Learned vs. Hand-Crafted Features for Pedestrian Gender Recognition

Grigory Antipov, Sid-Ahmed Berrani
Orange Labs – France Telecom
35512 Cesson-Sévigné, France
{grigory.antipov, sidahmed.berrani}@orange.com

Natacha Ruchaud, Jean-Luc Dugelay
Eurecom
06410 Biot, France
{natacha.ruchaud, jean-luc.dugelay}@eurecom.fr

ABSTRACT

This paper addresses the problem of image features selection for pedestrian gender recognition. Hand-crafted features (such as HOG) are compared with learned features which are obtained by training convolutional neural networks. The comparison is performed on the recently created collection of versatile pedestrian datasets which allows us to evaluate the impact of dataset properties on the performance of features. The study shows that hand-crafted and learned features perform equally well on small-sized homogeneous datasets. However, learned features significantly outperform hand-crafted ones in the case of heterogeneous and unfamiliar (unseen) datasets. Our best model which is based on learned features obtains 79% average recognition rate on completely unseen datasets. We also show that a relatively small convolutional neural network is able to produce competitive features even with little training data.

Categories and Subject Descriptors

I.4.7 [**Feature Measurement**]: Feature representation.

Keywords

Pedestrian gender recognition; CNN; HOG; image features.

1. INTRODUCTION

An ability to profile people based on their gender is a very important issue which has obvious applications in video surveillance and multimedia retrieval systems. However, quite often it is not possible to get a clear close-shot of a person's face and the gender should be estimated having only a general silhouette of a body. In this work, we address the problem of the gender recognition from still images of pedestrians taken in adverse conditions.

The choice of features to describe an object is crucial in computer vision. Existing image features can be roughly divided into 2 categories: the hand-crafted and the learned ones. By hand-crafted features we understand those which are extracted from separate images according to a certain *manually predefined algorithm* based on the expert knowledge. LBP [16], SIFT [14] and HOG [4] features are commonly known examples of hand-crafted features. Contrary to hand-crafted image features, the learned ones are derived from an image dataset by a training procedure in order to fulfill a certain task (e.g. gender recognition). Convolutional Neural Networks (CNN) [13] are examples of deep neural networks which can be used to extract learned features.

We consider 4 types of features in this study: 2 examples per each considered category. We choose person re-identification [12] features and HOG as examples of hand-crafted features. In order to obtain learned features, we train 2 different CNN architectures.

The rest of the paper is organized as follows: the bibliography overview and motivations for the choice of features to compare are presented in Section 2; the chosen features are introduced in details in Section 3; the collection of datasets which we have used for experiments is described in Section 4; the performed experiments and obtained results are presented in Section 5 and the main conclusions are highlighted in Section 6.

2. RELATED WORK

To the best of our knowledge, PETA datasets collection [5] is the largest open-access collection of pedestrian images with gender annotation. Authors of PETA propose a universal way to recognize a number of attributes in pedestrian images. Gender is one of the attributes they have considered. In their algorithm, authors of PETA employ re-identification features which were originally described in [12] for the purpose of automatic re-identification of humans in CCTV videos. We use PETA collection of datasets in our experiments. Following the authors of PETA we also choose re-identification features for comparison in our work.[1]

In numerous works, authors have found Histogram of Oriented Gradients (HOG) features to be the most appropriate for pedestrian gender recognition. In particular, in [1] authors claim obtaining 76% recognition rate on MIT dataset, while in [3] authors claim obtaining 80% recognition rate on VIPeR dataset.[2] Hence, we include HOG features in our comparison as the 2nd example of hand-crafted features.

Due to their exceptional success in recent years, deep CNN have become the first-choice solution for supervised learning

[1]We use the re-identification features implementation which was kindly provided to us by authors of PETA.
[2]Both MIT and VIPeR datasets are included in PETA.

on image datasets [19, 6, 11]. In [15] authors employ a CNN for gender recognition on a tiny MIT dataset having only 900 images. They claim obtaining 80% recognition rate. In this work, we design a compact CNN architecture and corresponding learned features contribute to our comparison.

