

When Person Re-identification Meets Changing Clothes

Fangbin Wan^{1*†} Yang Wu^{2*¶} Xuelin Qian^{3‡} Yixiong Chen^{1‡} Yanwei Fu^{1†§}

¹School of Data Science, Fudan University ²Kyoto University

³School of Computer Science, Fudan University

Abstract

Person re-identification (ReID) is now an active research topic for AI-based video surveillance applications such as specific person search, but the practical issue that the target person(s) may change clothes (clothes inconsistency problem) has been overlooked for long. For the first time, this paper systematically studies this problem. We first overcome the difficulty of lack of suitable dataset, by collecting a small yet representative real dataset for testing whilst building a large realistic synthetic dataset for training and deeper studies. Facilitated by our new datasets, we are able to conduct various interesting new experiments for studying the influence of clothes inconsistency. We find that changing clothes makes ReID a much harder problem in the sense of bringing difficulties to learning effective representations and also challenges the generalization ability of previous ReID models to identify persons with unseen (new) clothes. Representative existing ReID models are adopted to show informative results on such a challenging setting, and we also provide some preliminary efforts on improving the robustness of existing models on handling the clothes inconsistency issue in the data. We believe that this study can be inspiring and helpful for encouraging more researches in this direction. The dataset is available on the project web-site: <https://wanfb.github.io/dataset.html>.

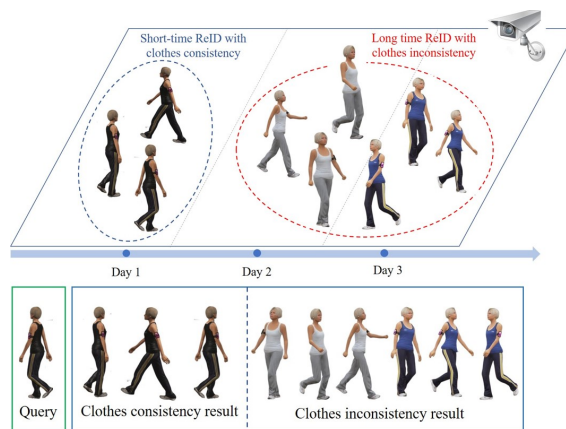


Figure 1. The appearance of the same person across different days. There usually exists clothes consistency in a short time-span, while in a longer period, the person usually changes his/her clothes. Given a query image, ReID shall target to find both the clothes consistency and clothes inconsistency results.

1. Introduction

Automatically searching for a specific person (e.g., a suspect) across multiple video surveillance cameras at different places is a popular imagination in AI-based science fiction movies and TV dramas. It is not only extremely valuable and impactful for public safety and security (e.g., finding the Boston Marathon bombing suspect efficiently), but also technically highly challenging. The key component of this task is referred to as person re-identification (ReID) in the literature. It is about fine-grained search over numerous similar-looking candidates while at the same time having to tolerate significant appearance changes of the same targets.

The identity of a person is usually determined by his/her biological traits/characteristics, such as facial appearance, ages, body shape and so on, rather than the outlook appearance like hairstyles, bags, shoes, and clothes *et al.* In contrast, the widely used benchmarks in ReID – Market1501 [28], CUHK03 [14], DukeMTMC-reid [30]) do not really contain significant appearance variance of the same person, which can easily mislead ReID models to learn appear-

*indicates equal contributions.

†indicates corresponding author.

‡Fangbin Wan, Xuelin Qian and Yixiong Chen are with Shanghai Key Lab of Intelligent Information Processing, Fudan University. Email: {fbwan18, xlqian15, 16307110231}@fudan.edu.cn.

§Yanwei Fu is with MOE Frontiers Center for Brain Science, and Shanghai Key Lab of Intelligent Information Processing, Fudan University. Email: yanweifu@fudan.edu.cn.

¶Yang Wu is with Department of Intelligence Science and Technology, Graduate School of Informatics, Kyoto University. Email: wu.yang.8c@kyoto-u.ac.jp.

ance features (*e.g.*, clothes) rather than the desired robust identity-sensitive features related to biological traits. This is particularly true, in the sense that, most successful DNN-based ReID models are learning from data. The clothes consistency on the same identities makes the models greatly rely on the visual appearance related to clothes, which usually takes the largest part of a human body. Thus, these models can easily fail when people wear similar clothes as others' or change their own clothes.

