

HAND-BASED PERSON IDENTIFICATION USING GLOBAL AND PART-AWARE DEEP FEATURE REPRESENTATION LEARNING

Nathanael L. Baisa, Bryan Williams, Hossein Rahmani, Plamen Angelov, Sue Black

ABSTRACT

In cases of serious crime, including sexual abuse, often the only available information with demonstrated potential for identification is images of the hands. Since this evidence is captured in uncontrolled situations, it is difficult to analyse. As global approaches to feature comparison are limited in this case, it is important to extend to consider local information. In this work, we propose hand-based person identification by learning both global and local deep feature representations. Our proposed method, Global and Part-Aware Network (GPA-Net), creates global and local branches on the conv-layer for learning robust discriminative global and part-level features. For learning the local (part-level) features, we perform uniform partitioning on the conv-layer in both horizontal and vertical directions. We retrieve the parts by conducting a soft partition without explicitly partitioning the images or requiring external cues such as pose estimation. We make extensive evaluations on two large multi-ethnic and publicly available hand datasets, demonstrating that our proposed method significantly outperforms competing approaches.

Index Terms— Person identification, Hand recognition, Deep representation learning, Global features, Part-level features

1. INTRODUCTION

There are many human parts or traits that can help to uniquely identify an individual such as face [1, 2], body [3, 4], hand [5, 6], fingerprint [7], Iris [8], gait [9] and voice [10]. A survey of person identification based on different human traits is given in [11, 12, 3]. These are primary biometric traits also referred to as biometric modalities. Personal attributes or ancillary information such as gender, age, (hair, skin, eye) color, ethnicity, height, weight, etc. which are also known as soft biometrics can be cheaply obtained from the primary biometric traits [11]. The identification of an individual can be improved by fusing different primary biometric traits such

as face and body or by fusing biometric modalities with soft biometrics such as face and gender [13], however, enough discriminative information may not be obtained by using only the soft biometric traits to uniquely identify a subject.

Often, the only available information relating to the identity of the perpetrator in cases of serious crime is images of the hand. Hand images, one of the primary biometric traits, provide distinctive features for person identification. They also have less variability when compared to other biometric modalities such as faces [11]. Facial expression has a considerable effect on the variability of the comparable features, whereas the hand has more limited flexibility, suggesting its potential as a more robust biometric trait for identification. Complex security systems rely on detailed images of the hand vein patterns with infrared imaging [14] but this is very rarely available in real cases. Evidence images of the hands are commonly captured with digital cameras [5] and there is a strong need to investigate the potential for identification from digital images of the hand.

A method in [5] trained a two-stream convolutional neural network (CNN) and used it as a feature extractor to obtain CNN-features which have been fed into a set of support vector machine (SVM) classifiers for a hand-based person identification, however, this method is not an end-to-end. In addition, this method treats images of a subject’s right and left hands as the same which is not usual, due to the disparity in comparable features including vein patterns [15]. The work in [16] used rather similar approach with additional data type for fusion, near-infrared (NIR) images. Researchers have also demonstrated that leveraging human part cues from external pose estimation to alleviate pose variations and learning feature representations can improve person re-identification based on images of the body [17]. However, the errors in the pose estimation can propagate to the identification which limits its performance. To alleviate this problem, uniform partitioning for learning part-level features has been proposed in [18], however, this method considered only the part-level features, partitioned in only horizontal stripes, and has been designed for body-based person re-identification.

In this work, we propose a hand-based person identification method, Global and Part-Aware Network (GPA-Net), by learning both global and local deep feature representations from hand images captured by digital cameras. The proposed method follows an end-to-end training. Our contributions can

Nathanael L. Baisa is with the School of Computer Science and Informatics, De Montfort University, Leicester LE1 9BH, UK. Email: nathanael.baisa@dmu.ac.uk.

