

Cloning Outfits from Real-World Images to 3D Characters for Generalizable Person Re-Identification

Yanan Wang¹

yanan.wang.cs@gmail.com

Xuezhi Liang²

xz.liang.cs@gmail.com

Shengcai Liao^{1*}

scliao@ieee.org

Inception Institute of Artificial Intelligence (IIAI)¹
 Mohamed bin Zayed University of Artificial Intelligence²
 Masdar City, Abu Dhabi, UAE

Abstract

Recently, large-scale synthetic datasets are shown to be very useful for generalizable person re-identification. However, synthesized persons in existing datasets are mostly cartoon-like and in random dress collocation, which limits their performance. To address this, in this work, an automatic approach is proposed to directly clone the whole outfits from real-world person images to virtual 3D characters, such that any virtual person thus created will appear very similar to its real-world counterpart. Specifically, based on UV texture mapping, two cloning methods are designed, namely registered clothes mapping and homogeneous cloth expansion. Given clothes keypoints detected on person images and labeled on regular UV maps with clear clothes structures, registered mapping applies perspective homography to warp real-world clothes to the counterparts on the UV map. As for invisible clothes parts and irregular UV maps, homogeneous expansion segments a homogeneous area on clothes as a realistic cloth pattern or cell, and expand the cell to fill the UV map. Furthermore, a similarity-diversity expansion strategy is proposed, by clustering person images, sampling images per cluster, and cloning outfits for 3D character generation. This way, virtual persons can be scaled up densely in visual similarity to challenge model learning, and diversely in population to enrich sample distribution. Finally, by rendering the cloned characters in Unity3D scenes, a more realistic virtual dataset called *ClonedPerson* is created, with 5,621 identities and 887,766 images. Experimental results show that the model trained on *ClonedPerson* has a better generalization performance, superior to that trained on other popular real-world and synthetic person re-identification datasets. The *ClonedPerson* project is available at <https://github.com/Yanan-Wang-cs/ClonedPerson>.

*Shengcai Liao is the Corresponding Author.

1. Introduction

The generalization of person re-identification has gained increasing attention in recent years. One way to improve generalization is to develop large-scale and diverse training datasets. However, collecting person images from surveillance videos is privacy sensitive, and the further data annotation is expensive. Therefore, recently, synthetic person re-identification datasets have been actively developed due to their advantages of no privacy concern and no annotation cost [3, 4, 25]. For example, RandPerson [29] automatically creates large-scale random 3D characters with 8,000 identities, rendered from simulation of surveillance environments in Unity3D [27]. It is also proved in [29] that large-scale synthetic datasets are very useful to improve generalization. Similar findings are also observed in the following work UnrealPerson [33]. However, synthesized persons in existing datasets are quite different from realistic persons, because synthesized persons are mostly cartoon-like and dress in random collocation. This clear domain gap limits the performance of models trained on such synthetic datasets.

On the other hand, some researchers proposed to generate 3D human body models from real-world person images [13, 28, 34], targeting at high-fidelity reconstruction. These methods try to generate 3D body shapes and the associated textures simultaneously, through deep neural networks. They can help reduce the gap between synthetic and realistic person images to some extent due to the input of real-world clothes textures. However, current methods are still not satisfactory as the results are usually blur, and there are many artifacts, e.g. in back views (see Fig. 8c).

Considering the above, in this work, an automatic approach is proposed to directly clone the whole outfits from real-world person images to virtual 3D characters. By doing so, we would like to achieve two goals¹: (1) the directly

¹However, high-fidelity 3D reconstruction of person bodies, for example, heads and 3D body shapes, is not our target. On the other hand, high-

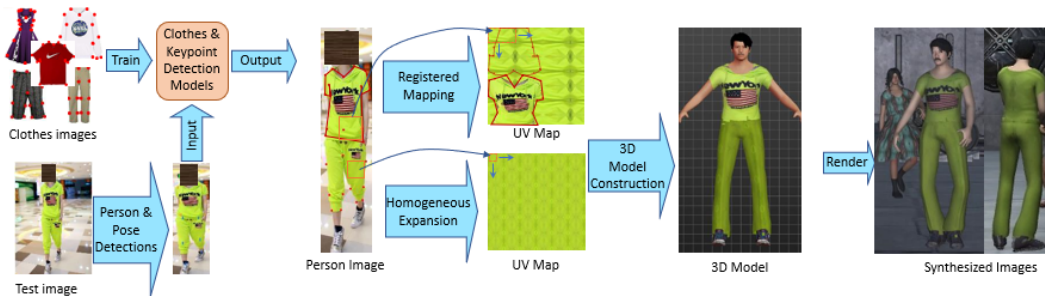


Figure 1. The proposed ClonedPerson pipeline, which automatically creates similarly dressed 3D characters from person images.

cloned clothes textures are clear and sharp in looking; and (2) by cloning the whole outfit, the virtual person thus created will appear very similar to its real-world counterpart, in similar clothes and dress collocation. Specifically, inspired from the UV texture mapping [5] method developed in RandPerson [29], in this work, two cloning methods for UV maps are designed, namely registered clothes mapping and homogeneous cloth expansion. Registered mapping targets at regular UV maps where clothes appear in regular shapes and structures. Based on clothes keypoints detected on real-world person images and labeled on UV maps, registered mapping applies perspective homography [26] to warp real-world clothes to the counterparts on the UV map. Homogeneous expansion is for invisible clothes parts and irregular UV maps. An optimization algorithm is proposed to find a large homogeneous area on clothes, use it as a realistic cloth pattern or cell, and expand the cell to fill the UV map. Fig. 1 shows the pipeline of the proposed method.