Though people do change clothes frequently, especially during different days, and crime suspects may even intentionally do that on the same day for more effective hiding, the **clothes inconsistency** problem has long been overlooked in the research community. There are mainly two reasons for that. Firstly, a traditional assumption about ReID is that it only covers a short period of a few minutes/hours, so that there only exists **clothes consistency** for the same identities. However, even in such a short period, one can still take on or take off his/her coat, grasp a bag, or get dressed in another suit, let alone in long periods like across days. This assumption greatly limits the generalization ability of ReID models towards real-world applications, such as finding crime suspects. To this end, it is imperative to investigate the ReID models in the settings that allow clothes inconsistency, as shown in Fig. 1. Secondly, it is very expensive to create a large-scale person ReID dataset with a significant amount of same-identity clothes inconsistency in real scenarios. Capturing person image or video data may not be as easy as before, since surveillance data is now under strict control and individual privacy is a highly valued issue in many countries. Besides, associating the same identities with clothes changes in such large amounts of data for getting ground-truth labels is even more challenging for human annotators. These difficulties discourage the study of person ReID with changing clothes to some extent. However, there are still very few efforts on this direction, *e.g.* RGB-D dataset [4] and the PRCC dataset [27]. Although these datasets enable some preliminary research of the clothes inconsistency issue in person ReID, they are quite limited in both the size and collected environment.

Despite these difficulties, in this paper, we provide benchmark datasets and systematically study the influence of clothes inconsistency problem for person ReID, which is the first time as far as we know. In more details, two benchmark datasets are built and released. One is a 3-day 4-camera dataset capturing real persons and scenarios around the entrance of a big office building (as shown in the upper row of Fig. 2). Due to the privacy concerns, currently we have only been able to capture data of 28 volunteers. Such difficulties motivated us to build another dataset – a synthetic dataset called **Virtually Changing-Clothes (VC-Clothes)** dataset using 3D human models with the help of a video game engine named GTA5 (samples shown in

Fig. 3). Thanks to the encouraging breakthroughs of computer graphical techniques and the prevalence of virtual reality games, we could get high-quality realistic data covering 512 identities of 19,060 images in 4 different scenes with significant clothes changes.

The clothes inconsistency problem makes person ReID much more challenging, which can be easily imagined. However, it is never validated and discussed experimentally in previous studies. In Sec. 4, we show that clothes changes not only significantly influence representation learning, but also challenge the generalization ability of ReID models in re-identifying the persons wearing unseen clothes. Moreover, we present some preliminary efforts on how to promote performance by learning more identity-sensitive and clothes-insensitive features, which can serve as baselines for future deeper studies.

Contributions. (1) We systematically study the influence of the same-identity clothes inconsistency issue for ReID, which is historically overlooked. (2) We build two new benchmark datasets, one real and one synthetic with significant clothes changes, for supporting the research. (3) We provide preliminary solutions for handling the clothes inconsistency issue, which can be baselines for inspiring further researches.

2. Related Work

2.1. Clothes Inconsistency in ReID.

There are few studies about the clothes inconsistency problem in person ReID. The biggest challenge is the lack of sufficient applicable and well-annotated data with clothes inconsistency. RGB-D [4] is the first applicable ReID dataset proposed to solve this problem using additional depth information. However, the depth cue is only applicable in indoor environments and the RGB-D dataset is not sufficient to train a deep learning model (with only 79 identities and 1-2 clothes per identity). iQIYI-VID [16] is a large scale multi-modal dataset collected from 600K video clips of 5,000 celebrities of various types. Although this dataset demonstrates great advantages in terms of size and variation in human poses, face quality, clothes, makeup, *etc.*, the data are collected from Internet videos, instead of surveillance cameras. Another possible way is to automatically generate images with clothes differences using Generative Adversarial Networks, *e.g.* DG-Net [29]. The encoders decompose each pedestrian image and project them into two latent spaces: an appearance space and a structured space. By doing so, one can change the color of clothes by switching the two latent space of different persons. However, this work did not dig into the importance of clothes-independent features, and the appearance information still contributes most to the performance improvement. Moreover, it can only change the color of clothes without altering



Figure 2. Some samples of “Real28” collected in different days of the same person and “RGB-D” captured in different days but grouped by different identity of persons.



Figure 3. Samples of VC-Clothes. The first row is the four environment maps, the following two rows are the same persons under the 4 cameras with at most 3 suits of clothes.

human poses or illuminations at the same time. The recently introduced PRCC [27] dataset contains 33,698 real human images of 221 different identities, which may be sufficient for research. However, it is only collected in indoor laboratory environment against simple backgrounds. Different from all the above works, we propose a large synthetic dataset which can better represent the real demands.

2.2. Adopting Synthetic Data.

Synthetic data is becoming increasingly popular in training deep learning models [22], because of the difficulty in building an ideal dataset in real environment and the requirements of huge human labor resources for data annotation. Comparatively, synthetic data can be generated in an automated way, therefore, promotes the development of many computer vision tasks like: semantic segmentation [18, 21], object detection [19, 9], pose estimation [1, 5, 6] *et al.* In ReID, only a few works [2, 3, 24] utilize synthetic data. Two of them [2, 24] are designed to investigate illumination varieties or viewpoint diversities in ReID instead of clothing changes. Although the other one SOMaset [3]

considers clothing changes, it does not discuss the influence of them in depth. Differently, our two datasets contain 512 identities with 1~3 different suits of clothes for each identity, and they can better serve the research on the clothes inconsistency issue in ReID.