There are more and more examples ([2, 17]) where a CNN trained on a huge general-purpose dataset is successfully fine-tuned for a very specific classification task with little extra data. Some authors even claim that today fine-tuning of a pre-trained CNN is a must-try method in all image classification tasks [17]. In our work we fine-tune the model which was trained by Krizhevsky et al. [11] on the ILSVRC dataset [18] containing about 1.3 million images.

Contrary to previous works, we do not focus on a single dataset, we rather compare considered features on a versatile collection of pedestrian datasets. Moreover, besides evaluating performance of features on separate datasets, we also compare how well they generalize by applying them on unfamiliar (i.e. completely unseen) datasets.

3. FEATURE REPRESENTATIONS

Below we present features that are compared in this study.

3.1 Person re-identification features

Person re-identification features of an image [12] is a 2784-dimensional vector which contains low-level colour and texture information. The complete vector is composed of six 464-dimensional vectors each of which is extracted from 6 equal sized horizontal strips from the image. Each strip uses 8 colour channels (RGB, HSV and YCbCr) and 21 texture filters (Gabor, Schmid) derived from the luminance channel. We use a bin size of 16 to describe each channel.

3.2 HOG features

In order to extract HOG features, we use 8-by-8 square cells which are organized in 2-by-2 blocks. The number of histogram bins is set to 9. When an image of 50 pixels width and of 150 height is given as input (which is the case in our experiments), the resulting HOG features vector with the described parameters has the dimension 2448.

3.3 Features learned by Mini-CNN

In order to avoid ambiguity, in this section and below, we refer to the CNN which is designed by ourselves as *Mini-CNN*, whereas the term *AlexNet-CNN* is employed to refer to the CNN which is designed by Alex Krizhevsky.

The architecture of *Mini-CNN* is presented in Figure 1. *Mini-CNN* has 2 convolutional layers ($C1$ and $C2$) with 5x5 kernels, each of which is followed by max-pooling with a stride of 2 pixels ($P1$ and $P2$) and Rectified Linear Unit (ReLU) activations.[3] *Mini-CNN* is ended by 2 fully connected layers ($F1$ and $F2$). The $F2$ layer has 2 neurons (corresponding to the number of classes). The loss is computed by a softmax loss function. As it is depicted in Figure 1, *Mini-CNN* takes 3 $(50, 150)$-dimensional feature maps (red, green and blue channels of an image) as an input.[4] The 1st and 2nd convolutional layers $C1$ and $C2$ have $n1$ and $n2$ feature maps respectively. Values of $n1$ and $n2$ depend on the experiment: in the 1st experiment in Section 5.1 we use $n1 = n2 = 10$, while in the 2nd experiment in Section 5.2 we use a slightly bigger architecture with $n1 = n2 = 20$ (because in the 2nd experiment we have more training images). The number of neurons $n3$ in the last but one fully connected layer $F1$ also varies depending on the experiment: we have $n3 = 25$ and $n3 = 100$ for the 1st and for the 2nd experiments respectively. In order to augment the training data, we use each training image alongside with its mirrored copy. Parameters of *Mini-CNN* have been chosen by trying several architectures and by choosing the most possibly compact one so that performance is not sacrificed.

[3]We have also tried using sigmoid activations but ReLU activations proved to be much faster.

[4]Before being treated by *Mini-CNN*, all input images are rescaled to $(50, 150)$ size.

Figure 2: Examples of images from PETA.

After training *Mini-CNN*, we use the obtained weights to calculate the values of neurons in the $F1$ layer for all testing images. These neuron values serve as *Mini-CNN* features for the testing images. Thus, *Mini-CNN* represents an image by either 25- or 100-dimensional vector (depending on the experiment), which is more than 100 times smaller than sizes of corresponding person re-identification or HOG vectors.