Furthermore, a general principle is established to scale up virtual 3D character creation, that is, it should expand both densely in similarity and diversely in population. The former one is to challenge discriminative model learning by providing similar persons, while the latter is to enrich the diversity in sample space. A similarity-diversity expansion strategy is thus proposed. Thanks to the proposed clothes cloning method, this can conveniently be achieved by clustering real-world person images and a controlled sampling of the clustered images for 3D character generation.

Eventually, the generated 3D characters are imported into Unity3D virtual environments to render a more realistic virtual dataset, called ClonedPerson, with 763,953 images from 4,826 characters for training, and 123,813 images of 795 characters for testing. Experimental results show that the similarity-diversity expansion strategy is effective, and the model trained on the ClonedPerson dataset has a better generalization performance, surpassing the models trained on various real-world and synthetic datasets.

In summary, our main contributions are: (1) We propose

fidelity reconstruction of identifiable biometric signatures, e.g. faces, may also raise privacy concerns.

Dataset	#ID	#Cam	#BBox	Sur	Real	Outfit
SOMAsset [4]	50	250	100,000	No	No	No
SyRI [3]	100	280	56,000	No	No	No
PersonX [25]	1,266	6	273,456	No	No	No
RandPerson ² [29]	8,000	19	1,801,816	Yes	No	No
UnrealPerson ³ [33]	6,799	34	1,256,381	Yes	Yes	No
ClonedPerson	5,621	24	887,766	Yes	Yes	Yes

Table 1. Statistics of some synthetic person re-identification datasets. “Sur”: surveillance simulation. “Real”: realistic clothes textures. “Outfit”: cloning full-body outfits from person images.

an automatic pipeline to clone outfits from real-world person images to virtual 3D characters, such that they look very similar to their real-world counterparts, with clear clothing textures; (2) two cloning methods, registered clothes mapping and homogeneous cloth expansion, are designed to fulfill this task; (3) a similarity-diversity expansion strategy is proposed, based on clustering of person images and controlled sampling, to scale up 3D character creation densely in similarity and diversely in population; and (4) a large-scale synthetic dataset called ClonedPerson, with 887,766 images of 5,621 characters, is created, which results in a better generalization performance than other popular real-world and virtual person re-identification datasets. All used and designed methods are listed in Table A of the Appendix.

2. RELATED WORK

Collecting and manually labeling real-world person re-identification datasets are expensive and privacy-sensitive. In contrast, the use of synthetic data can reduce the cost of manual labeling, and synthetic datasets do not have privacy issues. For synthetic datasets, SyRI [3] and PersonX [25] used limited hand-made characters to generate data. In contrast, RandPerson [29] proposed a clever way to generate new-looking clothes models by replacing UV maps of exist-

²Suggested subset from [29]: 8,000 characters with 132,145 images.

³Suggested subset from [33]: 3,000 characters with 120,000 images.

ing 3D clothes models with neutral images or random color and texture patterns, and designed an automatic pipeline in MakeHuman [6] to scale up character generation. Besides, similar to real-world environments, [29] simulated camera networks in Unity3D to render and record moving person videos. Moreover, [29] proved that models trained on synthetic data generalize well on real-world datasets. Following RandPerson, UnrealPerson [33] improved the accuracy by using real-world person images to create virtual characters, and rendering with the powerful Unreal Engine 4 (UE4) [1] with four large and realistic scenes. Specifically, it cropped blocks from segmented clothing images to directly replace UV maps of existing 3D clothes models. However, as shown in Fig. 7, this way still results in unrealistic-looking characters due to scale alignment issue. Statistics of some synthetic datasets are shown in Table 1.

On the other hand, one may consider using virtual try-on methods to generate synthesized persons. These methods aim to transfer a target clothing onto a reference person. However, existing virtual try-on methods [11,31] are mostly in 2D, which cannot generate 3D clothed human models, and thus cannot import them into virtual environments for comprehensive rendering. On the other hand, some existing methods, e.g. PIFu [23], targets at high-fidelity reconstruction of 3D persons from 2D images. However, such methods require ground-truth of 3D shapes and textures for training, which is quite expensive and limited in scale. Recently, some methods, e.g. HPBTT [34] tried training 3D reconstruction models from only 2D images. However, they are based on generative models, which usually result in blurred textures and artifacts. Besides, Pix2Surf [21] proposed to transfer texture from clothing images to 3D humans by neural networks. It achieved a good quality by training a specific model for every category of clothes. However, extending to other categories is costly. Furthermore, since HPBTT and Pix2Surf are both based on SMPL [20], they are not able to handle long skirts, as shown in Fig. 8.

Therefore, to further reduce the gap between virtual characters and realistic persons, we follow the way of RandPerson in repainting UV maps of existing 3D clothing models. However, different from RandPerson and UnrealPerson which directly replace existing UV maps by other images, we design two cloning methods for structure-aware, fine-grained repainting of UV maps.