3. New Benchmark Datasets

Current person ReID datasets are not sufficient for in-depth researches as they either lack clothes inconsistency or limit in data diversities such as missing outdoor scenarios. Therefore, we decided to build new benchmark datasets. However, it is difficult to build a large-scale dataset of real surveillance data with significant clothes changes. Thus, we built a small one with 28 student volunteers for just preliminary research and put more efforts in building a much larger dataset with realistic synthetic data generated by a powerful 3D video game engine. In this paper, we show that these two datasets are sufficient for supporting many valuable researches on the historically overlooked clothes inconsistency issue with our initial studies. Firstly, we give a brief introduction to the two datasets, especially the synthetic one, which is unique and good for developing new robust ReID models. We also prove its potential to serve as a new person ReID benchmark for advancing the related researches.

3.1. The Real28 Dataset

We name our real-scenario dataset “Real28”. As can be seen in Fig. 2 with the representative samples, the dataset is collected in 3 different days (with different clothing) by 4 cameras. It consists of totally 4,324 images from 28 different identities with 2 indoor scenes and 2 outdoors. Since the size is not big enough for training deep learning models, we suggest only using this dataset for performance evaluation and having the models trained on some other datasets, like the following synthetic dataset VC-Clothes. We only split



Figure 4. Examples of people with different clothes, various occlusions, background and illumination, as well as pose variance.

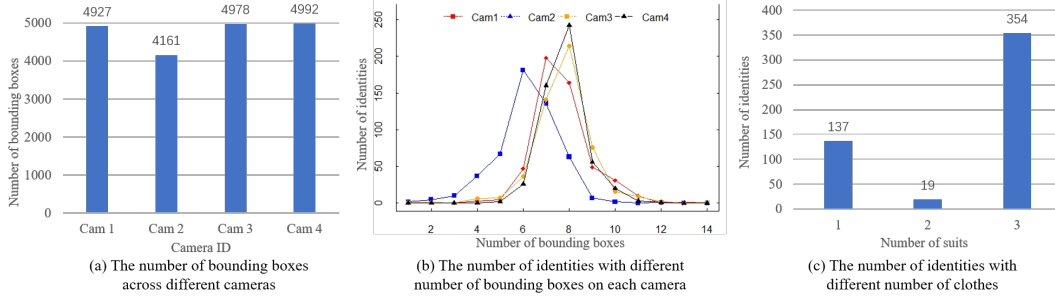


Figure 5. Statistical data of VC-Clothes. We show that the numbers of bounding boxes (i.e., person instances) from each camera are balanced (sub-figure (a)). The average number of images per identity are between 6 to 8 for each camera (sub-figure (b)). Totally, 354 people have 3 suits of clothes and 137 have only 1 suit of clothes (sub-figure (c)). All these numbers demonstrate the research values of our VC-Clothes dataset on the clothes changeable ReID task.

the whole dataset into a query set and a gallery set. We randomly select an image of each identity in different scenes and different dates as the query image, while the other images are treated as the gallery images. As a result, there are totally 336 images in the query set and 3,988 images in the gallery set.

3.2. The VC-Clothes Dataset

3.2.1 Dataset construction

The images are rendered by the Grand Theft Auto V¹ (GTA5), which is a famous action-adventure game, with HD and realistic game graphics. Critically, it allows very convenient configurations of the clothes of each avatar/character and supports user-defined environmental parameters, such as those for illumination, viewpoint and background. Fabbri *et al.* [6] even made a pedestrian tracking and pose estimation dataset on top of it. GTA5 has a very large environment map, including thousands of realistic buildings, streets and spots. We select 4 different scenes – 1 indoor scene and 3 outdoor scenes, *i.e.*, gate, street, natural scene and parking lot, virtually captured by 4 cameras (Cam1, Cam2, Cam3, Cam4), as shown in Fig. 3. We introduce additional illumination by changing not only the illumination of different scenes, but also the same scene under different time, to make the images more realistic. In general, outdoor scenes have more lighting than indoors.

In the process of image generation, each person walks

along a scheduled route, and cameras are set up and fixed at the chosen locations. We have different camera views of front-left, front-right, and right-after, thus the human faces may not be easily observed. However, each person still has very distinctive biological traits, as demonstrated in Fig. 4. Thus, it is imperative to identify these persons by the features of face, gender, age, body figure, hairstyle and so on, rather than the clothes. Notice we only change the clothes of persons, while maintaining their identities unchanged. In our VC-Clothes dataset, each identity has 1~3 suits of clothes and the clothes of each identity are kept unchanged when captured by the same camera. Furthermore, to facilitate experiments, we make all the identities wear the same clothes under Cam2 and Cam3.

