

UPAR: Unified Pedestrian Attribute Recognition and Person Retrieval

Andreas Specker^{1,2,3} Mickael Cormier^{1,2,3} Jürgen Beyerer^{2,1,3}

¹Karlsruhe Institute of Technology ²Fraunhofer IOSB ³Fraunhofer Center for Machine Learning

{firstname.lastname}@iosb.fraunhofer.de

Abstract

Recognizing soft-biometric pedestrian attributes is essential in video surveillance and fashion retrieval. Recent works show promising results on single datasets. Nevertheless, the generalization ability of these methods under different attribute distributions, viewpoints, varying illumination, and low resolutions remains rarely understood due to strong biases and varying attributes in current datasets. To close this gap and support a systematic investigation, we present UPAR, the Unified Person Attribute Recognition Dataset. It is based on four well-known person attribute recognition datasets: PA100K, PETA, RAPv2, and Market1501. We unify those datasets by providing 3.3M additional annotations to harmonize 40 important binary attributes over 12 attribute categories across the datasets. We thus enable research on generalizable pedestrian attribute recognition as well as attribute-based person retrieval for the first time. Due to the vast variance of the image distribution, pedestrian pose, scale, and occlusion, existing approaches are greatly challenged both in terms of accuracy and efficiency. Furthermore, we develop a strong baseline for PAR and attribute-based person retrieval based on a thorough analysis of regularization methods. Our models achieve state-of-the-art performance in cross-domain and specialization settings on PA100k, PETA, RAPv2, Market1501-Attributes, and UPAR. We believe UPAR and our strong baseline will contribute to the artificial intelligence community and promote research on large-scale, generalizable attribute recognition systems. The dataset is available here: https://github.com/speckean/upar_dataset

1. Introduction

Pedestrian analysis poses an essential task due to the fastly growing demand for intelligent video surveillance as well as online retailing. Relevant sub-tasks are Person Attribute Recognition (PAR), which aims to determine the semantic attributes of people given a person image, and attribute-based person retrieval, which searches for im-



Figure 1. **PAR datasets** – Sample images from different PAR datasets [4, 25, 22]. Each dataset shows different characteristics and scenarios. However, meaningful out-of-distribution evaluation was impossible due to a low number of common attributes across the datasets. Our UPAR dataset unifies 40 attributes and thus enables cross-domain investigations.

ages that show persons with the specified query characteristics. Semantic descriptions of a pedestrian’s attributes may serve as queries for attribute-based person retrieval or complementary information in common image-based re-identification in visual surveillance scenarios. In online retailing, attributes enable precisely filtering search parameters for clothing without requiring query images.

In recent years, deep learning-based visual models have shown impressive results in various tasks and also in PAR [2, 37, 15]. These models heavily rely on large-scale and well-annotated datasets. PAR datasets [43] offer different attribute annotations with different granularity. However, only a few attributes are identical across different datasets, which results in a lack of generalization analysis of PAR at a large scale. Moreover, the datasets are strongly biased regarding the captured scenarios. Some datasets only include indoor or outdoor scenes, the same weather conditions, a limited number of different cameras, and only one

ethnicity. As a result, models trained on these datasets tend to overfit and may not generalize well to out-of-distribution domains.

These limitations clearly indicate the demand for a large-scale and more comprehensive benchmark for generalizable PAR. To this aim, we construct the Unified Pedestrian Attribute Recognition and Person Retrieval (UPAR) dataset which is composed of the four popular datasets PA100K [25], PETA [4], RAPv2 [20], and Market1501 [22]. We unify those datasets by providing additional annotations to harmonize 40 important binary attributes over 12 categories and thus bring a large number of images covering extensive diversities w.r.t human subjects, data modalities, and capturing environments, as shown in Fig. 1. A total of 3.3M new binary color annotations are contributed to this dataset.

Besides extending the single datasets with previously unavailable attributes, we also enable detailed generalization evaluation for the first time. We propose two different evaluation protocols that represent different generalization scenarios. The first one includes multiple training datasets, while the more difficult one only allows one dataset for training and requires evaluation on several target domains.

In addition to introducing the UPAR dataset, we evaluate the performance of current CNN-based models on the unified dataset. However, due to the number of attributes and the diversity of the dataset, generalizable PAR and retrieval is challenging. Hence, we thoroughly examine design choices to construct a new strong baseline that achieves not only competitive results on the typical PAR task but also concerning attribute-based person retrieval on the UPAR dataset.

Contributions. UPAR contributes a standardized dataset to the community to launch investigations on realistic PAR in large-scale real-world PAR scenarios. Our contributions are summarized as follows.

- We propose a new unified dataset with 3.3M additional annotations to harmonize 40 important attributes across four datasets. UPAR is the first dataset that provides sufficient images and annotations to perform generalization experiments for PAR and attribute-based retrieval.
- We conduct a thorough analysis of design choices for generalizable PAR and attribute-based retrieval and provide detailed insights.
- We report state-of-the-art results on five PAR datasets and outperform more complex state-of-the-art approaches concerning PAR and attribute-based retrieval.

2. Related Work

Pedestrian Attribute Recognition. Research on PAR has shown significant progress over the last years. Early

deep learning-based works [17] propose considering the task a multi-label classification task. Moreover, the authors tackle the problem of imbalanced attribute distributions in the training dataset by introducing a weighted cross-entropy loss, which increases the loss for rare attributes. The JRL [40] method pursues a different approach and aims at leveraging attribute correlations through long-short-term-memory modules. A widely studied research direction is the use of attention mechanisms to focus on relevant regions and improve recognition of fine-grained attributes [25, 24, 19, 45, 13, 30]. It is also common to leverage features at multiple scales [47, 30, 38, 44, 49]. Most of the works build on large networks, which deteriorates practical applicability. Therefore, [16, 15] build on simple baseline models to achieve state-of-the-art PAR performance. While [16] introduces a strong baseline, [15] applies spatial and semantic regularization by exploiting that the same attribute is localized on similar positions across different images and attributes from the same semantic class with different sub-classes share similar visual features.