3. 3D Virtual Character Generation

3.1. Pipeline Overview

Fig. 1 shows the pipeline of the proposed ClonedPerson approach, which includes the following steps. Firstly, we apply pre-processing steps, including pedestrian detection, pose detection, clothes detection, and clothes keypoint detection to get qualified frontal-view person images and ob-

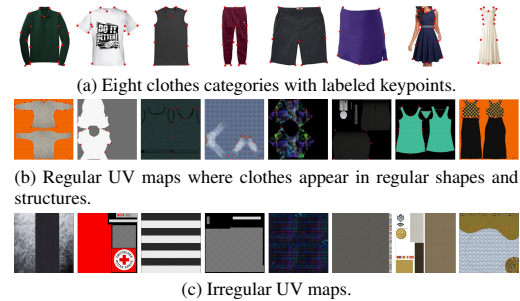


Figure 2. Different categories of clothes and UV texture maps of the corresponding 3D clothes models.

tain the clothes positions, categories, and clothes keypoints. Next, two cloning methods, registered mapping and homogeneous expansion, are applied to clone clothes from person images to UV texture maps and generate 3D characters. Finally, following RandPerson [29], these characters are imported into Unity3D to render synthesized person images.

Several pre-processing steps are implemented to fulfill our target. For example, to clone the full-body outfits from real-world person images to virtual 3D characters, we apply person detection to localize full-body person images, and remove standalone and occluded clothes. Besides, we design a number of rules based on pose detection to cherry-pick⁴ non-occluded frontal-view person images, since they best show clothing patterns and collocations. Due to space limits, pre-processing steps are introduced in Appendix B.

3.2. Registered Clothes Mapping

With 3D clothes models available in the MakeHuman community, we obtain some clothes models with regular UV maps, where clothes appear in regular shapes and structures, as Fig. 2b shows. With these regular UV maps, we apply perspective homography [26] to map real-world clothes textures to UV maps of 3D characters, so that the original texture structures in the clothes can be well kept, and will appear to be clear and sharp.

3.2.1 Perspective homography

Perspective homography is also known as perspective transformation [2,26]. Given a set of 2D points $\{\mathbf{p}_i\}$ and a corresponding set of points $\{\mathbf{p}'_i\}$, augmented with homogeneous coordinates (appending 1 as the z coordinate), perspective homography maps each \mathbf{p}_i to \mathbf{p}'_i by a homography matrix $\mathbf{H} \in \mathbb{R}^{3 \times 3}$, that is, $\mathbf{p}' = \mathbf{H}\mathbf{p}$.

Then, we can compute the homography matrix \mathbf{H} by solving the following optimization problem:

⁴No need to worry about dropping other images including some images in good conditions, since there are huge available sources.

$$\min_{\mathbf{H}} \sum_{i=1}^n \|\mathbf{p}'_i - \mathbf{H}\mathbf{p}_i\|_2^2, \quad (1)$$

where n is the number of the corresponding points. Eq. (1) defines a least squares problem and thus can be easily solved. In addition, the computed homography matrix \mathbf{H} can be refined with the Levenberg-Marquardt method [14] to further reduce the re-projection error.

3.2.2 Perspective warping

In our task, given labeled clothes keypoints \mathbf{p}_i on regular UV texture maps (e.g. Fig. 2b), and the corresponding detected clothes keypoints \mathbf{p}'_i on real-world person images (e.g. Fig. 2a), we can solve Eq. (1) and get the homography matrix \mathbf{H} . Then, each pixel location \mathbf{p} on the UV map will have a corresponding pixel location on the input image, by $[x, y, z]^T = \mathbf{H}\mathbf{p}$. Besides, we need to set the z coordinate of all the resulting points to 1 before the warping process, as the transformation operates on homograph coordinates. That is, $\mathbf{p}' = \begin{pmatrix} x \\ y \\ z \end{pmatrix}$, where \mathbf{p}' represents the corresponding point on the clothes image. Finally, the perspective warping can be done by bilinear interpolation on the clothes image and use the resulting pixel values to fill the UV map [26]. Specifically, by traversing \mathbf{p} on UV map in turn, a corresponding point $\mathbf{p}' = \mathbf{H}\mathbf{p}$ with float numbers of coordinates on the clothing image will be determined. Then, four pixels around \mathbf{p}' will be bilinearly interpreted into \mathbf{p} . An example is shown in Fig. 3, where the red dots represent the corresponding keypoints. To reduce background influence, we set the outer part of the clothes in black.

In this paper, with most regular UV maps we directly calculate perspective homography through all the clothes keypoints. However, as shown in Fig. 2b, the shapes of the long-sleeved and trousers on the UV map are quite different from those usually appear in person images. In these cases, we calculate the perspective homography on each part of them separately, then warp the clothes parts to the UV texture map and combine the results. For example, pants could be split into left and right sides.

3.3. Homogeneous Cloth Expansion

Registered clothes mapping can handle the clothes texture of the frontal side very well. However, the back side is usually different from the frontal side but invisible. Therefore, we further design the homogeneous cloth expansion method to find a homogeneous area on clothes as a realistic cloth cell, and expand the cell to fill the UV map. Besides, as Fig. 2c shows, the UV texture maps of some 3D clothes models are irregular, with unclear clothes structures. This also prevents the application of the registered clothes mapping. Therefore, we use the homogeneous cloth expansion

to handle irregular UV maps, and enable more clothes models and styles. In our experiments, we have 161 3D clothes models with regular UV maps, and 17 models with irregular UV maps, as illustrated in Fig. 2.