The main challenge of this dataset comes from the visual appearance of different clothes of the same identity. Other challenges also include different persons with similar clothes, various occlusion, illumination changes under different environments, and pose variance (see Fig. 4)

3.2.2 Dataset properties

VC-Clothes has 512 identities, 4 scenes (cameras) and on average 9 images/scene for each identity and a total number of 19,060 images. Mask RCNN [10] is employed to detect persons and automatically produce bounding boxes. The detailed statistics of VC-Clothes are summarized in Table. 1, in comparison with other widely used ReID datasets, including DukeMTMC-reid [30], Market1501 [28], CUHK03 [14], CUHK01 [13] and VIPeR [7]. More statistical analysis

¹<https://www.rockstargames.com/>

Table 1. Statistical comparison between VC-Clothes and other popular ReID datasets.

Dataset	DukeMTMC-reid	Market1501	CUHK03	VC-Clothes	CUHK01	VIPeR
No. of BBoxes	36,411	32,668	28,192	19,060	3,884	1,264
No. of Identities	1,812	1,501	1,467	512	971	632
No. of Cameras	8	6	2	4	10	2
Detector	hand	DPM	DPM, hand	Mask RCNN	hand	hand
Scene	outdoor	outdoor	indoor	indoor, outdoor	indoor	outdoor
Clothing change	X	X	X	✓	X	X

Table 2. The results of state-of-the-art methods on VC-Clothes and CUHK03 with labeled bounding boxes. Note that * indicates that PCB adopts a new training/testing protocol proposed in [31], thus yields worse performance than other methods.

Methods	VC-Clothes				CUHK03			
	mAP	R@1	R@5	R@10	mAP	R@1	R@5	R@10
LOMO+XQDA[15]	36.9	42.3	53.1	57.7	-	52.2	82.2	92.1
GOG+XQDA[17]	37.8	42.9	55.9	61.8	-	67.3	91	96
Baseline	47.4	50.1	72.6	80.2	-	68.9	79.7	84.5
MDLA[20]	76.8	88.9	92.4	93.9	-	76.8	96.1	98.4
PCB*[25]	74.6	87.7	91.1	92.8	57.5	63.7	80.6	86.9
Part-aligned[23]	79.7	90.5	94.3	96.7	-	88.0	97.6	98.6

of this dataset can be seen in Fig. 5.

We equally split the dataset by identities: 256 identities for training and the other 256 for testing. We randomly chose 4 images per person from each camera as query, and have the other images serve as gallery images. Eventually, we get totally 9,449 images in the training, 1,020 images as queries and 8,591 others in the gallery.

3.2.3 The benchmark

To make VC-Clothes dataset a qualified benchmark, we compare the results of some representative ReID methods on VC-Clothes with a representative benchmark, CUHK03.

Problem definition. We still follow the classical definition of person ReID task. Given a training dataset $D_{train} = \{\mathbf{I}_i, y_i\}_{i=1}^N$ of N images, where \mathbf{I}_i and y_i are the image data and its corresponding label, the goal of ReID is to learn an effective feature mapping $\Phi(\mathbf{I})$ for the input image \mathbf{I} , such that images of the same identity have smaller distances in the feature space than those of different identities. In VC-Clothes, we target to find both the clothes consistency results and clothes inconsistency results at the same time.

Evaluation metrics. We use two main evaluation metrics: the Mean Average Precision (mAP) and the accuracy values at the first few indicative ranks, *e.g.*, Rank-1 (in abbr. R@1), Rank-5 (R@5), and Rank-10 (R@10). MAP refers to the average accuracy rate of each relevant document retrieved by a query. The Rank- i accuracy means whether the first i query result contains the correct image.

Competitors. We compare several hand-crafted methods, including LOMO+XQDA [15], GOG+XQDA [17]; DNN methods: MDLA [20], PCB [25], Part-aligned [23]. We use the original setup of the 5 state-of-the-art methods

in our experiments. In addition, we also provide a baseline model using the widely adopted ResNet50 [11] network pre-trained on the ImageNet [12] dataset. During training, the batch-size is set to 128 and the model is trained for 200 epochs with an initial learning rate of 0.0001, which will decay by a factor of 0.01 every 50 epochs.

Benchmark results. As shown in Table 2, the performances trend of different methods on the VC-Clothes dataset is similar as on CUHK03, showing that VC-Clothes can reflect real-world performance. It also enables different design of sub settings, thus we can systematically study the influence of clothes-inconsistency problem. All of the above factors make VC-Clothes a rational benchmark for person ReID model evaluation.

4. Influence of Clothing Change

When designing experimental settings, we mainly consider whether people change their clothes or not, and have all identities maintaining their clothes over two cameras (Cam2 and Cam3) and changing their clothes over any other camera pairs. Such a setting enables the study of how can clothes inconsistency influences the performance of ReID.

4.1. Making ReID a Much Harder Problem

We first conduct a typical experiment on justifying the performance of ReID on two indicative settings: **Same Clothes** (Cam2 and Cam3) vs. **Change Clothes** (Cam3 and Cam4), using all the representative state-of-the-art methods. We intentionally select the images from specific cameras in training and testing set to serve as the data for change clothes and same clothes settings. As can be seen in Table 3, clothes inconsistency brings about 30% to 50% per-

Table 3. Result differences between wearing the same clothes and changing clothes (using two camera pairs as an example).