3.4 Features learned by AlexNet-CNN

The architecture of *AlexNet-CNN* is described in details in [11]. *AlexNet-CNN* is already a trained model. In our work, we only fine-tune it to recognize genders of pedestrians. After fine-tuning of *AlexNet-CNN* we use the weights from the last but one fully connected layer (which contains 4096 neurons) as features for input images. Thus, in the case of *AlexNet-CNN*, a features vector dimension is 4096.[5]

4. DATASETS

Originally, the PETA collection consists of 10 datasets of different sizes with a total amount of 19,000 images. Appearances of images hugely vary between different datasets of PETA in terms of image resolutions (from 17x39 to 169x365 pixels), camera angles (pictures are taken either by ground-based cameras or by surveillance cameras which are set at a certain height) and environments (indoors or outdoors). Examples of PETA images are presented in Figure 2.

Authors of PETA perform series of experiments on prediction of gender on the whole collection of 19,000 images [5] using re-identification features. They *randomly* split the total collection of images into 9,500 for training, 1,900 for validation and 7,600 for testing. They report obtaining be-

[5]Training and fine-tuning of all CNNs in this work is performed using Caffe deep learning framework [9].

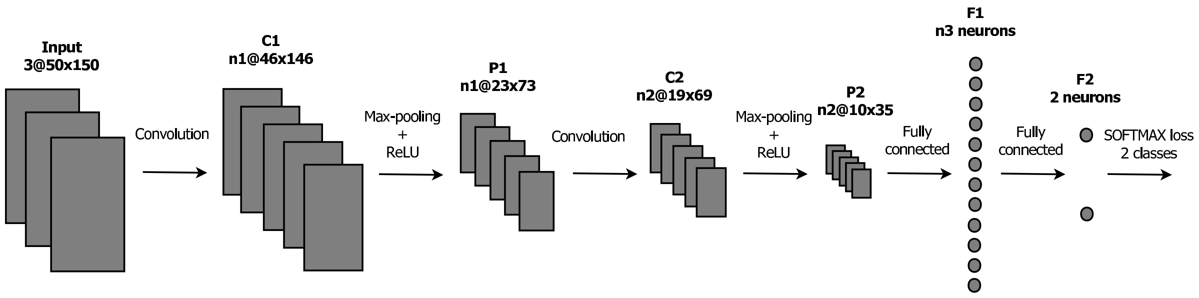

Figure 1: Architecture of *Mini-CNN*.

Dataset	Train size (♂+ ♀)	Test size (♂+ ♀)
CUHK	$3432 = (2420 + 1012)$	$377 = (189 + 188)$
PRID	$942 = (449 + 493)$	$101 = (50 + 51)$
GRID	$928 = (531 + 397)$	$100 = (50 + 50)$
MIT	$792 = (532 + 260)$	$84 = (42 + 42)$
VIPeR	$1138 = (556 + 582)$	$120 = (60 + 60)$
3DPeS	0	$100 = (50 + 50)$
CAVIAR	0	$68 = (34 + 34)$
i-LIDS	0	$100 = (50 + 50)$
SARC3D	0	$41 = (21 + 20)$
TownCentre	0	$42 = (21 + 21)$

Table 1: Training and testing parts per dataset.

tween 79.7% and 81.4% of male gender prediction rate (the results vary depending on the used classifier). We have successfully reproduced this result using several random splits of 19,000 images in the same proportions as it has been done by authors of PETA. However, the number of unique persons in the PETA collection is much smaller than 19,000 because PETA contains many images of the same persons which are taken few seconds away from each other by surveillance cameras. Images like that are almost identical and they can considerably bias the resulting prediction rates. After removal of all quasi-identical images from PETA the resulting prediction rates drop down to 63-65%.