3.3.1 Cloth segmentation

The homogeneous cloth expansion includes two steps, cloth segmentation and cloth expansion. For cloth segmentation, an optimization algorithm is proposed to find a large homogeneous area on clothes, and use it as a realistic cloth cell. As shown in Fig. 3, we first crop the clothes area, and use a model trained on MSMT17 [30] by QAConv 2.0 [18] to extract the layer2 feature map (48×16) of this clothes image. The purpose of the QAConv model is to extract discriminant feature maps to find homogeneous cells, so as to reduce the influence of image noises. Then, square blocks of various scales are defined on the feature map. Within each block, the average and standard deviation of the feature values are computed, as follows:

$$\mu^k = \frac{1}{n_k} \sum_{i=1}^{n_k} \mathbf{x}_i^k, \quad \sigma_j^k = \sqrt{\frac{1}{n_k - 1} \sum_{i=1}^{n_k} (x_{ij}^k - \mu_j^k)^2}, \quad (2)$$

where k denotes the k th block, n_k is the number of elements in that block, $\mathbf{x}_i^k \in \mathbf{R}^d$ is the feature vector of the i th element in block k , with $d = 512$ dimensions, and j denotes the j th dimension. Note that the standard deviation is computed per feature channel. This value estimates the variations within each block, and thus reflects how homogeneous the cloth is within that block. Furthermore, we would also like the selected block to be as large as possible. Therefore, we further compute the area A_k of each block k , and define a ratio R as our objective function for the optimization problem, as follows:

$$\min_{k=1}^K R_k = \frac{\frac{1}{d} \sum_{j=1}^d \sigma_j^k}{A_k}, \quad (3)$$

where K denotes the number of blocks. By optimizing the above objective, we obtain a cloth area, with textures within it as homogeneous as possible, and with the area as large as possible. Then, we locate this block on the input clothes image and crop it, resulting in a patch which we call cloth cell. Appendix Fig. I shows some cloth cells thus obtained.

3.3.2 Cloth Expansion

As described above, the homogeneous cloth expansion is applied for both regular UV maps and irregular UV maps. For regular UV maps, it is used to fill the back side of the clothes area, as well as the background. Since we already apply the registered clothes mapping for the frontal side of

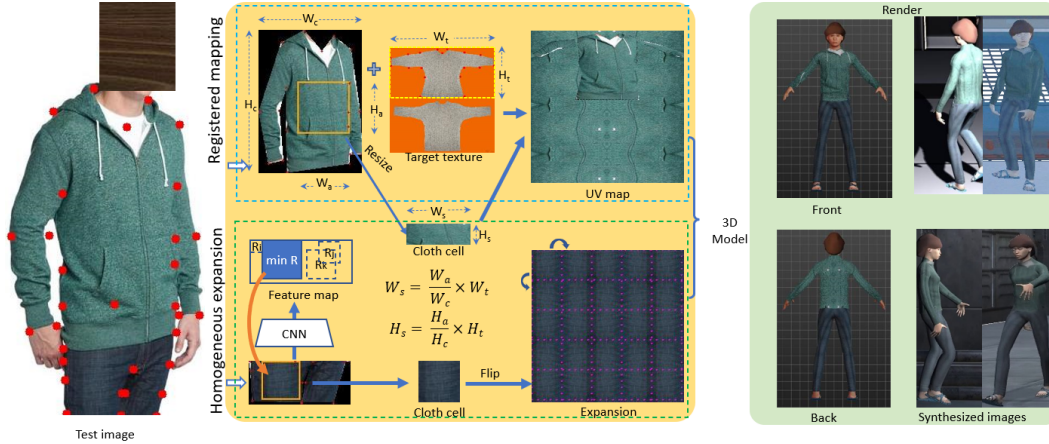


Figure 3. Process pipeline of registered clothes mapping (top) and homogeneous cloth expansion (bottom).

the clothes on regular UV maps, there exists a scale alignment problem for the cloth cell to be filled on the same UV map. Therefore, to maintain the consistency of the texture of the clothes, we need to scale the homogeneous cloth cell. As shown in Fig. 3, let W_c and H_c be the width and height, respectively, of the clothes image, W_a and H_a be the width and height, respectively, of the cropped cloth cell from the clothes image, W_t and H_t be the width and height, respectively, of the target area of the clothes after registered mapping, then, W_s and H_s , the width and height, respectively, of the cell to be scaled can be computed as follows:

$$W_s = \frac{W_a}{W_c} \times W_t, \quad H_s = \frac{H_a}{H_c} \times H_t. \quad (4)$$

Then, the scaled cloth cell is expanded on the whole UV map besides the registered clothes mapping area, by alternating flipping and tiling. As for irregular UV maps, since there is no reference of the scale, we simply use the original shape of the homogeneous cloth cell to flip and tile until fully fill the whole UV map, as shown in Fig. 3.

Note that besides the homogeneous cloth expansion, a simple way is to resize the cloth cell directly as a UV map, as done in RandPerson and UnrealPerson. However, simply resizing the cloth cells may result in blur textures and unrealistic patterns, as compared in Appendix D.