Methods	Same Clothes (1 suit/person, Cam2&Cam3)				Change Clothes (2 suits/person, Cam3&Cam4)			
	mAP	R@1	R@5	R@10	mAP	R@1	R@5	R@10
LOMO+XQDA [15]	83.3	86.2	93.8	95.1	30.9	34.5	44.2	49.7
GOG+XQDA [17]	86.7	90.5	94.5	95.7	31.3	35.7	48.3	54.2
Baseline	74.3	79.3	92.8	94.6	32.4	36.4	57.4	65.0
MDLA [20]	93.9	94.3	95.7	96.3	60.8	59.2	67.3	73.5
PCB [25]	94.3	94.7	96.3	96.7	62.2	62.0	68.2	72.9
Part-aligned [23]	93.4	93.9	96.0	96.7	67.3	69.4	77.5	84.5

Table 4. Generalization to unseen clothes (training with 2 suits/person using Cam3&Cam4).

Methods	Query in seen suits (Cam2) Gallery (Cam3&Cam4)				Query in unseen suits (Cam1) Gallery (Cam3&Cam4)			
	mAP	R@1	R@5	R@10	mAP	R@1	R@5	R@10
Baseline	28.8	27.5	53.1	63.9	19.5	18.4	37.5	49.6
MDLA[20]	59.0	57.7	75.2	81.1	43.2	40.5	63.5	73.5
PCB[25]	73.9	72.3	83.8	87.8	55.6	53.9	71.2	77.4
Part-aligned[23]	70.5	68.9	84.7	89.3	57.9	55.6	78.8	85.8

formance drop to DNN-based methods, and even more to hand-crafted methods, in terms of mAP and Rank-1 accuracy. Although camera diversities may also influence the results to some extent, the primary cause of the dramatic performance decrease should be the changing of clothes.

This indicates that clothes inconsistency is essentially a very subtle problem, which unfortunately has been overlooked in previous works. This is easy to understand as clothes usually cover most of the body, which makes it hard to extract identity-sensitive but clothes-independent appearance features. Note that the performance drop of the Part-aligned method [23] is the least in all the compared methods, indicating finer and more detailed modeling of appearance is needed for handling clothing changes.

4.2. Changing Learned Representations

In Fig. 6, we give a visualization of the appearance maps and part maps of the Part-aligned model [23] under the “Same Clothes” setting (Cam2 and Cam3) and “Change Clothes” (Cam1 and Cam2) settings (seeing Tab. 3). For a given input image (left), the global appearance (center) and body parts (right) are represented using different colors. The same color implies similar features and brighter locations denote more significant body parts.

In the “Same Clothes” setting where there is no clothes inconsistency, the appearance maps of body parts are much brighter than that in the “Change Clothes” setting and the colors are more scattered. Obvious attention on faces and feet can be seen in the part map of the “Change Clothes” setting, which means learning identity-sensitive and clothes-independent representations require exploration of the details of smaller areas instead of the large body parts.

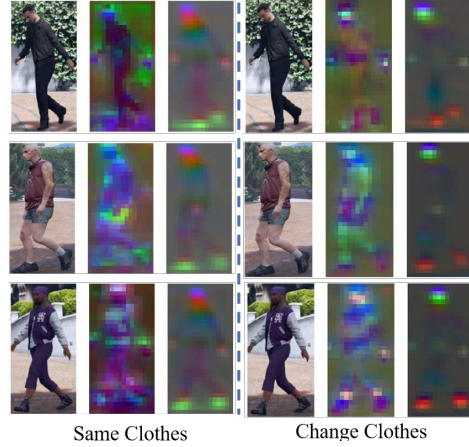


Figure 6. The feature maps of Part-Aligned method trained on “Same Clothes” setting (Cam2&Cam3) on the left and “Change Clothes” setting (Cam1&Cam2) on the right.

4.3. Generalization to Unseen Clothes

To evaluate the ReID models’ ability in handling clothes inconsistency of the same identities, training using the test data from the same camera pair (as shown in Tab. 3) is not enough, because the learned models may overfit specific camera viewpoints. Therefore, we propose to test by querying with data from another new camera.

In the VC-Clothes dataset, each camera corresponds to a type of suit for each identity. We can train models on Cam3 and Cam4, and test on two different cameras: Cam2 and Cam1. The two cameras represent two different settings. Choosing Cam2 means querying with human instances in **seen** suits, as the clothes of each person are the same for Cam2 and Cam3. Differently, Cam1 records a new **unseen**

Table 5. Different combination of models pre-trained and tested on “All clothes” or “Same clothes” settings.