This drastic drop in performance proves the importance of the manual filtering of PETA collection. Apart from filtering out images of the same people, we have also removed images with a very low resolution (height is less than 120 pixels or width is less than 40 pixels) and images where a person of interest is not clearly distinguishable (i.e. images of babies in strollers or images of several persons). Finally, we have been left with 8,365 images (see Table 1) which is less than half as many as the initial size of PETA.

5. EXPERIMENTS

In order to objectively compare the usefulness of considered features for the gender recognition problem and not to take into account a possible impact of a classifier on resulting prediction rates, we always use the same classifier for all considered features: the SVM classifier with a linear kernel.[6] Classification results are compared using Mean Average Precision (MAP) and Area Under ROC Curve (AUC) [7].[7]

[6]In particular, we use the publicly available SVM-light implementation of SVM [10].
[7]AUC is an important measure for us because of its invariance to the chosen decision threshold by an SVM-classifier.

Features	MAP		AUC	
	μ	σ	μ	σ
AlexNet-CNN	0.82	0.05	0.90	0.05
Mini-CNN	0.79	0.04	0.86	0.03
HOG	0.80	0.04	0.88	0.04
Re-identification	0.59	0.07	0.63	0.09

Table 2: Exp. 1: Results on 5 separate datasets (CUHK, PRID, GRID, MIT and VIPeR).

Features	Familiar		Unfamiliar	
	MAP	AUC	MAP	AUC
AlexNet-CNN	0.85	0.91	0.79	0.85
Mini-CNN	0.80	0.88	0.75	0.80
HOG	0.72	0.84	0.56	0.64
Re-identification	0.58	0.60	0.61	0.69

Table 3: Exp. 2: Results on "familiar" and "unfamiliar" datasets.

5.1 Exp. 1: evaluation on separate datasets

In the 1st experiment, we compare the performance of considered features on 5 separate datasets: CUHK, PRID, GRID, MIT and VIPeR. Sizes of training (which is used to train an SVM classifier) and testing parts for each of these datasets are given in Table 1. In the cases of *Mini-CNN* and *AlexNet-CNN* the features (i.e. CNN) are learned on the training images as well.

Results of the 1st experiment are summarized in Table 2. Mean values and standard deviations for 2 considered metrics are calculated over 5 datasets. Based on this experiment, we can conclude that re-identification features are hardly applicable for our problem. Successful results obtained by re-identification features in [5] might be explained by a significant number of quasi-identical images between training and testing datasets (as it is explained in Section 4). The other 3 features show very close performances and none of them can be favoured based only on the 1st experiment.

5.2 Exp. 2: evaluation of generalization

In the 2nd experiment, we evaluate the capability of compared features to generalize on heterogeneous and even completely unseen datasets.

Firstly, we train SVM classifiers with different features on the dataset which is composed of training parts of CUHK, PRID, GRID, MIT and VIPeR taken together. Then we test classifiers on the testing parts of the same datasets which are also taken together. Thus, we compare features on a single big dataset which is composed of non-homogeneous images (different camera angles, environments, etc.)

Results of this experiment are presented in Table 3 (in its "familiar" part). On one hand, we see that for learned features (i.e. *AlexNet-CNN* and *Mini-CNN*) both MAP and AUC metrics practically coincide with corresponding results in Table 2 (with respect to standard deviations). On the other hand, there is a significant drop of performance for HOG features. These results perfectly make sense because learned features have a possibility to adapt to the heterogeneity of the input data during the training phase, whereas HOG features are not trained and, therefore, less flexible. We do not consider re-identification features in the 2nd experiment due to their poor performance in the 1st one.

In the 2nd part of the 2nd experiment, we use the same SVM classifiers which have been trained on the collection of training images from CUHK, PRID, GRID, MIT and VIPeR but this time we apply them on the collection of images from completely unseen datasets: 3DPeS, CAVIAR, i-LIDS, SARC3D and TownCentre (see Table 1). In other words, classifiers are compared on "unfamiliar" datasets and, thus, we evaluate if they generalize well.