4. Similarity-Diversity Expansion

We use both DeepFashion (Apache License 2.0) [19] and DeepFashion2 [9] images for our virtual data creation. Through the pre-processing steps, there are still tens of thousands of images that are qualified and can be cloned to virtual characters. However, because of the enormous volume of images and repeating images of the same person, using these images directly to create characters is not efficient. To address this, two principles are considered. First,



Figure 4. Illustration of similarity-diversity expansion.




Figure 5. Clustering of different ϵ values. (a) Examples of two clusters (up and down) with $\epsilon=0.4$. (b) Examples of two clusters (left and right) with $\epsilon=0.5$, with person images (first row) and the generated characters (second row).

the more diverse samples, the better generalization performance should be. Second, similar person images are able to make the model training pay more attention to subtle differences. According to [33], similar characters as *hard samples* take a positive effect for person re-identification with large number of identities and cameras.

Therefore, we propose a similarity-diversity expansion strategy to scale up virtual character creation while improving along both the similarity and diversity aspects, as illustrated in Fig. 4. By clustering person images, we can create similar characters from the same cluster, while increase diversity by including more and more clusters. This way, the created characters can expand densely in similarity and diversely in population. Specifically, this strategy first applies DBSCAN [7] to cluster person images, then it samples a

certain number of images per cluster, and finally clones outfits from these images for 3D character generation. In this way, we can generate similar characters in the same cluster and diverse characters with different clusters⁵.

For the clustering, we use the same model trained on MSMT17 [30] by QAConv 2.0 [18] to extract feature maps and compute similarity scores between person images. Then, DBSCAN is applied, with different ϵ parameters to control the degree of similarity. Specifically, to remove repeating persons, we set $\epsilon=0.4$ to cluster the same person with the same outfits. Fig. 5a shows two examples where images from the same cluster are with the same person. Next, we select one image per cluster (closest to the cluster center) and remove other redundant images. Together with other images failed to be clustered (with label -1), the second round of clustering is performed, with $\epsilon=0.5$. As shown in Fig. 5b, this time images from the same cluster are visually similar but generally from different persons.

Finally, we select seven images (five for training and two for testing) per cluster to generate characters. Following RandPerson [29], these characters are imported into Unity3D to render synthesized person images. We implement some adjustments to improve the rendering, as detailed in Appendix E. Accordingly, we create 5,621 characters with 887,766 images, as the ClonedPerson dataset, with 763,953 images from 4,826 characters for training, and 123,813 images from 795 characters for testing. The Statistics of the dataset are shown in Table 1. We summarize the details and statistics of each step in our pipeline in Appendix F. Fig. 4 and Fig. 7c (more in the Appendix) show some characters created in ClonedPerson.

5. EXPERIMENTS

5.1. Datasets

Three real-world person re-identification datasets, CUHK03 [15], Market-1501 [35], and MSMT17 [30], are used for generalization evaluation. The CUHK03 dataset includes 14,097 images of 1,467 individuals. There are 7,365 images of 767 identities in the training set, and 6,732 images of 700 identities in the test set, according to the CUHK03-NP protocol [36]. The detected bounding boxes are used. The Market-1501 dataset includes 32,668 images of 1,501 identities captured from six cameras. 12,936 images of 751 identities are included in the training set, and the remaining 19,732 images of 750 identities are used for the test set. The MSMT17 dataset includes 126,441 images of 4,101 identities and divided into 32,621 images of 1,041 identities for training, and the remaining 93,820 images of 3,060 identities for testing.

We use two other synthetic datasets, RandPerson [29]

⁵Note that we only create one 3D character for one person image, since generating similar characters is already considered in the same cluster.

and UnrealPerson [33], for comparison, since they are shown to be superior to other synthetic datasets for generalizable person re-identification. RandPerson contains 8,000 identities and 1,801,816 images with 19 cameras. Both the full set and the suggested subset (132,145 images of the 8,000 identities) are used for our experiments. Besides, since some rendering setups are modified in this work, we further render RandPerson characters in our conditions for a fair comparison. This is denoted by RandPerson* (RP*). UnrealPerson releases 6,799 characters with 1,256,381 images. Both the full set and the suggested subset (120,000 images of 3,000 identities) are used for our experiments.

5.2. Methods

The validation of the proposed ClonedPerson is through person re-identification experiments. We mainly consider two tasks, generalizable person re-identification [12, 16, 32], and unsupervised domain adaptation (UDA). We apply QAConv 2.0 [18] and TransMatcher [17] for the former, and SpCL [10] for the latter. All of them are under the MIT License. We keep the same settings for each method.

All evaluations follow the single-query evaluation protocol [8]. We use the Cumulative Matching Characteristic (CMC) [22], especially the Rank-1 accuracy, and the mean Average Precision (mAP) [24] as the performance metrics.