Bryan Williams, Hossein Rahmani, Plamen Angelov and Sue Black are with the School of Computing and Communications, Lancaster University, Lancaster LA1 4WA, UK. Emails: {b.williams6, h.rahmani, p.angelov, sue.black}@lancaster.ac.uk.

be summarized as follows.

1. We extend the capability of uniform partitioning for hand-based person identification by including both horizontal and vertical partitioning.
2. We incorporate global representation learning in addition to the local (part-level) features proposing a combined approach that is flexible in terms of the backbone architecture for hand-based person recognition.
3. We make extensive evaluations on two large multi-ethnic and public hand datasets (11k hands [5] and HD [19] datasets).

The rest of the paper is organized as follows. The proposed method including the overall architecture of GPA-Net is described in Section 2 followed the experimental results in Section 3. The main conclusion along with suggestion for future work is summarized in Section 4.

2. PROPOSED METHOD

Our proposed method (GPA-Net) is summarized in Fig. 1. We use ResNet50 [20] pretrained on ImageNet as a backbone network due to its concise architecture with competitive performance though any network designed for image classification, for instance Google Inception [21], can be adapted by removing the hidden fully-connected layers. In our case, we make some modifications to the ResNet50 architecture to produce the GPA-Net. The structure of the original ResNet50 before the global average pooling (GAP) layer remains the same. When an input image passes through the backbone network, it becomes a 3D tensor \mathbf{T} of activations just before it enters the GAP layer. At this stage, we create two branches: Global and Local branches explained as follows.

Global branch: For this branch, the GAP layer is used to summarize the 3D tensor \mathbf{T} of activations to form a C-dimensional (2048-dimensional in case of a ResNet50 backbone network) column feature vector \mathbf{f}_g . C is the number of channels of the tensor (feature map) \mathbf{T} .

Local branch: In this case, we change from the GAP layer to conventional average pooling layer to create uniform partitions on the 3D tensor \mathbf{T} of activations along both horizontal and vertical directions to learn part-level features \mathbf{f}_l where $l \in [1, p]$; p is the total number of partitions.

For both branches, new fully-connected layers (FCs) are added to reduce the C-dimensional column feature vectors \mathbf{f} to 512-dimensional feature vectors \mathbf{r} . We also include a batch normalization and dropout with probability of 0.5 to reduce possible over-fitting. Each dimensionally reduced feature vector \mathbf{r} is given as input to a classification layer which is implemented using a FC layer followed by a softmax function. The classification layer predicts the identity (ID) of each input. Given the input image size ($[H, W]$), the spatial size

of the tensor \mathbf{T} ($[M, N]$) is determined by the spatial down-sampling ratio of the backbone network. To increase the size of the tensor \mathbf{T} which in turn increases the performance as observed in [18, 22], we remove the last spatial down-sampling operation by changing the last stride from 2 to 1 in the backbone network.

We optimize the GPA-Net during training by minimizing the sum of the cross-entropy losses over p pieces of ID predictions of the Local branch and ID prediction of the Global branch i.e. each classifier, either of local or global branch, predicts the identity of the input image. The total cross-entropy loss is given as

$$\mathcal{L}_{total} = \underbrace{\sum_{i=1}^K -q_{i,g} \log(p_{i,g})}_{Global\ loss} + \lambda \underbrace{\sum_{l=1}^p \sum_{i=1}^K -q_{i,l} \log(p_{i,l})}_{Local\ loss} \quad (1)$$

where K is the total number of classes (identities). $p_{i,\gamma}$ is the predicted probability of class i ; $\gamma = l$ in case of learning part-level features and $\gamma = g$ in case of learning global features. λ is the weight balancing the losses of the two branches. $\lambda = 1$ is our standard choice; $\lambda = \frac{1}{p}$ decreases the performance, for instance. We also use label smoothing [23] which prevents over-fitting and over-confidence with smoothing value (ϵ) of 0.1. The ground-truth distribution over labels $q_{i,\gamma}$ by including label smoothing can be given as

$$q_{i,\gamma} = \begin{cases} 1 - \frac{K-1}{K}\epsilon, & \text{if } i = y \\ \frac{1}{K}\epsilon, & \text{otherwise} \end{cases} \quad (2)$$

where y is ground-truth label. γ is equal to g for computing global loss or l otherwise.

