mathcal{W} : \mathbb{R}^{L_{id}} \times \mathbb{R}^3 &\rightarrow \mathbb{R}^K \\ (\mathbf{z}_{id}, \mathbf{x}) &\rightarrow \mathbf{w}, \end{aligned} \quad (7)$$

where \mathbf{w} is the point-wise blend weight of the canonical point \mathbf{x} . Then deformed point \mathbf{x}' is determined by the following convex combination:

$$\mathbf{x}' = \mathbf{d}(\mathbf{z}_{id}, \mathbf{x}, \theta) = \sum_{k=1}^K \mathbf{w}_k \mathbf{B}_k \mathbf{x}, \quad (8)$$

where \mathbf{w}_k is k -element in the vector \mathbf{w} , while \mathbf{B}_k denotes the k -element from the set of bone transformation matrices $\mathcal{B} = \{\mathbf{B}_k \in \mathbb{R}^{4 \times 4}\}_{k=1}^K$. Both \mathbf{x} and \mathbf{x}' in Eqn. 8 are represented in homogeneous coordinates.

Implicit Differentiable Skinning. An overview is illustrated in Fig. 4. While the goal is to determine $\mathbf{x}' \rightarrow \hat{\mathbf{x}}$, we only have direct access to the mapping defined by Eqn. 8, which is not invertible. Following [6, 7], the correspondence is calculated numerically by finding the root of the equation with Broyden’s method [3]:

$$\hat{\mathbf{x}} = \{\hat{\mathbf{x}} | \mathbf{d}(\mathbf{z}_{id}, \hat{\mathbf{x}}, \theta) - \mathbf{x}' = \mathbf{0}\}. \quad (9)$$

3.4. Implicit Rendering

During rendering, when 2D pixels are unprojected to 3D points, they are intrinsically mapped in the deformed 3D space. Given a pixel p of a masked input image, we construct a ray $\mathbf{x}' = \{\mathbf{c}_0 + t\mathbf{v} | t \geq 0\}$, where \mathbf{c}_0 represents the camera’s position and \mathbf{v} indicates the viewing direction based on the camera projection parameters Ω . t is the scalar distance along the ray. We then map the ray points to the canonical space, following Eqn. 9. The intersection point $\hat{\mathbf{x}}_p$ of the ray can be calculated by identifying the first change of clothed shape occupancy o_2 . Finally, the rendered color of the pixel p is calculated via Eqn. 6.

3.5. Semi-supervised Model Learning

While our model can perform self-supervised learning from real images without 3D labels, we first pre-train our joint implicit 3D model with 3D data in order to mitigate the inherent ambiguity.

3.5.1 Supervised Pre-training 3D clothed Model

Training Data. We combine CAPE [34] (3,000 scans) and THuman2.0 [65] (526 scans) to train our joint two-layer implicit model. THuman2.0 consists of 526 texture clothed 3D scans of 105 subjects. Following [6], for each scan, we obtain its SMPL naked shape code. CAPE provides 148,584 pairs of scans under clothing and SMPL naked body with rich pose variations of 15 subjects. We randomly sample 3,000 scans for training. Formally, each training sample can be represented as SMPL pose θ , identity label l_{3D} , n spatial points \mathbf{x}'_i , and their SDFs o_1^i, o_2^i , and color \mathbf{c}_i : $\{\theta, l_{3D}, \{\mathbf{x}'_i, o_1^i, o_2^i, \mathbf{c}_i\}_{i=1}^n\}$. With the autodecoding technique [39], we assign trainable identity shape code, clothing shape code, and texture code to each training sample.

Loss Function. We define the loss below for each sample:

$$\arg \min_{\mathcal{F}, \mathcal{C}, \mathcal{T}, \mathbf{z}_{id}, \mathbf{z}_{cloth}, \mathbf{z}_{tex}} \mathcal{L}_{id} + \mathcal{L}_{cloth} + \mathcal{L}_{tex} + \mathcal{L}_{cla}^{3D}(\mathbf{z}_{id}, l_{3D})$$

$$\mathcal{L}_{id} = \sum_{i=0}^n BCE(\mathcal{F}(\mathbf{z}_{id}, \hat{\mathbf{x}}_i), o_1^i) \quad (10)$$

$$\mathcal{L}_{cloth} = \sum_{i=0}^n BCE(\mathcal{C}(\mathbf{z}_{cloth}, \mathbf{f}_i, \hat{\mathbf{x}}_i), o_2^i) \quad (11)$$