These works provide impressive results for yet single datasets [22, 25, 4, 17, 20], with high focus on one type of scenario, *e.g.* outdoor [22, 25] or indoor [17, 20], on comparatively medium-sized datasets [22, 4, 17]. However, only a few attributes have an equivalent definition across different datasets, such as gender, upper-body clothing length, or backpack, and thus enable experiments w.r.t. generalization on unseen domains. Therefore, researchers only have few options to document the generalization ability of their proposed approaches. In this work, we provide the first unified dataset to perform generalization experiments at a large scale for 40 binary attributes and diverse scenarios.

Person Retrieval. Researchers approach the task of attribute-based person retrieval in different ways, either by PAR or by learning a cross-modal feature space. Methods that follow the first approach train a PAR model and compare the predicted attribute confidence scores with the query attribute description [39, 31, 18, 32, 35, 33, 34, 8]. This approach provides semantics, and thus retrieval results are explainable. However, recognizing fine-grained, local attributes in low-resolution surveillance imagery is challenging. Further methods align attribute descriptions and image embeddings in a shared cross-modal feature space. Approaches from the literature solve this by using high-dimensional hierarchical embeddings and an additional matching network [5]. Further works aim to match person attributes and images in a joint feature space [46, 1]. Adversarial training is applied to align the different modalities. Jeong et al. [14] argue that this procedure is often unstable and challenging due to the min-max optimization procedure. Thus, they propose an approach that does not

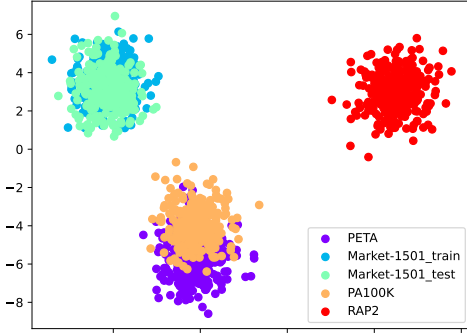


Figure 2. **UPAR Domains** – Distribution of image embeddings extracted using an ImageNet pre-trained Inception model. Train and test set embeddings from the same dataset overlap entirely. In contrast, our UPAR dataset as a combination of multiple sub-datasets with disjunct data distributions poses a more realistic and challenging problem and requires models to generalize well across different domains.

employ adversarial training but introduces a modality alignment loss function and a semantic regularization loss to leverage the relations between different attribute combinations explicitly [14].

Results in previous works [14] indicate that learning a shared embedding space between attributes and images achieves better performance for the attribute-based retrieval task. In contrast, our work clearly shows that state-of-the-art performance can be achieved by simply using a strong attribute classifier. Furthermore, semantics and explainability are provided and are often required from a real-world and application point of view.

3. The UPAR Dataset

Publicly available PAR datasets are heavily biased but also limited with respect to the diversity of the captured scenes. Except for PA100K [25] and RAPv2 [20], most of the existing public PAR datasets include less than 50,000 annotated pedestrians in either outdoor or indoor settings. For example, Market1501 only contains outdoor imagery captured during summer on campus. As a result, attributes are biased towards young Asian people wearing short summer clothing and thus do not reflect reality well. For this reason, findings on such datasets are limited and can hardly be transferred to other scenarios or applied in real-world applications. There are significant distribution shifts concerning attribute distributions as well as images across the datasets. The significant distribution shifts of the images are visualized in Fig. 2. It shows the distribution of image embeddings extracted using an ImageNet pre-trained Inception model. The outputs of the last feature layer are transformed into two-dimensional features using the Linear Discriminant Analysis method. One can observe that train and test set embeddings originating from the same dataset

overlap entirely, thus simplifying the PAR task greatly since models are only required to generalize well regarding comparatively small shifts of attribute distributions. In contrast, real-world applications require generalization to out-of-distribution domains, i.e., other datasets with different characteristics in our case. Out-of-distribution evaluation poses a much more complex and realistic task. To enable research on such realistic settings, we propose the UPAR dataset by unifying multiple existing PAR datasets concerning their attribute annotations and introducing two generalization evaluation protocols.

3.1. UPAR Composition

We construct our UPAR dataset by combining the images included in the four publicly available datasets PA100K [25], PETA [4], RAPv2 [20], and Market1501 [22, 48]. Due to several limitations regarding the creation of real-world surveillance datasets, particularly regarding privacy issues, we argue that already existing public datasets have much potential to be leveraged. Those datasets provide a large diversity concerning scenarios, ethnicities, and attributes, which makes UPAR less biased. The Market1501-Attribute [22, 48] dataset is an extended version of the Market1501 dataset augmented with the annotation of 27 attributes. The Market1501 dataset was collected outdoors in front of a supermarket using five high-resolution cameras and a low-resolution one. The dataset has 32,668 bounding boxes for 1,501 identities. PA100K [25] is currently the largest PAR dataset with 100,000 pedestrian images from various outdoor surveillance cameras with a large variance in image resolution, lighting conditions, and environments. The dataset provides annotation for 26 attributes. The Richly Annotated Pedestrian (RAP)[17, 20] dataset is currently the largest pedestrian attribute dataset composed of indoor scenes. There are two versions of this dataset. The first version, RAPv1 [17], was collected from a surveillance camera in shopping malls over three months and consists of 41,585 pedestrian images and 72 annotated attributes. RAPv2 [20] is an extension intended as a unified benchmark for both person retrieval and PAR in indoor real-world surveillance scenarios. The dataset contains 84,928 images for 2,589 person identities and 25 different scenes. The attributes annotated are the same as in RAPv1. Finally, the PEdesTrian Attribute (PETA) [4] dataset combines 19,000 pedestrian images from 10 publicly available datasets with indoor and outdoor scenes and 61 binary and four multi-class attributes. While widely used by the community for benchmarking PAR algorithms, those datasets have only a few attribute annotations in common, which significantly impairs generalization- and cross-evaluation.

To enable research on generalization, we unify 40 binary attributes over 12 categories between these datasets.