During testing, we concatenate all C-dimensional (2048-D in our case) feature vectors of p pieces of Local branch and global branch i.e. $\mathcal{F} = [\mathbf{f}_g, \mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_p]$, and then compare features of each query image with gallery features using cosine distance. Concatenating the global and part-level features in this way has a stable performance on all datasets, for instance, when compared to taking the mean of the part-level features and then concatenating to the global features.

3. EXPERIMENTAL RESULTS

3.1. Settings

Datasets: We use two datasets for evaluation, 11k hands¹ dataset [5] and Hong Kong Polytechnic University Hand Dorsal (HD)² dataset [19]. The 11k hands dataset has 11k images with 190 subjects (identities). It also has metadata such as the subject ID, gender, age, skin colour, and a set of information of the captured hand, i.e., right- or left-hand, hand side

¹<https://sites.google.com/view/11khands>

²<http://www4.comp.polyu.edu.hk/~csajaykr/knuck1eV2.htm>

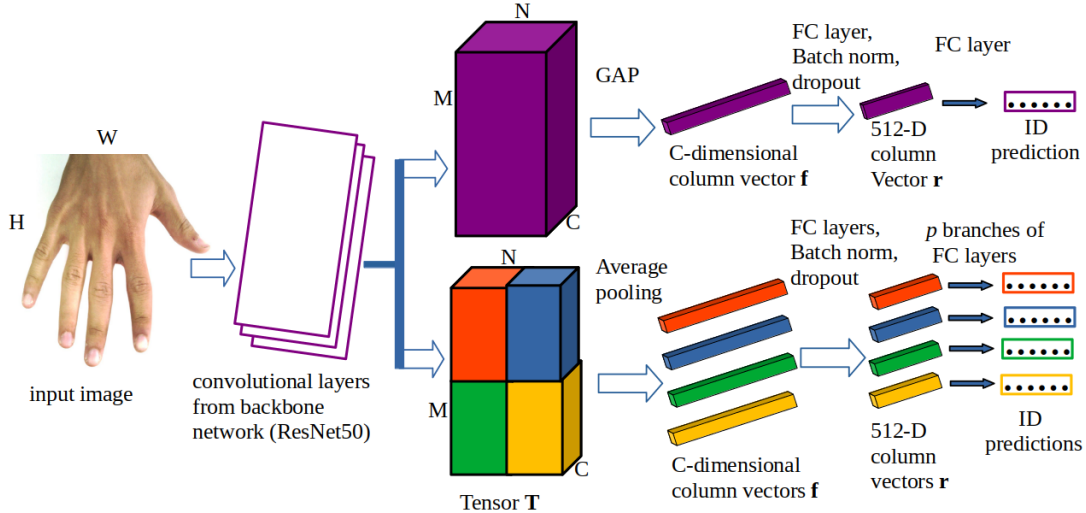


Fig. 1. Structure of GPA-Net. Given an input image, two separate 3D tensors (one for global branch and the other for local branch) are obtained by passing it through the stacked convolutional layers from the backbone network. Each classifier of the local or the global branch predicts the identity of the input image during training.