$$\mathcal{L}_{tex} = \sum_{i=0}^n \|\mathcal{T}(\mathbf{z}_{tex}, \mathbf{z}_{cloth}, \hat{\mathbf{x}}_i) - \mathbf{c}_i\|_2, \quad (12)$$

where $\mathcal{L}_{cla}^{3D}(\mathbf{z}_{id}, l_{3D})$ is the cross-entropy classification loss. We additionally add auxiliary loss \mathcal{L}_W to train network \mathcal{W} :

$$\arg \min_{\mathcal{W}} \mathcal{L}_W \quad (13)$$

$$\mathcal{L}_W = \sum_{k=0}^K \|\mathcal{W}(\mathbf{z}_{id}, \mathbf{J}_k) - \mathbf{w}_{\mathbf{J}_k}\|_2, \quad (14)$$

where $\mathbf{w}_{\mathbf{J}_k}$ is the pre-computed ground truth skinning weights of SMPL joints location \mathbf{J}_k .



Figure 5. Example images from CCDA, with one subject per row, showcasing the diversity of body poses, clothing styles and colors.

3.5.2 Self-supervised Joint Modeling and Fitting

Given a set of in-the-wild 2D images with body masks and identity labels $\{\mathbf{I}_i, \mathbf{M}_i, l_i\}_{i=1}^T$, the self-supervised identity disentanglement loss is:

$$\arg \min_{\mathcal{E}, \mathcal{T}} \sum_{i=1}^T \mathcal{L}_{sil} + \mathcal{L}_{rgb} + \mathcal{L}_{cla}, \quad (15)$$

where \mathcal{L}_{rgb} is the photometric loss, \mathcal{L}_{sil} is silhouette loss and \mathcal{L}_{cla} is the classification loss. Specifically, we denote \mathbf{I}_p and \mathbf{M}_p as the RGB and silhouette values of pixel $p \in P$. Here, P denotes the entire set of pixels in the input image \mathbf{I} . A subset of P , represented as P^{in} , corresponds to the pixels where an intersection between the rays and the body in the image has been detected. The photometric loss is defined as

$$\mathcal{L}_{rgb} = \frac{1}{|P|} \sum_{p \in P^{in}} |\mathbf{I}_p - \mathcal{T}(\mathcal{E}_{tex}(\mathbf{I}), \mathcal{E}_{cloth}(\mathbf{I}), \hat{\mathbf{x}}_p)|, \quad (16)$$

where the encoder \mathcal{E} estimates $\mathbf{z}_{id} = \mathcal{E}_{id}(\mathbf{I})$, $\mathbf{z}_{cloth} = \mathcal{E}_{cloth}(\mathbf{I})$ and $\mathbf{z}_{tex} = \mathcal{E}_{tex}(\mathbf{I})$ from image \mathbf{I} . $\hat{\mathbf{x}}$ is the intersection point (see Sec. 3.4).

We further define the silhouette loss as

$$\mathcal{L}_{sil} = \frac{1}{|P|} \sum_{p \in P^{out}} CE(\mathbf{M}_p, \hat{\mathbf{M}}_p), \quad (17)$$

where $\hat{\mathbf{M}}$ is the masked rendering, $P^{out} = P - P^{in}$ represents the indices in the mini-batch for which there is no ray-geometry intersection or $\mathbf{M}_p = 0$, and $CE(\cdot, \cdot)$ denotes the cross-entropy loss. We impose triplet loss and cross-entropy loss on the identity shape code $\mathbf{z}_{id} = \mathcal{E}_{id}(\mathbf{I})$ as our classification loss $\mathcal{L}_{cla}(\mathbf{z}_{id}, l)$.

3.6. Person ReID Inference

For person ReID inference, the encoder \mathcal{E} processes body images and extracts the identity shape features \mathbf{z}_{id} .

The Cosine similarity of z_{id} is then used to determine if the two images belong to the same person. It is worth noting that, the inference of our ReID does not incorporate the 3D reconstruction module, making it highly efficient.