Pre-trained	Test					
	All clothes			Same clothes		
	R@1	R@10	mAP	R@1	R@10	mAP
All clothes	90.8	96.7	82.9	92.9	96.1	92.4
Same clothes	63.1	69.7	43.1	93.9	96.7	93.4

suit for each identity. The two models are trained on the same setting but test on two different settings, thus making it possible to evaluate not only the performance on querying with a new camera, but also the generalization ability on handling unseen clothes, which is believed to be more important when clothing changes do happen.

The experimental results are presented in Table 4. Since DNN-based methods perform much better than hand-crafted ones, we only list the results of DNN-based methods. As the result shows, querying with a new camera is not a big issue when people keep their clothes unchanged, but wearing unseen suits do bring significant extra difficulties and limit the generalization ability of the models.

4.4. Benefits of Training with Clothes Inconsistency

To prove the advantage of models trained on the dataset with clothes inconsistency, we test the generalization ability of model pre-trained on “**All Clothes**” setting (images from all cameras) and “**Same Clothes**” setting (images for only Cam2 and Cam3) and make a confusion matrix of the two models tested on “All Clothes” setting and “Same Clothes” setting, respectively. Specifically, the “Same Clothes” setting imitates the general person ReID dataset, while the “All Clothes” setting is actually VC-Clothes dataset.

As shown in Table 5, the performance drop of model pre-trained on “All Clothes” setting and tested on “Same Clothes” setting is ignorable, about only 1% on VC-Clothes (in terms of R@1 and mAP, while R@10 has fewer differences), compared with the model pre-trained and tested both on “Same Clothes” setting. However, a dramatic decrease of performance appears when the pre-trained “Same Clothes” model is tested on the “All Clothes” setting (Rank-1 and mAP reduce by 27.7% and 39.8% respectively). Comparatively, models trained on the dataset with clothes inconsistency can perform well on both clothes consistency setting and clothes inconsistency setting, while models trained only on the dataset with only clothes consistency, does not show good generalization ability on clothes inconsistency setting.

5. Preliminary Solutions

In this section, we further discuss some preliminary solutions towards the robustness against clothes inconsistency. In particular, we present several straightforward attempts.

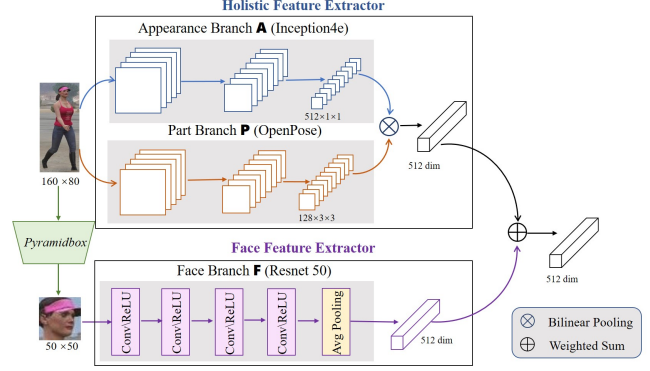


Figure 7. Overview of our model. The two 512 dimension feature vectors from holistic feature extractor and face feature extractor are concatenated using weighted sum.

We hope our experimental results can motivate and inspire further studies in such an important issue.

5.1. Enriching Training Data

One common strategy to improve the performance of DNN-based models is providing more training data. We treat the models trained on Cam3 and Cam4 as the basic settings, and add data from one more camera to enlarge the training set. By choosing either Cam1 or Cam2 as an additional camera, we can get two different settings. Adding Cam1 means adding one more unseen suit/person, while adding Cam2 means adding a seen suit/person, as Cam2 and Cam3 share the same clothes for each person.

The experimental results of these three different settings are presented in Table 6. Adding 1 unseen suit/person significantly boosts the performance of all methods, while the opposite effect can be observed when adding 1 seen suit/person. It may indicate that adding unseen suits can make the DNN-based models work harder on generating more clothes-independent representations, which is more effective for ReID with clothing changes. Therefore, adding more training data is proved to be effective when introducing new clothes.

5.2. Exploring Exposed Body Parts: the Faces

Overview of the fusion model. In order to improve the performance of current ReID models on top of VC-Clothes, we propose a 3-stream Appearance, Part and Face Extractor Network (3APF). As summarized in Fig. 7, it is composed of two main components: holistic feature extractor (ReID part) and a local face feature extractor (face part).

Holistic feature extractor. We directly utilize the Part-aligned network $f_1(\cdot)$ [23] as the holistic feature extractor following the original setting. In particular, this network learns the appearance $A(\cdot)$ and body part $P(\cdot)$ branches, which are further combined by a bi-linear pooling. Thus, this network computes the features as in Eq. 1.

Table 6. Improving the performance by enriching the training data.