3.2. UPAR Attributes & Statistics

The 40 unified attributes were selected regarding their importance for person retrieval in video surveillance scenarios. Our annotations cover a variety of different attribute types to allow broad investigations. Global-scale attributes such as age or gender and small-scale attributes, e.g., glasses, are included in the dataset. There are attributes belonging to twelve categories: age, gender, hair length, upper-body clothing length, upper-body clothing color, lower-body clothing length, lower-body clothing color, lower-body clothing type, accessory-backpack, accessory-bag, accessory-glasses, and accessory hat. A complete list of the attributes is given in the supplementary material.

As a result, after cleaning images from the original datasets that could not be satisfyingly annotated, the UPAR dataset contains 224,737 images with annotations for 40 binary attributes over 12 categories. The dataset follows the splits from the original datasets for a total of 148,048 train, 30,830 validation, and 45,859 test images. We provide annotations for 11 unique colors (plus other and mixture) for 100,000 images and two attributes for a total of 2,57M binary color annotations. Furthermore, we provide annotation for the lower-body clothing length for the PETA and RAPv2 datasets. Annotation for glasses are contributed to PA100K, PETA and Market1501, age attribute to PETA and hair length to PA100K and Market1501 resulting in further 0.8M new annotations, for a total of 3,3M manually labeled and validated new binary annotations. With these new annotations, we enable multiple splits for leave-one-out evaluation to assess generalizability (see Tab. 1). Moreover, the UPAR can be used as a whole for the evaluation of specialization approaches.

We attempt to mitigate several issues regarding PAR datasets with our dataset. First, these are highly unbalanced regarding some attributes. In UPAR, even less represented attributes show at least 446 examples. However, the gap between the less and best-represented attributes remains significant. This is because of the data context and camera configuration in which the datasets were captured. The environment (market, mall, campus, etc.) greatly impacts the distribution of clothing and demographics. Publicly available datasets tend to have biases towards ethnicity and cultures, e.g., mostly representing western or south-east Asian cultures. We slightly mitigate this by unifying four large datasets. Future works should focus explicitly on representing larger and more diverse communities with diverse clothing. While not specifically addressing this, the increased diversity of our dataset tends to mitigate the involuntary learning of biometrics. Second, researchers are confronted with privacy issues and human rights violations when gathering a real-world dataset. From a research point of view, PAR datasets should be captured by real-world surveillance cam-



Figure 3. **Annotation challenges** – For the first three images, the real color of the clothing differs from the color seen. In the fourth image, the upper body is missing. Sometimes the primary person is difficult to tell, as can be seen in the first two images of the second row. Do leggings make the lower-body clothing long? (yes) Do long boots make the lower body clothing long? (no)

eras, with subjects unaware of being filmed and behavior as natural as possible. From an ethical perspective, however, pedestrians should consent before their images are annotated and publicly distributed. This important dilemma also touches already existing datasets. For example, the Duke-MTMC [29] and Celeb1M [10] datasets were withdrawn due to privacy violation issues. In this paper, we argue that already existing datasets are enough for our purposes and that these are not used to their full potential due to lacking annotation. Therefore, we believe that extending already existing and widely accepted datasets instead of acquiring more personal data is a viable strategy for scenarios such as PAR.

3.3. Data Annotation Process

The PA100K dataset does not provide color annotation, which is the reason why we asked 16 paid annotators to manually define the color of the upper-body clothing and lower-body clothing for eleven unique colors plus additional classes to indicate multiple colors or colors not included in the color list. The annotation is performed for each image individually, i.e., when describing the attributes of the same person in different images, only the visible information is taken into account. Thus, different attribute values may be assigned to the same person in different images depending on external factors such as lighting, shadows cast by objects or persons, or the clothing area visible in the frame, as illustrated in Fig. 3. In general, annotators are confronted with a variety of challenges. For instance, some images depict multiple persons and; thus, it is often difficult to define the primary person in the image due to occlusion. Additionally, colors might not be unambiguously assignable to unique color classes. In such cases, images are assigned to the "other" collection class, which, among others, includes metallic colors or the color beige. If the corresponding body area of the primary person is not visi-

Split ID	Leave-one-out		Cross-validation	
	Training	Evaluation	Training	Evaluation
0	PA100K, PETA, RAPv2	Market1501	Market1501	PA100K, PETA, RAPv2
1	Market1501, PETA, RAPv2	PA100K	PA100K	Market1501, PETA, RAPv2
2	Market1501, PA100K, RAPv2	PETA	PETA	Market1501, PA100K, RAPv2
3	Market1501, PA100K, PETA	RAPv2	RAPv2	Market1501, PA100K, PETA

Table 1. **UPAR splits** – Split definitions for the two UPAR generalization evaluation schemes. Cross-validation (CV) protocol is more challenging since only sub-dataset is used for training. In both cases, evaluation is performed on unseen domains.

ble in the frame or if the primary person is depicted in false colors, e.g., grayscale, an assignment to the class unknown is made, and the image is discarded from the UPAR dataset. For more info, we refer to the supplementary materials.

3.4. Benchmark and Metrics

We exploit two schemes to assess generalization ability on UPAR: Cross-validation (CV) and Leave-one-out (LOO). Different training and testing splits are given in Tab. 1.

The CV protocol is the more challenging one since it measures the generalization performance on multiple domains given training data from a single distribution. Metrics are calculated by evaluating the test datasets separately, computing the average across these test datasets for each split, and finally determining the mean and standard deviation across the splits. This procedure ensures that the datasets contribute equally to the final results and not proportionally to their number of test images.

The second generalization evaluation scheme uses LOO to simulate the case when training data from several domains are available, but data from the target distribution is missing. I.e., three of the four sub-datasets are leveraged for training and the remaining one for evaluation. Similar to the first protocol, there are four splits, and each sub-dataset is used for testing once. The final evaluation scores are the average and standard deviation across the four splits.

Last but not least, the dataset can be used in total to perform evaluation analogous to the original datasets. In this case, the training images and the test images from all sub-datasets are combined.