4. CCDA Person ReID Dataset

We construct a new dataset with diverse human activities and clothing changes for evaluating LT-ReID. Specifically, we collect data for popular athletes in soccer, tennis, and basketball, and popular artists, such as fashion models and singers. We crawl whole body images of each subject on Google Image¹ with athlete/artist names. We collect two sets of images per subject: ‘challenging’ and ‘normal’ body poses. As an example, for a basketball player, the ‘challenging’ set includes images of players’ actions on the court, while the ‘normal’ set contains standing or walking poses. We then crop the body region from the original image via the detected bounding box and resize it to 256×128 . Finally, the annotator verifies the identity of each image. In total, 1,555 images of 100 subjects are retained. For each subject, we randomly select one image from ‘normal’ images for the gallery set, while the remaining 1,455 images comprise the query set. Fig. 5 shows examples of images in the CCDA dataset.

5. Experimental Results

Implementation Details. Our training process includes two stages: 1) Networks \mathcal{F} , \mathcal{C} , \mathcal{T} , \mathcal{W} are pre-trained on 3D data. 2) \mathcal{E} and \mathcal{T} are trained or fine-tuned with real images. The encoder \mathcal{E} is implemented as a ResNet-50. Networks \mathcal{F} , \mathcal{C} , \mathcal{T} and \mathcal{W} are MLPs. In experiments, we set $L_{cloth} = L_{tex} = 512$, $L_{id} = 4,096$, $L_f = L_h = 256$, $K = 24$, $n = 200,000$. We implement in Pytorch and use Adam optimizer in both stages.

5.1. Person ReID

Metric. For person ReID, we follow the standard retrieval accuracy metrics, namely the Cumulative Matching Characteristics (CMC) and mean average precision (mAP).

Baseline. We compare our method with eight SoTA person ReID methods: Two-Stream [66], MLFN [4], HACNN [26], Part-Aligned [47], PCB [48], TriNet [16], MGN [51], DG-Net [64], and five SoTA cloth-changing re-ID methods: ReIDCaps [20], 3DSL [5], RCSAnet [19], FSAM [17] and CAL [14].

5.1.1 Results on Cloth-changing Person ReID datasets

Datasets. We test on four popular cloth-changing ReID datasets: Celeb-reID/Celeb-reID-light [18, 20], PRCC [56], LTCC [46] and the recent CCVID dataset [14, 61].

¹All collected images are under Creative Commons licenses.

Method	Backbone	Celeb-reID		Celeb-reID-light	
		mAP	Rank1	mAP	Rank1
Two-Stream [66]	ResNet-50	7.8	36.3	-	-
MLFN [4]	*	6.0	41.4	6.3	10.6
HACNN [26]	*	9.5	47.6	11.5	16.2
Part-Aligned [47]	GoogLeNet $\times 2$	6.4	19.4	-	-
PCB [48]	ResNet-50	8.2	37.1	-	-
MGN [51]	ResNet-50	10.8	49.0	13.9	21.5
DG-Net [64]	ResNet-50	10.6	50.1	12.6	23.5
<i>cloth-changing person ReID methods</i>					
ReIDCaps- [20]	DenseNet-121	9.8	51.2	11.2	20.3
ReIDCaps [20]	DenseNet-121 $\times 6$	15.8	63.0	19.0	33.5
RCSAnet [19]	DenseNet-121 $\times 2$	11.9	55.6	16.7	29.5
CAL [14]	ResNet-50	13.7	59.2	18.5	33.6
3DInvarReID	ResNet-50	11.8	55.2	15.0	30.1
3DInvarReID [#]	ResNet-50	15.2	61.2	21.8	37.0
ReIDCaps+ 3DInvarReID [#]	-	18.4	65.5	25.7	42.2

Table 2. Comparison with SoTA on Celeb-reID and Celeb-reID-light datasets (%). ‘*’ indicates that the backbone is designed by the authors. The red number means the total number of models. **3DInvarReID**[#]’s weights are initialized using the CAL model.

Method	Backbone	LTCC		PRCC	
		mAP	Rank1	mAP	Rank1
HACNN [26]	*	9.3	21.6	-	21.8
PCB [48]	ResNet-50	10.0	23.5	38.7	41.8
3DSL [5]	ResNet-50 $\times 2$	14.8	31.2	-	51.3
FSAM [17]	ResNet-50	16.2	38.5	-	54.5
CAL [14]	ResNet-50	18.0	40.1	55.8	55.2
3DInvarReID	ResNet-50	16.7	37.8	52.5	51.6
CAL+ 3DInvarReID	-	18.9	40.9	57.2	56.5
CAL	ResNet-50	2.8	3.8	20.3	31.6
3DInvarReID	ResNet-50	2.8	5.1	20.1	34.6
3DInvarReID [#]	ResNet-50	2.8	5.6	21.4	40.7

Table 3. Comparison with SoTA cloth-changing person ReID methods on the LTCC and PRCC datasets (%). Models highlighted in pink are trained on the Celeb-reID dataset. **3DInvarReID**[#]’s weights are initialized using the CAL model.