5.3. Generalizable Person Re-Identification

The mAP results of direct cross-dataset evaluation are shown in Table 2 comparing real-world datasets with QAConv 2.0, and Table 4 comparing synthetic datasets with QAConv 2.0 and TransMatcher. Rank-1 results are reported in Appendix Table C. In overall, ClonedPerson achieves the best performance, surpassing existing datasets of both synthetic and real-world. The better performance over existing real-world datasets further confirms the findings in [29] and [33]. Besides, ClonedPerson is much better than UnrealPerson on CUHK03 and Market-1501, while they are comparable on MSMT17. Note that scenes used by UnrealPerson are more large and realistic than ours, due to the powerful UE4 engine. Besides, UnrealPerson has ten more cameras than ClonedPerson. Note also that, by comparing to both full set and subset results of RandPerson and UnrealPerson, it is clear that ClonedPerson’s better performance is not because it is larger, but because of its capability of cloning the whole outfits from person images.

Moreover, by comparing RandPerson* to RandPerson, the new rendering settings are more effective. Besides, compared to RandPerson* with the same rendering setting, ClonedPerson has gained an averaged improvement of about 2% in mAP. This is encouraging since ClonedPerson has only 4,826 identities, compared to 8,000 in RandPerson.

Furthermore, with the same learned QAConv models in Tables 2 and 4, we also evaluate their performances on

Dataset	CUHK03-NP		Market-1501		MSMT17	
	Rank-1	mAP	Rank-1	mAP	Rank-1	mAP
CUHK03	-	-	65.5	34.5	42.3	13.4
Market-1501	15.1	15.1	-	-	41.0	14.6
MSMT17	18.5	19.2	76.3	47.9	-	-
ClonedPerson	22.6	21.8	84.5	59.9	49.1	18.5

Table 2. Direct cross-dataset evaluation results of models trained on different datasets with QAConv 2.0.

Training Data	QAConv		SpCL	
	Rank-1	mAP	Rank-1	mAP
CUHK03	29.1	2.8	11.5	1.0
Market-1501	40.3	5.9	12.0	1.1
MSMT17	39.8	6.3	10.2	0.9
ClonedPerson	91.1	68.9	10.6	0.9

Table 3. Evaluation results on the ClonedPerson testing set with different tasks. Green: Cross-dataset evaluation. Gray: Within-dataset evaluation. Blue: UDA. Pink: Unsupervised Learning.

the ClonedPerson testing set, with results shown in Table 3. First, with within-dataset evaluation, QAConv achieves 91.1% in Rank-1 and 68.9% in mAP, indicating that this synthetic domain can be reasonably fitted by a representative method. Second, with cross-dataset evaluation, it can be seen that all models trained on real-world datasets perform not satisfactory on ClonedPerson, indicating that ClonedPerson is quite different and challenging. Nevertheless, it can still be observed that MSMT17 and Market-1501 are more diverse for generalization than CUHK03.

5.4. Unsupervised Domain Adaptation

As for UDA, we conducted experiments with SpCL, using ClonedPerson as source training data or target testing data. Since the whole training set of ClonedPerson is too large for SpCL to handle, e.g. in its clustering stage, we also selected a training subset of ClonedPerson for SpCL, with one image per camera, and 75,830 images in total from the 4,826 subjects. With ClonedPerson as source training data, the results are shown in Table 4. It shows that on average ClonedPerson outperforms both RandPerson and UnrealPerson, especially, with a large margin on CUHK03.

With ClonedPerson as target dataset, the results are shown in Table 3. Similar as cross-dataset evaluation, the UDA results on ClonedPerson are also poor. Furthermore, we also conduct an unsupervised learning task on ClonedPerson by SpCL, as shown in Table 3. Again, the results are poor. Therefore, it appears that, for SpCL, real-world source training data does not help much in domain adaptation to ClonedPerson, and thus the poor performance

Method	Dataset	#ID	#Imgs	CUHK	Market	MSMT	Avg
QAConv	RP	8,000	1,801k	16.0	46.9	14.0	25.6
	RP	8,000	132k	15.1	45.9	13.8	24.9
	RP*	8,000	1,239k	20.1	56.4	17.6	31.4
	UP	6,799	1,256k	17.2	56.1	17.5	30.3
	UP	3,000	120k	17.8	55.9	19.3	31.0
	CP	4,826	763k	21.8	59.9	18.5	33.4
TransM	RP	8,000	1,801k	18.7	49.6	16.4	28.2
	RP	8,000	132k	16.9	49.0	15.8	27.2
	RP*	8,000	1,239k	22.9	58.0	20.9	33.9
	UP	6,799	1,256k	19.7	60.2	18.4	32.8
	UP	3,000	120k	19.6	59.4	21.6	33.5
	CP	4,826	763k	24.4	62.3	20.8	35.8
SpCL	RP	8,000	132k	4.7	67.2	27.2	33.0
	UP	3,000	120k	5.3	71.7	28.4	35.1
	CP	4,826	75k	12.0	72.7	24.2	36.3

Table 4. mAP results with different datasets for different tasks. TransM: TransMatcher. RP: RandPerson. RP*: RandPerson* in new rendering settings. UP: UnrealPerson. CP: ClonedPerson.