Typically, label-based and instance-based metrics are used to evaluate PAR approaches [20]. Label-based metrics, such as the Mean Accuracy, focus on the recognition accuracy of single attributes. The classification accuracy is computed separately for positive and negative samples of the attributes, and the mean is calculated for each attribute. Then, the mA is the average across all attributes. In contrast, instance-based metrics, i.e., the F1 score in our case, focus on the attribute description assigned to persons. The precision and recall values for each test set sample is calculated as the basis for the F1 score, which is the harmonic mean of those values. In this case, the average is computed over all instances in the test set. Concerning attribute-based person

retrieval, we rely on the typically used Mean Average Precision (mAP) and Rank-1 Accuracy (R-1) metrics. While the mAP reflects the quality of the entire retrieval ranking, the R-1 score is the accuracy of the first position in the ranking.

Note that attributes without a positive sample in the training data of the split are excluded. Furthermore, attribute queries for the retrieval are determined by computing all unique sets of attributes occurring in the respective test set.

4. Proposed Strong Baseline

This work aims to develop a strong baseline universally applicable to the PAR and attribute-based person retrieval tasks. In contrast to PAR, there is no simple and strong baseline specifically tailored and evaluated for attribute-based person retrieval to achieve state-of-the-art comparable performance. Hence, we conduct a detailed analysis of different design aspects to construct a baseline approach that achieves state-of-the-art performance on UPAR as well as all sub-datasets for both tasks. We apply current backbone models and optimizers, improve the training batch size, and leverage several regularization methods resulting in an approach robust against overfitting and overconfidence.

Architecture. We follow the typical baseline architecture, which consists of a backbone model with a fully-connected classification head. We make use of only one classification layer that follows the global pooling layer of the backbone. Our experiments showed that average pooling works best, which is why it is applied throughout all investigations.

Backbones. Most of the approaches still rely on ResNet50 [11] as the backbone model since it achieved competitive performance at reasonable computational costs on various tasks for years. However, several new network architectures have been proposed recently. Starting with the ViT [6], transformer-based models gained increasing importance and outperformed CNNs in various research fields. Furthermore, improved CNN models are introduced, e.g., the ConvNeXt [27] architecture which resembles the structure of modern vision transformers but with convolutional blocks. In our work, we compare the ResNet50 [11] model with vision transformers Swin-S [26] and PVTv2 [42], and ConvNeXt-B [27].

Exponential Moving Average. Research on various model architectures such as EfficientNet [36] has shown that using EMAs of a model’s trainable parameters for evaluation may significantly improve the results. An additional EMA model is maintained during training. In each iteration, the original model is updated first. Subsequently, the EMA is computed for each parameter based on the current EMA values and the new parameter values of the original model.

Batch size. Our experimental investigations indicate that the training batch size influences the final evaluation performance. The reason is that if the batch size is too large, rarely

occurring attributes only have little impact on batch updates of the model’s parameters. However, if the batch size is small, the model tends to overfit on some samples, which may result in unstable training behavior and poor generalization performance. Furthermore, regularizations such as EMA may require fewer images per batch to learn meaningful features for fine-grained or rare attributes, especially on small datasets.

Dropout. Dropout is a well-studied regularization method to avoid overfitting. In particular, this improves the generalization ability on data from unseen domains, i.e., the results on the UPAR evaluation protocols. However, it is also beneficial for specialization tasks within the same domain since attribute datasets include huge biases, and models start to overfit after a few epochs.

Label smoothing. Two common problems with deep learning-based approaches for image classification tasks are overconfidence and overfitting. Applying label smoothing tackles both problems simultaneously. If a model is overconfident, output predictions do not reflect the accuracy, i.e., output scores for frequently occurring attributes may be significantly higher than the accuracy.

Typically, label smoothing is applied when the cross-entropy loss is used as an optimization target for multi-class classification. However, we adopt the same idea to the multi-label attribute recognition problem as follows:

$$y_{ls} = (1 - \alpha) * y + \alpha * (1 - y). \quad (1)$$

In Equation 1, y is the original, binary attribute label, α represents the hyperparameter to control the impact of label smoothing, and y_{ls} stands for the resulting smoothed training label.

Optimizer. Many works use adaptive optimizers such as Adam to train their neural networks since they require less fine-tuning of hyperparameters and typically lead to faster training than stochastic gradient descent. However, models trained with Adam have shown to generalize worse. Loschilov *et al.* [28] indicate that the reason is the less effective L2 weight regularization for adaptive optimization methods and proposed the AdamW optimizer, which fixed the problem. Nevertheless, current PAR baselines [16] still rely on the vanilla version of Adam and therefore suffer from poor generalization. In contrast, we apply the AdamW optimizer in our approach and show significant improvement. Other algorithms, e.g., RAdam [23], were investigated but achieved worse results than AdamW.

Data augmentation. Another commonly used method to reduce overfitting is data augmentation. We conducted experiments with various augmentation methods and finally found that Random Erasing (RE) [50] and AugMix (AM) [12] augmentation algorithms improve the results. Besides, random flipping and cropping established them-

selves as standard augmentation methods and are thus always applied in our experiments.

5. Experiments

In this section, we present and discuss experimental results. First, general information about training and evaluation is given, followed by thorough studies regarding generalization on the UPAR dataset. Second, we compare our baseline model to state-of-the-art methods on single datasets.

5.1. Training Setup

In our work, we leverage the work of Jia *et al.* [16] as our baseline. We train all models using a learning rate of $1e-4$, weight decay of $5e-4$, and the Adam optimizer if otherwise stated. Regarding the learning rate schedule, we apply a plateau scheduler that reduces the learning rate by a factor of 0.1 if the validation results do not improve for four epochs. Batches of size 64 are used per default, and gradient clipping is applied after computing the gradients based on the weighted cross-entropy loss function [17]. The backbone networks are initialized with pre-trained weights from the ImageNet dataset [3]. More details can be found in the supplementary material.

We optimize all models in this work, focusing on retrieval performance since it is the relevant measure for most applications. Hence, obtained PAR results are not optimal, especially concerning the mA. However, we are still able to outperform state-of-the-art approaches.

We train our models using PyTorch 1.11 and CUDA 11.3. To speed up the trainings, we leverage adaptive mixed precision scaling of the models’ trainable parameters. The experiments were done on a server with NVIDIA GeForce RTX 3090 GPUs.