Results on Celeb-reID and Celeb-reID-light. As reported in Tab. 2, our **3DInvarReID (Ours)** outperforms all the general person re-ID baselines on both datasets. Considering methods using a *single* network, we also outperform the cloth-changing baseline, ReIDCaps- [20]. More importantly, to investigate the complementarity between our learned 3D shape features and existing 2D features, we fuse our method with cloth-changing baselines by simple summation at the score level. By fusing with ReIDCaps [20], our method improves the Rank1 accuracy on Celeb-reID from 63.0% to 65.5% and from 33.5% to 42.2% on Celeb-reID-light. These results clearly demonstrate that the 3D shape features learned from our method are both discriminative and complementary to the 2D features, indicating the effectiveness of our proposed approach for person ReID, particularly under cloth-changing scenarios.

Results on LTCC, PRCC and CCVID. The compari-

Method	Backbone	General		Cloth-changing	
		mAP	Rank1	mAP	Rank1
TriNet [16]	ResNet-50	78.1	81.5	77.0	81.1
CAL [14]	ResNet-50	81.3	82.6	79.6	81.7
3DInvarReID	ResNet-50	66.1	70.8	65.4	70.2
CAL+ 3DInvarReID	-	82.6	83.9	81.3	84.3

Table 4. Comparison with SoTA methods on CCVID (%).

Method	Backbone	mAP	Rank1	Rank5
ReIDCaps [20]	DenseNet-121×6	10.9	6.5	20.2
CAL [14]	ResNet-50	19.3	10.0	26.7
3DInvarReID	ResNet-50	21.7	11.1	30.5

Table 5. Comparison with SoTA methods on CCDA dataset (%).

son of the LTCC, PRCC and CCVID datasets is shown in Tabs. 3 and 4. Similarly, by fusing the best baseline CAL [14], our method achieves additional improvements. For instance, when evaluated in the cloth-changing setting, our method achieves a significant improvement of 2.6% in Rank1 accuracy by fusing with CAL on the CCVID dataset. These findings highlight the effectiveness of **3DInvarReID**, with the 3D shape features being shown to be both discriminative and complementary to 2D features. Our approach stands out for its superior performance compared to other 3D feature extraction methods for person ReID. This is particularly noteworthy in comparison to the 3DSL [5] (Tab. 3). Additionally, we assess our model under a *cross-domain setting*—that is, the model is trained with the Celeb-reID dataset and tested using the LTCC and PRCC datasets. Table 3 shows that our models outperform the baseline CAL, indicating a superior discriminative feature representation.

5.1.2 Results on LT-ReID dataset (CCDA)

Given that both our CCDA and Celeb-reID datasets are obtained from the Internet and share a similar image style, we choose trained models on Celeb-reID and evaluate them on CCDA. We choose the SoTA cloth-changing methods, ReIDCaps- [20] and CAL [14] as baselines. The results in Tab. 5 demonstrate that our **3DInvarReID** outperforms the baselines, providing strong evidence of its effectiveness in handling person ReID with challenging body poses and cloth-changing variations.

5.1.3 Results on Short-term Person ReID datasets

Despite our method is designed for LT-ReID, we additionally compare with SoTA methods on two conventional short-term ReID datasets: Market-1501 [62] and MSMT17 [54], in Tab. 6. By fusing with CAL [14], we observe an average improvement of 2.7% in Rank1 accuracy on both datasets, demonstrating the complementary nature

Method	Backbone	Market-1501		MSMT17	
		mAP	Rank1	mAP	Rank1
PCB [48]	ResNet-50	81.6	93.8	40.4	68.2
3DSL [5]	ResNet-50×2	87.3	95.0	—	—
FSAM [17]	ResNet-50	85.6	94.6	—	—
CAL [14]	ResNet-50	87.5	94.7	57.3	79.7
3DInvarReID	ResNet-50	85.5	94.2	55.1	76.3
CAL+ 3DInvarReID	-	87.9	95.1	59.1	80.8

Table 6. Comparison on short-term ReID datasets (%).