Results of this experiment are presented in Table 3 (in its "unfamiliar" part). *AlexNet-CNN* performs almost equally well as in Table 2 (with respect to standard deviations). *Mini-CNN* experiences a subtle drop in performance with respect to the 1st experiment. However, it still shows quite satisfactory recognition rate of 75% despite having only 100 features. The SVM-classifier learned on HOG features performs poorly on "unfamiliar" datasets showing results which are not so far away from random guessing.

6. CONCLUSIONS

In this work, we have compared 2 hand-crafted features (HOG and re-identification) and 2 learned features (obtained by *Mini-CNN* and by *AlexNet-CNN*) in the frame of pedestrian gender recognition problem. Our findings are:

1. On small-sized homogeneous datasets, HOG and learned features perform equally well. It complies with previous works ([1, 3, 15]) on pedestrian gender recognition.

2. Learned features *significantly* outperform HOG features in the case of heterogeneous data.

3. Contrary to hand-crafted features, *Mini-CNN* and *AlexNet-CNN* features generalize well to completely unseen datasets: MAPs of 75% and 79% respectively.

4. Even a relatively small CNN (like *Mini-CNN*) trained on little data is able to produce compact features which generalize almost as good as features produced by much bigger pre-trained networks (like *AlexNet-CNN*).

In our future work, different learned features will be tried in order to recognize other important attributes of soft biometrics (like age, clothing details etc.). Moreover, we would like to verify if our findings about learned and hand-crafted features hold in other domains of computer vision. Being encouraged by the results of this paper and by the recent work presented in [8], we are also planning to further investigate the possibility of approaching the expressiveness of complex models using architectures of moderate complexity.

7. REFERENCES

[1] L. Cao, M. Dikmen, Y. Fu, and T. S. Huang. Gender recognition from body. In *ACM MM*, Canada, 2008.

[2] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. *CoRR*, abs/1405.3531, 2014.

[3] M. Collins, J. Zhang, P. Miller, and H. Wang. Full body image feature representations for gender profiling. In *ICCV*, Japan, 2009.

[4] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, USA, 2005.

[5] Y. Deng, P. Luo, C. C. Loy, and X. Tang. Pedestrian attribute recognition at far distance. In *ACM MM*, USA, 2014.

[6] I. J. Goodfellow, Y. Bulatov, J. Ibarz, S. Arnoud, and V. Shet. Multi-digit number recognition from street view imagery using deep convolutional neural networks. *CoRR*, abs/1312.6082, 2013.

[7] J. A. Hanley and B. J. McNeil. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 1982.

[8] G. E. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning Workshop*, Canada, 2014.

[9] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *CoRR*, abs/1408.5093, 2014.

[10] T. Joachims. Making large scale SVM learning practical. 1999.

[11] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, USA, 2012.

[12] R. Layne, T. M. Hospedales, S. Gong, et al. Person re-identification by attributes. In *BMVC*, UK, 2012.

[13] Y. LeCun and Y. Bengio. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 1995.

[14] D. G. Lowe. Object recognition from local scale-invariant features. In *ICCV*, Canada, 1999.

[15] C.-B. Ng, Y.-H. Tay, and B.-M. Goi. A convolutional neural network for pedestrian gender recognition. In *ISNN*. Springer, 2013.

[16] T. Ojala, M. Pietikäinen, and D. Harwood. A comparative study of texture measures with classification based on featured distributions. *Pattern recognition*, 1996.

[17] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. *CoRR*, abs/1403.6382, 2014.

[18] O. Russakovsky, J. Deng, H. Su, J. Krause, and S. S. et al. ImageNet Large Scale Visual Recognition Challenge. *CoRR*, abs/1409.0575, 2014.

[19] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *CVPR*, USA, 2014.