Methods	Train: 2 suits/person (Cam3&Cam4) Test: 2 suits/person (Cam3&Cam4)				Train: add 1 unseen suit/person (add Cam1) Test: 2 suits/person (Cam3&Cam4)				Train: add 1 seen suit/person (add Cam2) Test: 2 suits/person (Cam3&Cam4)			
	mAP	R@1	R@5	R@10	mAP	R@1	R@5	R@10	mAP	R@1	R@5	R@10
Baseline	32.4	36.4	57.4	65.0	34.4	38.8	63.9	73.5	35.0	36.7	47.6	52.5
MDLA	60.8	59.2	67.3	73.5	62.6	63.5	73.1	79.2	54.2	55.1	61.6	68.2
PCB	62.2	62.0	68.2	72.9	66.9	66.3	73.3	79.2	58.1	57.5	64.7	70.6
Part-aligned	67.3	69.4	77.5	84.5	68.3	69.4	79.0	85.9	63.9	64.1	73.7	80.0

$$f_1(\mathbf{I}) = \text{pooling}(A(\mathbf{I}) \otimes P(\mathbf{I})) \quad (1)$$

Local face feature extractor. We utilize Pyramidbox [26] for face detection; and only 80% person images have detectable faces, because of non-frontal faces. The confidence of the face bounding boxes is set to 0.8. Further, we expand the detected bounding boxes of faces by 15 pixels in four directions (up, down, left and right, respectively) to cover more facial area. All face images are then resized to the same size of 50×50 pixels. As for the images with undetectable faces, we use only the feature from a holistic feature extractor for identification.

We use the widely adopted Resnet50 network $F(\cdot)$ [11] pre-trained on MS-1M [8] dataset to extract face features $f_2(\mathbf{I}) = F(G(\mathbf{I}))$, where G is a face detector. The final feature vector is the weighted mean result of two features as in Eq. 2, where w is the hyper-parameter to determine the proportion of each part.

$$f = w \cdot f_1 + (1 - w) \cdot f_2 \quad (2)$$

Improvement brought by faces The visualization of Fig. 6 demonstrates the significance of face, when there exists clothes inconsistency. To describe the influence of face to the ReID part in a quantitative way, we employ the 3APF model to control the proportion of face and ReID part by adjusting the weight. The larger the weight, the more important of ReID part in our model. As can be seen in Table. 7, when we adjust the weight to a certain value (0.95), a significant improvement on mAP (82.1%) can be witnessed over single face extractor and holistic extractor.

5.3. Values of VC-Clothes to Real Dataset

To evaluate whether our pre-trained models on VC-Clothes (synthetic dataset with clothes inconsistency) has some real-world applications and prove its advantages over Market1501 (real dataset without clothes inconsistency). We adopt 2 different ways to transfer the model pre-trained on VC-Clothes to real dataset Real28. One is directly applying the 3APF models trained on VC-Clothes or Market1501 to Real28 or RGB-D. Another way is using fine-tuning to learn better representations of real data with clothes inconsistency. “VC-Clothes+FT” means pre-trained on VC-Clothes and fine-tuned on Market1501, while “Market1501+FT” denotes the opposite.

Table 7. The results of our 3APF model on VC-Clothes.

Model	mAP	R@1	R@5	R@10
Face extractor	60.8	63.5	87.0	91.4
Holistic extractor	79.7	90.5	94.3	96.7
Ours	82.1	90.2	94.7	96.8

Table 8. Comparison of three different ways to transfer models to real test dataset with clothes inconsistency.

Method	Real28				RGB-D	
	Different clothes		All Clothes		All clothes	
	R@1	mAP	R@1	mAP	R@1	mAP
VC-Clothes	5.2	3.6	22.7	8.2	17.2	9.4
Market1501	8.2	6.1	48.1	19.1	27.9	20.0
VC-Clothes+FT	11.0	10.6	55.3	24.3	26.6	16.0
Market1501+FT	12.4	7.6	54.0	21.8	27.9	20.0

As can be seen in Table. 8, directly applying the model trained on VC-Clothes is hard to get good results, but we can get much better results (VC-Clothes+FT) by simply fine-tuning it on a third real dataset. Moreover, our synthetic data can also help model pre-trained on real datasets to generate better performance on the tasks with clothes inconsistency (Market1501+FT). This phenomenon is especially significant on “Real28”, where more clothes changes exist.

6. Conclusions and Discussions

In this paper, we studied the clothes inconsistency problem in person ReID. Such a setting is much realistic and helpful in tackling the longtime person ReID task. Because of the huge difficulty in collecting a large-scale dataset in real scenes, we just collect a small one for testing. Alternatively, we proposed to use the GTA5 game engine to render a large-scale synthetic dataset “VC-Clothes” with desired properties. With these two benchmark datasets, we conducted a series of pilot studies, which helps us better understand the influence and importance of this new problem. Finally, we also investigated some straightforward solutions for improving the performance on VC-Clothes and ways to transfer the model to the real test dataset. It is hopefully to motivate and inspire further studies on this important issue.

Acknowledgement. This work was supported in part by Microsoft Research Asia Collaborative Research Grant.