	LVD [9]	SHAPY [8]	Ours
CD- L_2	0.654	0.632	0.610

Table 7. Comparison of 3D boy reconstruction on HBW.

Model Type	mAP	Rank1	Rank5
3DInvarReID	17.5	36.3	56.4
3DInvarReID -w/o 3D	14.6	28.4	48.3
3DInvarReID -w/o Pre-training	15.1	29.3	50.7
3DInvarReID -w/o 3D clothing	16.0	32.1	52.8

Table 8. Ablation studies on CCDA dataset (%).

of our 3D shape feature, even on short-term datasets.

5.2. 3D Reconstruction

Most 3D body reconstruction methods focus more on pose estimation than shape estimation. Recently, SHAPY [8] releases a dataset (HBW) that contains ground-truth 3D body scans and the corresponding in-the-wild images, which enables us to test the accuracy of our reconstructed 3D naked body shapes. We thus evaluate our methods on the validation set of HBW, which contains 237 in-the-wild images of 10 subjects. Our baseline includes LVD [9] and SHAPY [8], which are recent pipelines for discriminative identity shape fitting. Following [9], we evaluate the reconstruction accuracy with Chamfer distance (CD- L_2), by uniformly sampling 10,000 points on both ground-truth and predicted meshes in the canonical space. As shown in Tab. 7, our method outperforms the two baselines. We also visualize 3D reconstructions in Fig. 6. Our reconstructions resemble the ground truth better than the baselines. These results demonstrate the superiority of the proposed method in reconstructing naked 3D body shapes. Fig. 7 shows qualitative comparisons with ICON [55] and ClothWild [37]. Our approach achieves comparable clothed 3D body reconstructions.

5.3. Ablation Study

In this section, all models are trained on Celeb-reID and tested on the CCDA dataset.

Effect of the 3D module. We compare our full model with an ablated version that only incorporates the \mathcal{L}_{cla} loss in its training, disregarding 3D modules. The results in




Figure 6. Qualitative comparisons with LVD [9] and SHAPY [8] on 3D naked body reconstruction. Our approach recovers more accurate 3D body shapes from the images.

Tab. 8 show that our **3DInvarReID** significantly improves the recognition accuracy, leading to a Rank1 accuracy increase from 28.4 to 36.3.

Effect of Our Two-layer Implicit Model. Our primary goal is to disentangle the identity feature from non-identity features in 3D shape space. To evaluate the effectiveness of our 3D disentanglement module, we train a model (**Ours-w/o** 3D clothing) by replacing our 3D body model with SMPL shape bases and omitting the modeling of the 3D clothing component and rendering layer. The results in Tab. 8 demonstrate the advantages of modeling clothing shape and texture for person ReID (Rank1: 32.1→36.3).

Effect of Pre-training. Tab. 8 shows that removing the pre-training stage results in a lower Rank1 accuracy, with a score of 29.3 compared to our 36.3. This finding highlights the effectiveness of pre-training in addressing the inherent ambiguity of disentanglement.

6. Conclusions

This paper tackles the challenging setting of long-term person ReID, which allows a wider range of real-world human activities and accounts for cloth-changing scenarios. To address this problem, we present a joint two-layer implicit representation to model textured 3D clothed humans

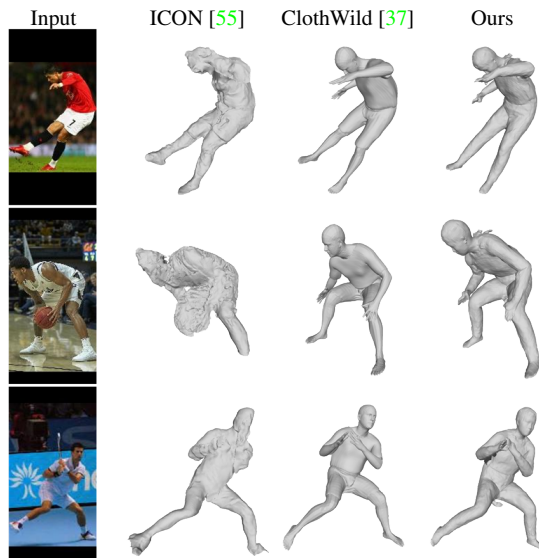


Figure 7. Qualitative comparisons with ICON [55] and ClothWild [37] on 3D clothed body reconstruction.

together with a discriminative fitting module, enabling us to disentangle identity and non-identity features for real-world images. We collect a new LT-ReID dataset, CCDA, with diverse human activities and clothing changes, facilitating future research on real-world scenarios. Experimental results demonstrate the effectiveness of our method in disentangling identity and non-identity features in 3D clothed body shapes, thereby contributing to LT-ReID.