5.2. UPAR Generalization

CV Results. Results for training on one sub-dataset and evaluating on the others are provided in Tab. 2. Regarding approaches from literature, one can observe that the strong baseline [16], which also serves as our baseline, achieves the best results. Especially, more complex methods that, e.g., apply GANs [1] achieve poor results for generalizing on out-of-distribution domains. Furthermore, using ResNet50 as the backbone model performs significantly worse than using more recent CNN- or transformer-based models. Best results are obtained by the ConvNeXt-B architecture. This applies to both tasks PAR and retrieval. As a result, we rely on this backbone model for our method. All the proposed regularization techniques improve the scores. Adding Dropout to the last feature layer of the backbone and using AdamW instead of vanilla Adam has the most considerable impact on performance. Interesting findings are that EMA evaluation is only beneficial

Method	Backbone	mA	F1	mAP	R-1
VAC [9]†	ResNet50	64.3 ± 1.8	71.4 ± 6.1	6.2 ± 3.3	7.4 ± 4.1
ALM [38]†	BN-Inception	66.3 ± 2.0	71.0 ± 5.4	5.5 ± 3.0	7.3 ± 4.0
SAL [1]†	ResNet50	—	—	3.7 ± 1.5	4.8 ± 2.0
Baseline [16]†	ResNet50	65.5 ± 2.2	71.2 ± 5.6	6.6 ± 3.4	8.7 ± 4.5
	Swin-S	68.4 ± 2.2	74.5 ± 5.2	8.6 ± 3.5	9.9 ± 4.5
	PVTv2	68.1 ± 2.1	75.0 ± 4.2	8.6 ± 3.4	10.4 ± 4.5
	ConvNeXt-B	70.2 ± 1.2	76.8 ± 4.2	11.1 ± 3.9	12.9 ± 5.1
+ EMA	ConvNeXt-B	69.3 ± 1.0	78.0 ± 3.5	11.8 ± 4.3	13.1 ± 5.1
+ optimal BS	ConvNeXt-B	69.4 ± 1.1	78.3 ± 3.0	12.0 ± 4.3	13.4 ± 5.3
+ dropout	ConvNeXt-B	69.8 ± 1.4	79.1 ± 3.1	12.6 ± 4.2	14.3 ± 4.9
+ label smoothing	ConvNeXt-B	70.0 ± 1.5	79.2 ± 3.0	12.7 ± 4.0	14.5 ± 4.7
+ AdamW	ConvNeXt-B	70.8 ± 1.7	79.5 ± 3.1	13.4 ± 4.3	15.4 ± 5.4
+ AugMix (Ours)	ConvNeXt-B	70.5 ± 1.9	80.1 ± 2.7	13.7 ± 4.2	15.7 ± 5.0
Ours	ResNet50	67.0 ± 2.5	74.2 ± 4.5	8.3 ± 4.1	10.2 ± 5.0

Table 2. **UPAR CV** – Generalization results when using data from only one domain for training. The proposed optimizations of the baseline approach increase the performance regarding all metrics. † Results were produced using the official implementation.

when enough training data is available. Splits 0 and 2, i.e., leveraging the small Market1501 or PETA datasets for training, achieve better results without EMA computation of weights. In contrast, regularization with EMA requires a lower batch size of 32 in this generalization setting (splits 1 & 3). Otherwise, small-scale features for highly localized attributes are hardly learned. Similar differences between splits with more and less training data are observed for the label smoothing hyperparameter α . The two smaller splits need more regularization, i.e., $\alpha = 0.1$, while $\alpha = 0.05$ is sufficient for the other splits.

It is also noteworthy that applying AdamW before adding regularization does not improve the results. The benefits of AdamW can not compensate for the model’s tendency to overfit. Concerning data augmentation, only AM improves the results. The reason is that RE may lead to overfitting on the training data since if relevant image regions are erased, some attributes may not be determinable, which raises the risk of recognizing attributes for this image based on irrelevant background characteristics. In conclusion, the experiments indicate that the improvements are universal, but the parameterization highly depends on the amount of available training data when focusing on generalization to unseen domains. If few training images (<15,000 images in our case) are available, more regularization will be needed, and EMA is not advantageous. In contrast, training with more data (>50,000 images in our case) benefits from EMA, but the batch size should be reduced to keep the influence of images showing rare attributes.

Our final results greatly outperform the state-of-the-art and show that current algorithms with a simple baseline architecture are well-suited for generalization tasks. Finally, we apply the same bag of tricks with identical hyperparameters to the ResNet50 backbone, although the parameters might not be optimal. However, results indicate that there also is a significant improvement, and the state-of-the-art is outperformed.

Method	Backbone	mA	F1	mAP	R-1
VAC [9]†	ResNet50	64.3 ± 1.8	71.4 ± 6.1	6.2 ± 3.3	7.4 ± 4.1
ALM [38]†	BN-Inception	72.0 ± 2.3	78.0 ± 2.7	11.4 ± 4.2	13.9 ± 9.0
SAL [1]†	ResNet50	—	—	7.9 ± 3.4	10.5 ± 7.2
Baseline [16]	ResNet50	71.5 ± 1.9	78.7 ± 2.9	13.4 ± 5.0	16.0 ± 10.5
	Swin-S	72.9 ± 1.9	80.4 ± 2.6	14.6 ± 4.4	17.1 ± 9.3
	PVTv2	72.5 ± 1.7	80.7 ± 2.8	15.0 ± 4.4	17.0 ± 8.7
	ConvNeXt-B	73.2 ± 1.8	81.8 ± 2.7	17.2 ± 3.7	20.1 ± 9.4
+ EMA	ConvNeXt-B	72.4 ± 1.9	83.1 ± 2.1	19.3 ± 5.3	21.3 ± 9.9
+ dropout	ConvNeXt-B	73.4 ± 2.4	83.5 ± 2.2	20.2 ± 5.2	22.1 ± 9.1
+ label smoothing	ConvNeXt-B	74.4 ± 2.2	83.5 ± 2.1	20.2 ± 5.1	22.2 ± 8.9
+ AdamW	ConvNeXt-B	75.3 ± 2.0	84.1 ± 2.0	20.9 ± 5.3	22.9 ± 9.5
+ AugMix (Ours)	ConvNeXt-B	75.4 ± 1.9	84.2 ± 2.0	21.4 ± 5.4	23.4 ± 9.1
Ours	ResNet50	72.6 ± 2.4	81.4 ± 2.3	16.0 ± 5.4	18.5 ± 10.3

Table 3. **UPAR LOO** – Generalization results when using data from multiple domains for training. The proposed optimizations of the baseline approach increase the performance regarding all metrics. † Results were produced using the official implementation.