(dorsal or palmar), and logical indicators referring to whether the hand image contains accessories, nail polish, or irregularities. We exclude any hand image containing accessories from the training data to avoid any potential bias as in [5]. Unlike the work in [5] which treats both right and left dorsal or palmar hand of a subject as the same, which is not usually the case and is even not possible to compare to, we treat a subject to be identified by either right or left dorsal, or right or left palmar for robust identification, rather than mixing them since right and left hands show disparity in their vein patterns [15]. Accordingly, we separate the dataset into right dorsal, left dorsal, right palmar and left palmar sub-datasets to train a person identification model. After excluding accessories and separating the dataset, right dorsal has 143 identities, left dorsal has 146, right palmar has 143 and left palmar has 151 identities. The first half and the second half based on the order of identity of each sub-dataset are used for training and testing, respectively. For instance, for right dorsal, the first 72 identities are used for training and the last 71 identities are used for testing. Similarly, the first 73, 72 and 76 identities are used for training phase for left dorsal, right palmar and left palmar, respectively. The remaining identities of each sub-dataset (73 for left dorsal, 71 for right palmar, 75 for left palmar) are used for testing. From each identity of the test set of each sub-dataset, we randomly choose one image and put in a common gallery for all the sub-datasets. The remaining images of each identity of the test set of each sub-dataset are used as a query set for that sub-dataset. Accordingly, the gallery has 290 images and the query has 971 images for right dorsal, 988 images for left dorsal, 917 images for right palmar and 948 images for left palmar. The

HD dataset has 502 subjects (identities). The first half and the second half of the dataset based on identity are used for training and testing, respectively i.e. identities of 1 - 251 for training and 252 - 502 for testing. From each identity of the test set, one image is randomly chosen and put in a gallery and the rest are used as a query. There are also additional images of 213 subjects in the HD dataset which lack clarity or do not have second minor knuckle patterns. These additional images of 213 subjects are added to the gallery. Accordingly, the gallery for the HD dataset has 1593 images and the query has 1992 images. A randomly chosen image of each identity of the training set of both datasets is used as a validation for monitoring the training process. This procedure is repeated for 10 times and the average performance is reported. Typical exemplar image (query) of each (sub-)dataset is shown in Fig. 3 with ranked results retrieved from a gallery.

Implementation details: The GPA-Net is implemented using PyTorch deep learning framework and trained on NVIDIA GeForce RTX 2080 Ti GPU. We train the model using Cross-Entropy loss and mini-batch Stochastic Gradient Descent (SGD) optimizer with momentum; the Nesterov momentum is also enabled. The mini-batch size, the momentum and the weight decay factor for L2 regularization are set to 20, 0.9 and 5×10^{-4} , respectively. The model is trained for 60 epochs with an initial learning rate of 0.02. We divide this initial learning rate by 10 for the existing layers of the backbone network i.e. bigger initial learning rate is given to the newly added layers (FC layers and batch normalization) with appropriate weight and bias initializations. We also use the learning rate scheduler with a decay factor of 0.1 after every 30 epochs. The input images are resized to 384×384

and then augmented by random horizontal flip, normalization and color jittering. A random order of images are used by reshuffling the data set. During testing, only normalization is used.

Evaluation metrics: We use Cumulative Matching Characteristics (CMC) [24] (Rank-1 matching accuracy) and mean Average Precision (mAP) [24] to evaluate the proposed identification method.

3.2. Performance Evaluation

In this section, we evaluate our model on 11k [5] and HD [19] datasets, and we report the results using rank-1 accuracy and mAP. We use input image size of 384×384 which means the spatial size of \mathbf{T} is 24×24 . After investigating both horizontal and vertical partitions, its ablation analysis is given Table 3. We find out that partitioning the tensor T into 3 horizontally gives the best performance. In addition to same-domain performance, we also report cross-domain performance in order to prevent being misled by over-fitting. Since there is not much work on person identification based on the full hand, we compare with the baselines ResNet50 [20] and VGG with batch normalization (VGG-bn) [25] pretrained on ImageNet by changing the number of classes from 1000 (ImageNet) to the number classes (identities) in our datasets for transfer learning. The training hyperparameter values such as learning rate, batch size, etc. are kept the same for all methods. We use the 2048-dimensional features after the GAP layer of the ResNet50 and the 4096-dimensional output of fc7 of the VGG19-bn as visual representations to compare to our proposed method.