Limitations & Potential Negative Impacts. Our work tackles the challenge of disentangling clothing and body shape in 3D shape representation. The clothing reconstruction task remains challenging as evidenced by the visual quality of the published models and our models. Our results show that the task of body-clothing disentanglement brings benefit in the recognition task, a finding which opens new possibilities to multi-task learning across 2D-3D modalities. Like most person ReID methods, one potential negative impact of our approach is that it could be used for unethical surveillance and invasion of privacy.

Acknowledgments. This research is based upon work supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2022-21102100004. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [1] Ejaz Ahmed, Michael Jones, and Tim K Marks. An improved deep learning architecture for person re-identification. In *CVPR*, 2015. 1
- [2] Thiemo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3D reconstruction of humans wearing clothing. In *CVPR*, 2022. 2
- [3] Charles G Broyden. A class of methods for solving nonlinear simultaneous equations. *Mathematics of computation*, 1965. 4
- [4] Xiaobin Chang, Timothy M Hospedales, and Tao Xiang. Multi-level factorisation net for person re-identification. In *CVPR*, 2018. 6
- [5] Jiaying Chen, Xinyang Jiang, Fudong Wang, Jun Zhang, Feng Zheng, Xing Sun, and Wei-Shi Zheng. Learning 3D shape feature for texture-insensitive person re-identification. In *CVPR*, 2021. 2, 6, 7
- [6] Xu Chen, Tianjian Jiang, Jie Song, Jinlong Yang, Michael J Black, Andreas Geiger, and Otmar Hilliges. gDNA: Towards generative detailed neural avatars. In *CVPR*, 2022. 2, 4, 5
- [7] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. SNARF: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, 2021. 2, 4
- [8] Vasileios Choutas, Lea Müller, Chun-Hao P Huang, Siyu Tang, Dimitrios Tzionas, and Michael J Black. Accurate 3D body shape regression using metric and semantic attributes. In *CVPR*, 2022. 3, 7, 8
- [9] Enric Corona, Gerard Pons-Moll, Guillem Alenyà, and Francesc Moreno-Noguer. Learned vertex descent: A new direction for 3D human model fitting. In *ECCV*, 2022. 3, 7, 8
- [10] Enric Corona, Albert Pumarola, Guillem Alenyà, Gerard Pons-Moll, and Francesc Moreno-Noguer. SMPlicit: Topology-aware generative model for clothed people. In *CVPR*, 2021. 2, 3
- [11] Yixiao Ge, Dapeng Chen, and Hongsheng Li. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. In *ICLR*, 2020. 1, 2
- [12] Yixiao Ge, Feng Zhu, Dapeng Chen, Rui Zhao, et al. Self-paced contrastive learning with hybrid memory for domain adaptive object re-id. In *NeurIPS*, 2020. 1, 2
- [13] Shaogang Gong, Tao Xiang, Shaogang Gong, and Tao Xiang. *Person re-identification*. Springer, 2011. 1
- [14] Xinqian Gu, Hong Chang, Bingpeng Ma, Shutao Bai, Shiguang Shan, and Xilin Chen. Clothes-changing person re-identification with RGB modality only. In *CVPR*, 2022. 1, 2, 6, 7
- [15] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. ARCH++: Animation-ready clothed human reconstruction revisited. In *ICCV*, 2021. 2, 3