LOO Results. Results for the LOO generalization evaluation protocol are given in Tab. 3. Observations are similar to those seen in the previous section. Since more training data is available in each split, EMA always led to an improvement, and adjusting the batch size was unnecessary. Again, our approach significantly outperforms methods from the literature regardless of the backbone model used.

Comparing both generalization protocols shows that attribute-based person retrieval primarily benefits from additional and more diverse training data. However, there is still much room for improvement. Considering real-world applications, a R-1 score of 15.7% might not be sufficient. Even with training data from multiple domains, further research is required. While a positive match at the first ranking position in almost every fourth case seems sufficient in some cases, the mAP indicates that many matches occur in late ranking positions. As a result, relevant people might be missed. This finding indicates that current attribute classifiers can be transferred to other domains and used for retrieval, but they only work well for simple cases so far. Besides, PAR models seem to be better suited for generalization tasks than methods aiming at learning a cross-modal feature space using adversarial training.

5.3. Specialization Results

Ablation. We also evaluate the regularization methods for the specialization case, i.e., training and testing data originate from the same distribution, in Tab. 4. We observed that in contrast to the generalization experiments, optimal hyperparameters and improvements by the optimizations greatly depend on the dataset. This finding indicates that the datasets are biased and that models optimized using the same data source for training and testing may result in poor performance in real-world applications. Also, RE leads to improvements in the specialization case, which hurts generalization ability, as was observed in the UPAR experiments.

Method	Backbone	PETA				PA100K				RAPv2				Market-1501				UPAR			
		mA	F1	mAP	R-1	mA	F1	mAP	R-1	mA	F1	mAP	R-1	mA	F1	mAP	R-1	mA	F1	mAP	R-1
Baseline [16] [†]	ResNet50	84.0	86.3	20.7	21.3	81.6	88.1	23.8	31.1	77.4	78.5	17.0	12.1	76.5	83.6	23.8	39.5	83.2	87.3	21.2	23.5
	Swin-S	86.6	87.7	24.1	24.1	83.2	88.5	25.7	32.9	80.0	80.3	21.0	15.0	78.2	84.7	26.3	39.5	81.4	86.9	19.9	21.4
	PVTv2	84.1	86.3	20.5	20.3	82.1	88.7	25.8	33.5	78.0	78.8	17.2	12.0	77.5	83.8	25.5	39.5	83.7	88.6	26.1	28.5
	ConvNeXt-B	86.1	88.1	24.4	24.4	82.2	88.5	26.2	34.5	79.3	80.0	20.5	14.4	80.7	85.7	31.6	47.7	83.9	89.2	26.9	28.7
+ EMA	ConvNeXt-B	85.6	88.4	25.4	24.7	83.7	89.7	29.3	36.3	78.5	80.9	22.2	15.9	77.5	86.2	29.9	45.3	83.9	89.2	27.3	29.3
+ optimal BS	ConvNeXt-B	87.0	88.5	26.7	26.5	83.7	89.7	29.3	36.3	78.5	80.9	22.2	15.9	79.6	86.7	35.7	51.0	83.9	89.2	27.3	29.3
+ dropout	ConvNeXt-B	87.1	88.4	26.5	26.8	84.0	89.7	30.0	37.6	79.9	81.0	22.3	16.3	79.3	86.6	36.5	49.8	84.3	89.4	27.7	29.6
+ label smoothing	ConvNeXt-B	87.4	88.8	27.7	27.7	84.8	90.0	30.0	38.5	79.5	81.1	21.7	15.8	79.0	86.8	37.4	50.2	84.9	89.4	27.9	30.1
+ AdamW	ConvNeXt-B	88.1	89.7	30.4	29.3	84.8	90.2	30.5	39.5	79.8	81.1	20.4	14.9	81.4	87.7	40.3	52.5	85.7	90.2	30.7	32.7
+ Random Erasing	ConvNeXt-B	88.4	89.9	30.2	29.7	84.6	90.4	30.7	38.8	79.7	81.1	22.5	16.4	81.5	87.6	40.6	55.4	85.9	90.2	31.6	33.9

Table 4. **Specialization results** – Analogous to the generalization protocols on UPAR, the proposed optimizations of the baseline method also lead to significant improvements for specializing on a single domain. [†] Results were produced using the official implementation.

Method	Backbone	PETA		PA100K		RAPv2		Market1501	
		mA	F1	mA	F1	mA	F1	mA	F1
MsVAA[30]	ResNet101	84.6	86.5	–	–	–	–	–	–
VAC [9]	ResNet50	83.6	86.2	79.0	86.8	–	–	–	–
ALM [38]	BN-Inception	86.3	86.9	80.7	86.5	–	–	–	–
JLAC [37]	ResNet50 [†]	87.0	87.5	82.3	87.6	–	–	–	–
VFA [2]	ResNet50	86.5	87.3	81.3	87.0	–	–	–	–
MSCC [49]	ResNet50	–	–	82.1	86.8	80.2	79.1	–	–
SB [16]	ResNet50	84.0	86.4	80.2	87.4	78.5	78.7	–	–
SB [16] [†]	ResNet50	84.0	86.3	81.6	88.1	77.4	78.5	76.5	83.6
SSC _{soft} [15]	ResNet50	86.5	87.0	81.9	86.9	–	–	–	–
Ours	ResNet50	87.1	87.7	82.2	88.5	78.8	80.0	79.5	85.4
	ConvNeXt-B	88.4	89.9	84.8	90.2	79.9	81.0	81.5	87.6

Table 5. **State-of-the-art PAR** – Our models achieve state-of-the-art results on every benchmark. Only on PA100K and RAPv2 only few approaches from the literature report higher mA scores. The reason is that our models were optimized for attribute-based retrieval. **Red** and **blue** colors highlight the best and second-best results, respectively. [†] Results were produced using the official implementation.