Same-domain performance: In this setting, the model is trained on a specific dataset and then tested on the test set of that dataset. As shown in Table 1, our method (GPA-Net) outperforms the other methods by a large margin in both rank-1 accuracy and mAP on all datasets. For instance, it outperforms the baseline ResNe50 by 6.06% rank-1 accuracy and 5.52% mAP and the baseline VGG-bn by 12.10% rank-1 accuracy and 11.27% mAP on right dorsal (D-r) sub-dataset of 11k dataset. Matching rate versus rank comparison of our method with other methods using CMC plot on right dorsal of 11k dataset is shown in Fig. 2.

Cross-domain performance: In this case, the model is trained on the training set of one dataset and then evaluated on the test set of the other dataset. We conduct a cross-domain performance on the right dorsal (D-r) of 11k and HD datasets as shown in Table 2. D-r (11k) \rightarrow HD shows the model is trained on the training set of D-r (11k) dataset and then evaluated on the test set of HD dataset. Our proposed method outperforms the other methods by a large margin on both datasets which shows our method has more generalizability. For instance, our method outperforms the baseline ResNet50 by 24.74% rank-1 accuracy and 22.57% mAP and the baseline VGG-bn by 40.42% rank-1 accuracy and 37.82%

mAP on the D-r (11k) \rightarrow HD.

3.3. Ablation Study

As described in Section 2, our method considers both horizontal and vertical partitions on the 3D tensor \mathbf{T} of activations in addition to the global 3D tensor. We make an ablation analysis of our method on both partitions and components as follows.

Ablation analysis on partitions: In this case, we investigate how partitioning in horizontal and vertical directions affects the performance of our method. As shown in Table 3, we increase the horizontal partitions (H_p) and the vertical partitions (V_p) from 1 to 4 and from 1 to 3, respectively. The values at $H_p = 1$ and $V_p = 1$ (rank-1 accuracy of 91.19% and mAP of 92.14%) show the global component of our method (only 1 classifier is trained). Similarly, the values at $H_p = 2$ and $V_p = 2$ show the GPA-Net performance with 1 global classifier and 4 local classifiers. Accordingly, horizontal partitioning with $H_p = 3$ (1 global and 3 local classifiers) gives the best performance with 95.83% rank-1 accuracy and 96.31% mAP.

Ablation analysis on components: Our proposed method (GPA-Net) has global and local components which learn global and part-level features, respectively, for robust discriminability. The learned part-level features outperforms the learned global features. Our method which integrates these components outperforms the components when treated separately, in both evaluation metrics, as shown in Table 4. When evaluated on the right palmar (P-r) of 11k hands dataset, for instance, using only the global features gives 91.19% rank-1 accuracy and 92.14% mAP, using only the local features gives 94.33% rank-1 accuracy and 94.91% mAP, and using the proposed GPA-Net gives 95.83% rank-1 accuracy and 96.31% mAP.

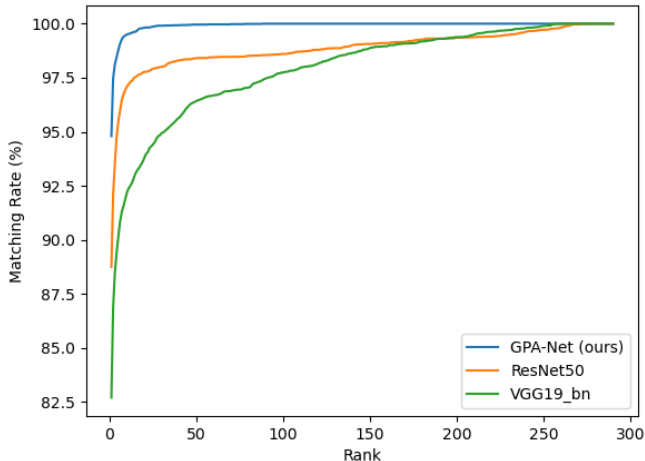


Fig. 2. Matching rate vs rank comparison of our method with other methods using CMC plot on right dorsal of 11k dataset.