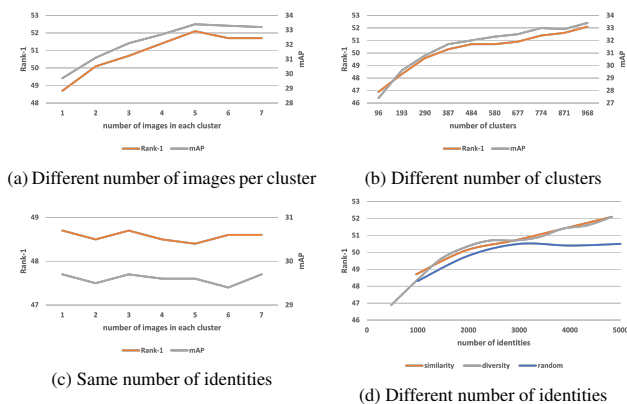


Figure 6. Performance of different character scaling up methods.

is mainly due to the unique challenges in ClonedPerson for clustering-based identity label reasoning. For example, there are a large number of diverse cameras and a lot of similar persons created by the proposed similarity-diversity expansion strategy. Consequently, though ClonedPerson is a synthesized dataset, it may provide a good test bed for both domain generalization and domain adaptation, and challenge researchers in developing more effective algorithms.

5.5. Comparison of Different Generation Settings

Fig. 6 shows performance (averaged Rank-1 and mAP of the three real-world testing datasets) with different character scaling up methods, including different settings of the proposed similarity-diversity expansion strategy, and a straightforward random scaling up strategy.

After the clustering procedure described in Sec. 4, we obtain 968 clusters. Then, first, we use all the clusters for maximum diversity, and select different numbers of images per cluster for experiments, indicating increasing similarity. Fig. 6a shows the performance. As the number of selected images increases, the performance clearly increases. Therefore, it proves that creating similar persons is indeed important for discriminant model learning, since it has to pay more attention to fine details of characters. However, it is saturating when the number of images per cluster reaches five. Therefore, to avoid data redundancy and improve efficiency, we select five images per cluster for the training set, and treat the remaining as a separate testing set.

Next, we keep the similarity level consistent, with five images per cluster, and select different numbers of clusters for experiments, indicating increasing diversity. As Fig. 6b shows, the performance increases as the number of clusters increases, which aligns with our expectations that performance raises as the diversity increases.

However, the adjustment of diversity and similarity will inevitably cause changes in identities, which might influence performance. Therefore, we keep the number of identities consistent by balancing the variation of similarity and diversity. That is, the selected clusters are gradually reduced when the number of images per cluster increases. The results are shown in Fig. 6c. The performance fluctuates in a small range within 0.3%, indicating that both similarity and diversity are important in our virtual data creation.

Finally, we also compare the random creation method that randomly selects person images for texture cloning and character creation before our clustering step. Fig. 6d shows the comparison of this random creation method to our strategies in Fig. 6a and Fig. 6b. From the results it is clear that with random creation after 3,000 characters the performance is saturating. In contrast, the proposed similarity-diversity expansion strategy is much more efficient in scaling up the virtual character creation, especially with larger number of identities.

Therefore, the above analysis proved that the similarity-diversity expansion is effective and efficient in scaling up the virtual character creation, and is potentially useful in creating an even larger and effective dataset when more person image sources are considered, considering the trend in Fig. 6b. In contrast, in UnrealPerson [33] the conclusion is that it can only achieve the best performance with 3,000 characters, but more characters do not help. This is also verified in our experiments with UnrealPerson in Tab. 4.

5.6. Qualitative Comparisons

Fig. 7 shows some images of characters created by three different methods, RandPerson, UnrealPerson, and the proposed ClonedPerson. As can be seen, RandPerson is the most cartoon-like. As for UnrealPerson, though it also uses



Figure 7. Examples from different synthetic datasets.




Figure 8. Qualitative comparison of different synthesis methods.

real clothes textures, most of its created characters do not match real-life clothes due to the scale alignment issue of cloth patterns. In contrast, thanks to the designed cloning pipeline, the ClonedPerson characters are more realistic and dressed more like real-life persons.

Furthermore, Fig. 8 shows a qualitative comparison between models created by our method and HPBTT [34]. It can be observed that characters created by the proposed method have clearer and sharper clothes textures, and better back-view looking of the clothes, than that generated by HPBTT. Besides, from the results shown in the first row, it can be seen that in ClonedPerson the clothes category is preserved, while HPBTT fails to deal with long skirts.

6. CONCLUSION

This paper contributes an automatic approach to clone the whole outfits from real-world person images to virtual 3D characters. Two critical cloning methods are proposed, registered clothes mapping and homogeneous cloth expansion. As a result, these characters bridge the gap between synthesized and realistic persons, and so models trained by our synthesized persons have better generalization ability for person re-identification. In addition, a similarity-diversity expansion strategy is proposed to scale up virtual characters. We show that similarity can help improve model’s discrimination, while diversity can improve the generalization ability of the model. In the future, we could exploit more in developing different types of clothes models and exploit more data sources. Moreover, we show some limitations of this research in the Appendix.