These findings clearly highlight the necessity of a general benchmark such as UPAR consisting of data from different domains with varying image and attribute distributions. Otherwise, developing generalizable and applicable algorithms is hardly possible.

Comparison with State-of-the-art. Tables 5 and 6 provide comparisons with the current state-of-the-art for PAR and attribute-based person retrieval. Regarding both tasks, our approach sets a new state-of-the-art regardless of the choice of the backbone architecture.

6. Discussion of Potential Societal Impact

The fields of PAR and attribute-based retrieval, which encompass fashion attributes as well as attributes related to visual surveillance, have several potential uses in real-world scenarios. Intended scenarios may include the use of retrograde PAR and retrieval systems by law enforcement authorities to identify suspects based on recalled testimonies. Another application may be group detection and monitoring during large events in public transportation. However, malicious use would imply identifying persons and groups based on their ethnicity, skin color, religion, and or cultural accessories or associated clothing. Mitigation strate-

Method	PETA		PA100K		RAPv2		Market1501	
	mAP	R-1	mAP	R-1	mAP	R-1	mAP	R-1
DeepMAR [18]	–	–	–	–	–	–	8.9	13.2
DCCAE [41]	–	–	15.6	21.2	–	–	9.7	8.1
2WayNet [7]	–	–	10.6	19.5	–	–	7.8	11.3
CMCE [21]	–	–	13.1	25.8	–	–	22.8	35.0
AAIPR [46]	–	–	–	–	–	–	20.7	40.3
AIHM [5]	–	–	–	–	–	–	24.3	43.3
SAL [1] [†]	–	–	15.0	22.7	–	–	29.4	44.4
CWDD [33]	–	–	–	–	21.1	14.5	27.5	42.1
SB [16] [†]	20.7	21.3	23.8	31.1	17.0	12.1	23.8	39.5
ASMR [14]	–	–	20.6	31.9	–	–	31.0	49.6
Ours (ResNet50)	23.0	22.9	25.6	32.5	19.5	13.8	32.3	45.0
Ours (ConvNeXt-B)	30.2	29.7	30.5	39.5	21.7	15.8	40.6	55.4

Table 6. **State-of-the-art attribute-based retrieval** – Our models achieve state-of-the-art results on every benchmark. PAR-based and feature alignment-based methods are significantly outperformed independent of the backbone model. **Red** and **blue** colors highlight the best and second-best results, respectively. [†] Results were produced using the official implementation.

gies should preclude the integration of cultural and religious attributes into public datasets. Furthermore, it is currently unclear how well PAR models distinguish between clothing and skin color in low-resolution surveillance images, *i.e.*, whether the model could confuse a dark-skinned person with short lower-body clothing and a person with long black lower-body clothing (*e.g.*, leggings). We think that our UPAR dataset offers more diversity and lowers the impact of biases in datasets. Further work should study latent biases in public datasets and UPAR.

7. Conclusions

In this work, we proposed a unified dataset named UPAR to allow generalization experiments for 40 attributes across four PAR datasets along with standard evaluation schemes with 148,048 train, 30,830 validation, and 45,859 test images to measure the generalization ability of PAR and attribute-based person retrieval approaches. Besides, we develop a robust baseline that outperforms state-of-the-art approaches in generalization and specialization problems of PAR and retrieval. Based on UPAR, we believe the research community will develop new large-scale generalizable algorithms for attribute recognition and attribute-based person retrieval.

References

- [1] Yu-Tong Cao, Jingya Wang, and Dacheng Tao. Symbiotic adversarial learning for attribute-based person search. In *Proc. European Conference on Computer Vision (ECCV)*, 2020.
- [2] Ming Chen, Guijin Wang, Jing-Hao Xue, Zijian Ding, and Li Sun. Enhance via decoupling: Improving multi-label classifiers with variational feature augmentation. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 1329–1333. IEEE, 2021.
- [3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [4] Yubin Deng, Ping Luo, Chen Change Loy, and Xiaoou Tang. Pedestrian attribute recognition at far distance. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 789–792, 2014.
- [5] Qi Dong, Shaogang Gong, and Xiatian Zhu. Person search by text attribute query as zero-shot learning. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.
- [7] Aviv Eisenschlat and Lior Wolf. Linking image and text with 2-way nets. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [8] Lucas Florin, Andreas Specker, Arne Schumann, and Jürgen Beyerer. Hardness prediction for more reliable attribute-based person re-identification. In *2021 IEEE 4th International Conference on Multimedia Information Processing and Retrieval (MIPR)*, pages 418–424. Institute of Electrical and Electronics Engineers (IEEE), 2021.
- [9] Hao Guo, Kang Zheng, Xiaochuan Fan, Hongkai Yu, and Song Wang. Visual attention consistency under image transforms for multi-label image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 729–739, 2019.
- [10] Yandong Guo, Lei Zhang, Yuxiao Hu, X. He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *ECCV*, 2016.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [12] Dan Hendrycks, Norman Mu, Ekin D Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. *arXiv preprint arXiv:1912.02781*, 2019.
- [13] Ruibing Hou, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Temporal complementary learning for video person re-identification. In *Proceedings of the European Conference on Computer Vision*, 2020.
- [14] Boseung Jeong, Jicheol Park, and Suha Kwak. Asmr: Learning attribute-based person search with adaptive semantic margin regularizer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12016–12025, 2021.
- [15] Jian Jia, Xiaotang Chen, and Kaiqi Huang. Spatial and semantic consistency regularizations for pedestrian attribute recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 962–971, 2021.
- [16] Jian Jia, Houjing Huang, Xiaotang Chen, and Kaiqi Huang. Rethinking of pedestrian attribute recognition: A reliable evaluation under zero-shot pedestrian identity setting. *arXiv preprint arXiv:2107.03576*, 2021.
- [17] Dangwei Li, Xiaotang Chen, and Kaiqi Huang. Multi-attribute learning for pedestrian attribute recognition in surveillance scenarios. In *ACPR*, pages 111–115, 2015.
- [18] Dangwei Li, Xiaotang Chen, and Kaiqi Huang. Multi-attribute learning for pedestrian attribute recognition in surveillance scenarios. In *IAPR Asian Conference on Pattern Recognition (ACPR)*, 2015.
- [19] Dangwei Li, Xiaotang Chen, Zhang Zhang, and Kaiqi Huang. Pose guided deep model for pedestrian attribute recognition in surveillance scenarios. In *2018 IEEE International Conference on Multimedia and Expo*, pages 1–6. IEEE, 2018.
- [20] D. Li, Z. Zhang, X. Chen, and K. Huang. A richly annotated pedestrian dataset for person retrieval in real surveillance scenarios. *IEEE Transactions on Image Processing*, 28(4):1575–1590, 2019.
- [21] Shuang Li, Tong Xiao, Hongsheng Li, Wei Yang, and Xiaogang Wang. Identity-aware textual-visual matching with latent co-attention. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [22] Yutian Lin, Liang Zheng, Zhedong Zheng, Yu Wu, Zhi-lan Hu, Chenggang Yan, and Yi Yang. Improving person re-identification by attribute and identity learning. *Pattern Recognition*, 2019.
- [23] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019.
- [24] Pengze Liu, Xihui Liu, Junjie Yan, and Jing Shao. Localization guided learning for pedestrian attribute recognition. *arXiv preprint arXiv:1808.09102*, 2018.
- [25] Xihui Liu, Haiyu Zhao, Maoqing Tian, Lu Sheng, Jing Shao, Junjie Yan, and Xiaogang Wang. Hydraplus-net: Attentive deep features for pedestrian analysis. In *Proceedings of the IEEE international conference on computer vision*, pages 1–9, 2017.
- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *International Conference on Computer Vision (ICCV)*, 2021.
- [27] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