Method	D-r of 11k		D-l of 11k		P-r of 11k		P-l of 11k		HD	
	rank-1	mAP	rank-1	mAP	rank-1	mAP	rank-1	mAP	rank-1	mAP
VGG19-bn	82.70	84.45	88.64	89.82	85.09	86.88	89.98	91.31	86.90	88.04
ResNet50	88.74	90.20	91.08	92.12	91.09	92.18	93.30	94.02	89.67	90.63
GPA-Net (ours)	94.80	95.72	94.87	95.93	95.83	96.31	95.72	96.20	94.64	95.08

Table 1. Same-domain performance comparison of our method with other methods on right dorsal (D-r) of 11k, left dorsal (D-l) of 11k, right palmar (P-r) of 11k, left palmar (P-l) of 11k and HD datasets. Rank-1 accuracy (%) and mAP (%) are shown.

Method	D-r (11k) → HD		HD → D-r (11k)	
	rank-1	mAP	rank-1	mAP
VGG19-bn	42.60	46.83	58.87	62.24
ResNet50	58.28	62.08	69.18	72.68
GPA-Net (ours)	83.02	84.65	86.84	89.06

Table 2. Cross-domain performance comparison on right dorsal (D-r) of 11k and HD datasets. Rank-1 accuracy (%) and mAP (%) are shown.

$H_p \setminus V_p$	1	2	3
1	91.19 (92.14)	93.45 (94.40)	92.16 (93.20)
2	93.05 (94.01)	94.06 (94.83)	93.60 (94.21)
3	95.83 (96.31)	95.39 (95.95)	94.07 (94.80)
4	95.42 (95.95)	94.41 (95.08)	93.81 (94.50)

Table 3. Ablation analysis of our method across both horizontal (H_p) and vertical (V_p) partitioning on right palmar (P-r) of 11k hands dataset. The results are shown in *rank-1* (mAP) format in %.

Method	rank-1 (%)	mAP (%)
Global only (ours)	91.19	92.14
Local only (ours)	94.33	94.91
GPA-Net (ours)	95.83	96.31

Table 4. Component analysis of our method on right palmar (P-r) of 11k hands dataset using rank-1 (%) and mAP (%).

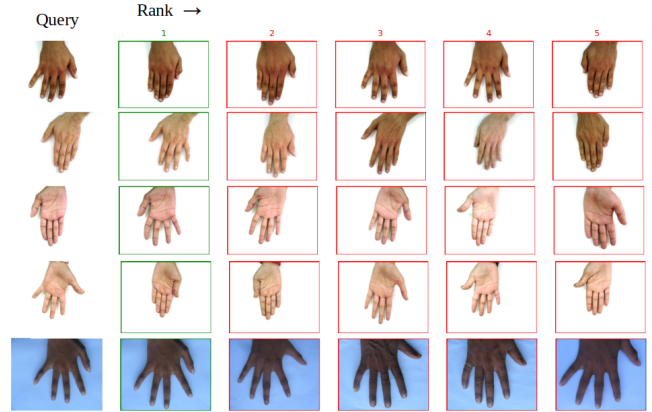


Fig. 3. Some qualitative results of our method using query vs ranked results retrieved from gallery. From top to bottom row are right dorsal of 11k, left dorsal of 11k, right palmar of 11k, left palmar of 11k and HD datasets. The green and red bounding boxes denote the correct and the wrong matches, respectively.

4. CONCLUSION

In this paper, we introduce a Global and Part-Aware Network (GPA-Net) to learn both global and local (part-level) deep feature representations for person identification based on the hand. The combination of global and part-level feature representations allows a deeper study of the features of the hand in less controlled situations and our extended uniform partitioning helps to obviate the need for explicit pose estimation that can result in compound error. The experimental results on two public hand datasets demonstrate the superiority of the proposed method over the competing methods, giving it strong potential for robust identification of the perpetrators of serious crime.