References

- [1] Mathieu Aubry, Daniel Maturana, Alexei A Efros, Bryan C Russell, and Josef Sivic. Seeing 3d chairs: exemplar part-based 2d-3d alignment using a large dataset of cad models. In *CVPR*, pages 3762–3769, 2014.
- [2] Slawomir Bak, Peter Carr, and Jean-Francois Lalonde. Domain adaptation through synthesis for unsupervised person re-identification. In *ECCV*, pages 189–205, 2018.
- [3] Igor Barros Barbosa, Marco Cristani, Barbara Caputo, Alexander Rognhaugen, and Theoharis Theoharis. Looking beyond appearances: Synthetic training data for deep cnns in re-identification. *CVIU*, 167:50–62, 2018.
- [4] Igor Barros Barbosa, Marco Cristani, Alessio Del Bue, Loris Bazzani, and Vittorio Murino. Re-identification with rgb-d sensors. In *ECCV*, pages 433–442. Springer, 2012.
- [5] Pau Panareda Busto, Joerg Liebelt, and Juergen Gall. Adaptation of synthetic data for coarse-to-fine viewpoint refinement. In *BMVC*, pages 14–1, 2015.
- [6] Matteo Fabbri, Fabio Lanzi, Simone Calderara, Andrea Palazzi, Roberto Vezzani, and Rita Cucchiara. Learning to detect and track visible and occluded body joints in a virtual world. In *ECCV*, 2018.
- [7] Douglas Gray and Hai Tao. Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *ECCV*, pages 262–275. Springer, 2008.
- [8] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. MS-Celeb-1M: A dataset and benchmark for large scale face recognition. Springer, 2016.
- [9] Hironori Hattori, Vishnu Naresh Boddeti, Kris M Kitani, and Takeo Kanade. Learning scene-specific pedestrian detectors without real data. In *CVPR*, pages 3819–3827, 2015.
- [10] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012.
- [13] Wei Li, Rui Zhao, and Xiaogang Wang. Human reidentification with transferred metric learning. In *ACCV*. Springer, 2012.
- [14] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deep-reid: Deep filter pairing neural network for person re-identification. In *CVPR*, 2014.
- [15] Shengcai Liao, Yang Hu, Xiangyu Zhu, and Stan Z Li. Person re-identification by local maximal occurrence representation and metric learning. In *CVPR*, 2015.
- [16] Yuanliu Liu, Peipei Shi, Bo Peng, He Yan, Yong Zhou, Bing Han, Yi Zheng, Chao Lin, Jianbin Jiang, Yin Fan, et al. iqiivid: A large dataset for multi-modal person identification. *arXiv preprint arXiv:1811.07548*, 2018.
- [17] Tetsu Matsukawa, Takahiro Okabe, Einoshin Suzuki, and Yoichi Sato. Hierarchical gaussian descriptor for person re-identification. In *CVPR*, 2016.
- [18] John McCormac, Ankur Handa, Stefan Leutenegger, and Andrew J Davison. Scenenet rgb-d: Can 5m synthetic images beat generic imagenet pre-training on indoor segmentation? In *ICCV*, pages 2678–2687, 2017.
- [19] Xingchao Peng, Baochen Sun, Karim Ali, and Kate Saenko. Learning deep object detectors from 3d models. In *ICCV*, pages 1278–1286, 2015.
- [20] Xuelin Qian, Yanwei Fu, Yu-Gang Jiang, Tao Xiang, and Xiangyang Xue. Multi-scale deep learning architectures for person re-identification. In *ICCV*, 2017.
- [21] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, pages 3234–3243, 2016.
- [22] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. Learning from simulated and unsupervised images through adversarial training. In *CVPR*, pages 2107–2116, 2017.
- [23] Yumin Suh, Jingdong Wang, Siyu Tang, Tao Mei, and Kyoung Mu Lee. Part-aligned bilinear representations for person re-identification. In *ECCV*, 2018.
- [24] Xiaoxiao Sun and Liang Zheng. Dissecting person re-identification from the viewpoint of viewpoint. In *CVPR*, pages 608–617, 2019.
- [25] 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.
- [26] Xu Tang, Daniel K Du, Zeqiang He, and Jingtuo Liu. Pyramidbox: A context-assisted single shot face detector. In *ECCV*, pages 797–813, 2018.
- [27] Qize Yang, Ancong Wu, and Wei-Shi Zheng. Person re-identification by contour sketch under moderate clothing change. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [28] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, 2015.
- [29] Zhedong Zheng, Xiaodong Yang, Zhiding Yu, Liang Zheng, Yi Yang, and Jan Kautz. Joint discriminative and generative learning for person re-identification. In *CVPR*, pages 2138–2147, 2019.
- [30] Zhedong Zheng, Liang Zheng, and Yi Yang. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *ICCV*, 2017.
- [31] Zhun Zhong, Liang Zheng, Donglin Cao, and Shaozi Li. Re-ranking person re-identification with k-reciprocal encoding. In *CVPR*, pages 1318–1327, 2017.