- [16] Alexander Hermans, Lucas Beyers, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. 6, 7
- [17] Peixian Hong, Tao Wu, Ancong Wu, Xintong Han, and Wei-Shi Zheng. Fine-grained shape-appearance mutual learning for cloth-changing person re-identification. In *CVPR*, 2021. 6, 7
- [18] Yan Huang, Qiang Wu, Jingsong Xu, and Yi Zhong. Celebrities-ReID: A benchmark for clothes variation in long-term person re-identification. In *IJCNN*, 2019. 1, 2, 6
- [19] Yan Huang, Qiang Wu, Jingsong Xu, Yi Zhong, and Zhaoxiang Zhang. Clothing status awareness for long-term person re-identification. In *ICCV*, 2021. 1, 6
- [20] Yan Huang, Jingsong Xu, Qiang Wu, Yi Zhong, Peng Zhang, and Zhaoxiang Zhang. Beyond scalar neuron: Adopting vector-neuron capsules for long-term person re-identification. *TCSVT*, 2019. 1, 2, 6, 7
- [21] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. ARCH: Animatable reconstruction of clothed humans. In *CVPR*, 2020. 2, 3
- [22] Boyi Jiang, Juyong Zhang, Yang Hong, Jinhao Luo, Ligang Liu, and Hujun Bao. Bcnet: Learning body and cloth shape from a single image. In *ECCV*, 2020. 2
- [23] Xin Jin, Tianyu He, Kecheng Zheng, Zhiheng Yin, Xu Shen, Zhen Huang, Ruoyu Feng, Jianqiang Huang, Zhibo Chen, and Xian-Sheng Hua. Cloth-changing person re-identification from a single image with gait prediction and regularization. In *CVPR*, 2022. 1, 2
- [24] Minxian Li, Xiatian Zhu, and Shaogang Gong. Unsupervised person re-identification by deep learning tracklet association. In *ECCV*, 2018. 1, 2
- [25] Minxian Li, Xiatian Zhu, and Shaogang Gong. Unsupervised tracklet person re-identification. *PAMI*, 2019. 1, 2
- [26] Wei Li, Xiatian Zhu, and Shaogang Gong. Harmonious attention network for person re-identification. In *CVPR*, 2018. 1, 6
- [27] Yu-Jhe Li, Xinshuo Weng, and Kris M Kitani. Learning shape representations for person re-identification under clothing change. In *WACV*, 2021. 1, 2
- [28] Yutian Lin, Xuanyi Dong, Liang Zheng, Yan Yan, and Yi Yang. A bottom-up clustering approach to unsupervised person re-identification. In *AAAI*, 2019. 2
- [29] Feng Liu, Ryan Ashbaugh, Nicholas Chimits, Najmul Hassan, Ali Hassani, Ajay Jaiswal, Minchul Kim, Zhiyuan Mao, Christopher Perry, Zhiyuan Ren, et al. Farsight: A physics-driven whole-body biometric system at large distance and altitude. *arXiv preprint arXiv:2306.17206*, 2023. 1
- [30] Feng Liu and Xiaoming Liu. 2d gans meet unsupervised single-view 3d reconstruction. In *ECCV*, 2022. 3
- [31] Feng Liu, Luan Tran, and Xiaoming Liu. Fully understanding generic objects: Modeling, segmentation, and reconstruction. In *CVPR*, 2021. 3
- [32] Feng Liu, Ronghang Zhu, Dan Zeng, Qijun Zhao, and Xiaoming Liu. Disentangling features in 3D face shapes for joint face reconstruction and recognition. In *CVPR*, 2018. 2
- [33] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *TOG*, 2015. 2, 4
- [34] Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, and Michael J Black. Learning to dress 3D people in generative clothing. In *CVPR*, 2020. 2, 5