References

- [1] <https://www.unrealengine.com/en-US/>. 3
- [2] Alex M Andrew. Multiple view geometry in computer vision. *Kybernetes*, 2001. 3
- [3] Slawomir Bak, Peter Carr, and Jean-Francois Lalonde. Domain adaptation through synthesis for unsupervised person re-identification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 189–205, 2018. 1, 2
- [4] Igor Barros Barbosa, Marco Cristani, Barbara Caputo, Aleksander Rognhaugen, and Theoharis Theoharis. Looking beyond appearances: Synthetic training data for deep cnns in re-identification. *Computer Vision and Image Understanding*, 167:50–62, 2018. 1, 2
- [5] Christian Chang. *Modeling, UV Mapping, and Texturing 3D Game Weapons*. Wordware Publishing, Inc., 2006. 2
- [6] MakeHuman Community. *MakeHuman: Open Source Tool for Making 3D Characters*, 2020. <http://www.makehumancommunity.org>. 3
- [7] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. Density-based spatial clustering of applications with noise (dbscan). In *Proc. of the Second International Conference on Knowledge Discovery and Data Mining*, pages 226–231, 1996. 5
- [8] Michela Farenzena, Loris Bazzani, Alessandro Perina, Vittorio Murino, and Marco Cristani. Person re-identification by symmetry-driven accumulation of local features. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2360–2367. IEEE, 2010. 6
- [9] Yuying Ge, Ruimao Zhang, Lingyun Wu, Xiaogang Wang, Xiaoou Tang, and Ping Luo. A versatile benchmark for detection, pose estimation, segmentation and re-identification of clothing images. *CVPR*, 2019. 5
- [10] Yixiao Ge, Feng Zhu, Dapeng Chen, Rui Zhao, and Hongsheng Li. Self-paced contrastive learning with hybrid memory for domain adaptive object re-id. In *Advances in Neural Information Processing Systems*, 2020. 6
- [11] Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, and Larry S Davis. Viton: An image-based virtual try-on network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7543–7552, 2018. 3
- [12] Yang Hu, Dong Yi, Shengcai Liao, Zhen Lei, and Stan Z Li. Cross dataset person re-identification. In *Asian Conference on Computer Vision*, pages 650–664. Springer, 2014. 6
- [13] Verica Lazova, Eldar Insafutdinov, and Gerard Pons-Moll. 360-degree textures of people in clothing from a single image. In *2019 International Conference on 3D Vision (3DV)*, pages 643–653. IEEE, 2019. 1
- [14] KENNETH LEVENBERG. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, 1944. 4
- [15] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 152–159, 2014. 6
- [16] Shengcai Liao and Ling Shao. Interpretable and generalizable person re-identification with query-adaptive convolution and temporal lifting. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 456–474. Springer, 2020. 6
- [17] Shengcai Liao and Ling Shao. Transmatcher: Deep image matching through transformers for generalizable person re-identification. *Advances in Neural Information Processing Systems*, 34, 2021. 6
- [18] Shengcai Liao and Ling Shao. Graph Sampling Based Deep Metric Learning for Generalizable Person Re-Identification. *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 4, 6
- [19] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 5
- [20] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. 3
- [21] Aymen Mir, Thiemo Alldieck, and Gerard Pons-Moll. Learning to transfer texture from clothing images to 3d humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7023–7034, 2020. 3
- [22] P. Jonathon Phillips, Patrick Grother, and Ross Micheals. Evaluation methods in face recognition. In S. Z. Li and A. K. Jain, editors, *Handbook of Face Recognition*, chapter 21, pages 551–574. Springer, 2nd edition, Sep. 2011. 6
- [23] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2304–2314, 2019. 3
- [24] T. Sobh and K. Elleithy. *Innovations in Computing Sciences and Software Engineering*. Springer Netherlands, 2010. 6
- [25] Xiaoxiao Sun and Liang Zheng. Dissecting person re-identification from the viewpoint of viewpoint. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 608–617, 2019. 1, 2
- [26] Richard Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010. 2, 3, 4
- [27] Unity Technologies. *Unity3D: Cross-platform game engine*, 2020. <https://unity.com>. 1
- [28] Jian Wang, Yunshan Zhong, Yachun Li, Chi Zhang, and Yichen Wei. Re-identification supervised texture generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11846–11856, 2019. 1
- [29] Yanan Wang, Shengcai Liao, and Ling Shao. Surpassing real-world source training data: Random 3d characters for generalizable person re-identification. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 3422–3430, 2020. 1, 2, 3, 6
- [30] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *Proceedings of the IEEE conference on*

computer vision and pattern recognition, pages 79–88, 2018. 4, 6

- [31] Han Yang, Ruimao Zhang, Xiaobao Guo, Wei Liu, Wang-meng Zuo, and Ping Luo. Towards photo-realistic virtual try-on by adaptively generating-preserving image content. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 7850–7859, 2020. 3
- [32] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Deep metric learning for person re-identification. In 2014 22nd International Conference on Pattern Recognition, pages 34–39. IEEE, 2014. 6
- [33] Tianyu Zhang, Lingxi Xie, Longhui Wei, Zijie Zhuang, Yongfei Zhang, Bo Li, and Qi Tian. Unrealperson: An adaptive pipeline towards costless person re-identification. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 11506–11515, 2021. 1, 2, 3, 5, 6, 8
- [34] Fang Zhao, Shengcai Liao, Kaihao Zhang, and Ling Shao. Human parsing based texture transfer from single image to 3d human via cross-view consistency. Advances in Neural Information Processing Systems, 33, 2020. 1, 3, 8
- [35] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In Proceedings of the IEEE international conference on computer vision, pages 1116–1124, 2015. 6
- [36] Zhun Zhong, Liang Zheng, Donglin Cao, and Shaozi Li. Re-ranking person re-identification with k-reciprocal encoding. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1318–1327, 2017. 6