5. REFERENCES

- [1] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, “ArcFace: Additive angular margin loss for deep face recognition,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 4685–4694.
- [2] S. Zhang, X. Zhu, Z. Lei, H. Shi, X. Wang, and S. Z. Li,

- “ S^3FD : Single shot scale-invariant face detector,” in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 192–201.
- [3] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven C. H. Hoi, “Deep learning for person re-identification: A survey and outlook,” 2020.
- [4] Nathanael L. Baisa, “Occlusion-robust online multi-object visual tracking using a GM-PHD filter with CNN-based re-identification,” *Journal of Visual Communication and Image Representation*, vol. 80, pp. 103279, 2021.
- [5] Mahmoud Afifi, “11k hands: gender recognition and biometric identification using a large dataset of hand images,” *Multimedia Tools and Applications*, 2019.
- [6] Abdelouahab Attia, Zahid Akhtar, and Youssef Chahir, “Feature-level fusion of major and minor dorsal finger knuckle patterns for person authentication,” *Signal, Image and Video Processing*, Feb. 2021.
- [7] J. J. Engelsma, K. Cao, and A. K. Jain, “Learning a fixed-length fingerprint representation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2019.
- [8] Z. Zhao and A. Kumar, “Towards more accurate iris recognition using deeply learned spatially corresponding features,” in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 3829–3838.
- [9] D. Muramatsu, A. Shiraiishi, Y. Makihara, M. Z. Uddin, and Y. Yagi, “Gait-based person recognition using arbitrary view transformation model,” *IEEE Transactions on Image Processing*, vol. 24, no. 1, pp. 140–154, 2015.
- [10] Joon Son Chung, Arsha Nagrani, and Andrew Senior, “Voxceleb2: Deep speaker recognition,” *Interspeech 2018*, Sep 2018.
- [11] A. Dantcheva, P. Elia, and A. Ross, “What else does your biometric data reveal? A survey on soft biometrics,” *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 3, pp. 441–467, 2016.
- [12] Mei Wang and Weihong Deng, “Deep face recognition: A survey,” 2020.
- [13] Maneet Singh, Richa Singh, and Arun Ross, “A comprehensive overview of biometric fusion,” *Information Fusion*, vol. 52, pp. 187 – 205, 2019.
- [14] Sanchit, M. Ramalho, P. L. Correia, and L. D. Soares, “Biometric identification through palm and dorsal hand vein patterns,” in *2011 IEEE EUROCON - International Conference on Computer as a Tool*, 2011, pp. 1–4.
- [15] Yiding Wang, Heng Cao, Xiaochen Jiang, and Yuanyan Tang, “Recognition of dorsal hand vein based bit planes and block mutual information,” *Sensors*, vol. 19, no. 17, pp. 3718, 2019.
- [16] Yimin Yuan, Chaoying Tang, Shuhang Xia, Zhou Chen, and Tong Qi, “HandNet: Identification based on hand images using deep learning methods,” in *Proceedings of the 2020 4th International Conference on Vision, Image and Signal Processing*, New York, NY, USA, 2020, ICVISP 2020, Association for Computing Machinery.
- [17] C. Su, J. Li, S. Zhang, J. Xing, W. Gao, and Q. Tian, “Pose-driven deep convolutional model for person re-identification,” in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 3980–3989.
- [18] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [19] A. Kumar and Z. Xu, “Personal identification using minor knuckle patterns from palm dorsal surface,” *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 10, pp. 2338–2348, 2016.
- [20] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [21] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A. Alemi, “Inception-v4, inception-resnet and the impact of residual connections on learning,” in *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*. 2017, AAAI’17, p. 4278–4284, AAAI Press.
- [22] H. Luo, Y. Gu, X. Liao, S. Lai, and W. Jiang, “Bag of tricks and a strong baseline for deep person re-identification,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2019, pp. 1487–1495.
- [23] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 2818–2826.
- [24] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, and Q. Tian, “Scalable person re-identification: A benchmark,” in *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 1116–1124.

- [25] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun, Eds., 2015.