- [28] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. 2018.
- [29] Ergys Ristani, Francesco Solera, Roger S. Zou, R. Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *ECCV Workshops*, 2016.
- [30] Nikolaos Sarafianos, Xiang Xu, and Ioannis A Kakadiaris. Deep imbalanced attribute classification using visual attention aggregation. In *Proceedings of the European Conference on Computer Vision*, pages 680–697, 2018.
- [31] Walter J Scheirer, Neeraj Kumar, Peter N Belhumeur, and Terrance E Boulton. Multi-attribute spaces: Calibration for attribute fusion and similarity search. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [32] Arne Schumann, Andreas Specker, and Jürgen Beyerer. Attribute-based person retrieval and search in video sequences. In *2018 15th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–6. IEEE, 2018.
- [33] Andreas Specker and Jürgen Beyerer. Improving attribute-based person retrieval by using a calibrated, weighted, and distribution-based distance metric. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 2378–2382. IEEE, 2021.
- [34] Andreas Specker, Arne Schumann, and Jürgen Beyerer. An interactive framework for cross-modal attribute-based person retrieval. In *16th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), 2019, Taipei, Taiwan, 18-21 Sept. 2019*, page 8909832. Institute of Electrical and Electronics Engineers (IEEE), 2019.
- [35] Andreas Specker, Arne Schumann, and Jürgen Beyerer. An evaluation of design choices for pedestrian attribute recognition in video. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 2331–2335. IEEE, 2020.
- [36] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.
- [37] Zichang Tan, Yang Yang, Jun Wan, Guodong Guo, and Stan Z Li. Relation-aware pedestrian attribute recognition with graph convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12055–12062, 2020.
- [38] Chufeng Tang, Lu Sheng, Zhaoxiang Zhang, and Xiaolin Hu. Improving pedestrian attribute recognition with weakly-supervised multi-scale attribute-specific localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4997–5006, 2019.
- [39] Daniel A Vaquero, Rogerio S Feris, Duan Tran, Lisa Brown, Arun Hampapur, and Matthew Turk. Attribute-based people search in surveillance environments. In *Proc. Winter Conference on Applications of Computer Vision (WACV)*, 2009.
- [40] Jingya Wang, Xiatian Zhu, Shaogang Gong, and Wei Li. Attribute recognition by joint recurrent learning of context and correlation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 531–540, 2017.
- [41] Weiran Wang, Raman Arora, Karen Livescu, and Jeff Bilmes. On deep multi-view representation learning. In *Proc. International Conference on Machine Learning (ICML)*, 2015.
- [42] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pvt2: Improved baselines with pyramid vision transformer. *arXiv preprint arXiv:2106.13797*, 2021.
- [43] Xiao Wang, Shaofei Zheng, Rui Yang, Aihua Zheng, Zhe Chen, Jin Tang, and Bin Luo. Pedestrian attribute recognition: A survey. *Pattern Recognition*, 121:108220, 2022.
- [44] Jie Yang, Jiarou Fan, Yiru Wang, Yige Wang, Weihao Gan, Lin Liu, and Wei Wu. Hierarchical feature embedding for attribute recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 13055–13064, 2020.
- [45] Wenjie Yang, Houjing Huang, Zhang Zhang, Xiaotang Chen, Kaiqi Huang, and Shu Zhang. Towards rich feature discovery with class activation maps augmentation for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1389–1398, 2019.
- [46] Zhou Yin, Wei-Shi Zheng, Ancong Wu, Hong-Xing Yu, Hai Wan, Xiaowei Guo, Feiyue Huang, and Jianhuang Lai. Adversarial attribute-image person re-identification. In *Proc. International Joint Conferences on Artificial Intelligence (IJCAI)*, 2018.
- [47] Kai Yu, Biao Leng, Zhang Zhang, Dangwei Li, and Kaiqi Huang. Weakly-supervised learning of mid-level features for pedestrian attribute recognition and localization. *arXiv preprint arXiv:1611.05603*, 2016.
- [48] Liang Zheng, Liye 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*, 2015.
- [49] Jiabao Zhong, Hezhe Qiao, Lin Chen, Mingsheng Shang, and Qun Liu. Improving pedestrian attribute recognition with multi-scale spatial calibration. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2021.
- [50] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 13001–13008, 2020.