- [35] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *CVPR*, 2019. 4
- [36] Gyeongsik Moon, Hongsuk Choi, and Kyoung Mu Lee. Accurate 3D hand pose estimation for whole-body 3d human mesh estimation. In *CVPR*, 2022. 3
- [37] Gyeongsik Moon, Hyeongjin Nam, Takaaki Shiratori, and Kyoung Mu Lee. 3D clothed human reconstruction in the wild. In *ECCV*, 2022. 2, 3, 7, 8
- [38] Pablo Palafox, Aljaž Božič, Justus Thies, Matthias Nießner, and Angela Dai. Npms: Neural parametric models for 3D deformable shapes. In *ICCV*, 2021. 2
- [39] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 5
- [40] Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. Tailornet: Predicting clothing in 3D as a function of human pose, shape and garment style. In *CVPR*, 2020. 2
- [41] Xi Peng, Xiang Yu, Kihyuk Sohn, Dimitris N Metaxas, and Manmohan Chandraker. Reconstruction-based disentanglement for pose-invariant face recognition. In *ICCV*, 2017. 2
- [42] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 2, 3
- [43] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *CVPR*, 2020. 2, 3
- [44] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. SCANimate: Weakly supervised learning of skinned clothed avatar networks. In *CVPR*, 2021. 2
- [45] Ruizhi Shao, Hongwen Zhang, He Zhang, Mingjia Chen, Yan-Pei Cao, Tao Yu, and Yebin Liu. Doublefield: Bridging the neural surface and radiance fields for high-fidelity human reconstruction and rendering. In *CVPR*, 2022. 2
- [46] Xiujun Shu, Xiao Wang, Xianghao Zang, Shiliang Zhang, Yuanqi Chen, Ge Li, and Qi Tian. Large-scale spatio-temporal person re-identification: Algorithms and benchmark. *TCSVT*, 2021. 1, 2, 6
- [47] Yumin Suh, Jingdong Wang, Siyu Tang, Tao Mei, and Kyoung Mu Lee. Part-aligned bilinear representations for person re-identification. In *ECCV*, 2018. 6
- [48] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *ECCV*, 2018. 6, 7
- [49] Garvita Tiwari, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Neural-GIF: Neural generalized implicit functions for animating people in clothing. In *ICCV*, 2021. 2
- [50] Fangbin Wan, Yang Wu, Xuelin Qian, Yixiong Chen, and Yanwei Fu. When person re-identification meets changing clothes. In *CVPRW*, 2020. 2
- [51] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. Learning discriminative features with multiple granularities for person re-identification. In *ACMMM*, 2018. 6
- [52] Jingya Wang, Xiatian Zhu, Shaogang Gong, and Wei Li. Transferable joint attribute-identity deep learning for unsupervised person re-identification. In *CVPR*, 2018. 2
- [53] Shaofei Wang, Marko Mihajlovic, Qianli Ma, Andreas Geiger, and Siyu Tang. Metaavatar: Learning animatable clothed human models from few depth images. In *NeurIPS*, 2021. 2
- [54] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer GAN to bridge domain gap for person re-identification. In *CVPR*, 2018. 7
- [55] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J Black. Icon: Implicit clothed humans obtained from normals. In *CVPR*, 2022. 2, 3, 7, 8
- [56] Qize Yang, Ancong Wu, and Wei-Shi Zheng. Person re-identification by contour sketch under moderate clothing change. *PAMI*, 2019. 1, 2, 6
- [57] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven CH Hoi. Deep learning for person re-identification: A survey and outlook. *PAMI*, 2021. 1
- [58] Hong-Xing Yu, Wei-Shi Zheng, Ancong Wu, Xiaowei Guo, Shaogang Gong, and Jian-Huang Lai. Unsupervised person re-identification by soft multilabel learning. In *CVPR*, 2019. 2
- [59] Shijie Yu, Shihua Li, Dapeng Chen, Rui Zhao, Junjie Yan, and Yu Qiao. COCAS: A large-scale clothes changing person dataset for re-identification. In *CVPR*, 2020. 1, 2
- [60] Yunpeng Zhai, Shijian Lu, Qixiang Ye, Xuebo Shan, Jie Chen, Rongrong Ji, and Yonghong Tian. Ad-cluster: Augmented discriminative clustering for domain adaptive person re-identification. In *CVPR*, 2020. 2
- [61] Ziyuan Zhang, Luan Tran, Xi Yin, Yousef Atoum, Xiaoming Liu, Jian Wan, and Nanxin Wang. Gait recognition via disentangled representation learning. In *CVPR*, 2019. 6
- [62] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, 2015. 1, 7
- [63] Liang Zheng, Yi Yang, and Alexander G Hauptmann. Person re-identification: Past, present and future. *arXiv preprint arXiv:1610.02984*, 2016. 1
- [64] Zhedong Zheng, Xiaodong Yang, Zhiding Yu, Liang Zheng, Yi Yang, and Jan Kautz. Joint discriminative and generative learning for person re-identification. In *CVPR*, 2019. 1, 6
- [65] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3D human reconstruction from a single image. In *ICCV*, 2019. 5
- [66] Zhedong Zheng, Liang Zheng, and Yi Yang. A discriminatively learned CNN embedding for person reidentification. *TOMM*, 2017. 6
- [67] Zhedong Zheng, Nenggan Zheng, and Yi Yang. Parameter-efficient person re-identification in the 3D space. *arXiv preprint arXiv:2006.04569*, 2020. 1