IMAVIS Special Issue on Biometrics in the Wild
Editorial

Abstract

This document contains the editorial for the special issue.

Biometric recognition from data captured in unconstrained settings, commonly referred to as biometric recognition in the wild, represents a challenging and highly active area of research. The interest in this area is fueled by the numerous application domains that deal with unconstrained data acquisition conditions, such as forensics, surveillance, social media, consumer electronics or border control. While the existing biometric technology has matured to a point where excellent performance is nowadays regularly reported for various tasks in ideal laboratory-like settings, many problems related to so-called *in-the-wild* scenarios still require further research and novel ideas. The goal of this special issue was to collect the most advanced work related to these types of scenarios and to introduce novel solutions to open biometrics-related problems.

The special issue was organized in conjunction with the 2^{nd} *International Workshop on Biometrics in the Wild 2017* and accepted submissions from workshop participants, but also other researchers working in this area. In response to the call-for-papers, the special issue received a total of 32 submissions that underwent a rigorous peer-reviewing process overseen by Guest Editors. A total of 12 highest quality papers were finally selected for publication. These 12 papers cover a wide variety of topics ranging from face recognition, facial trait classification and 3D face reconstruction to iris matching, gait recognition and re-identification. The selection of papers also includes two survey papers that review: 1) existing work on facial aging and kinship recognition, and 2) methods that exploit contextual information to aid biometric recognition systems. In accordance with recent trends in computer vision and related areas, the majority of papers focus on deep learning methodologies, which seem particularly suitable for solving challenges typically encountered in unconstrained data-acquisition settings.

The first special issue paper by *Trigueros*, *Meng* and *Hartnett* titled "Enhancing Convolutional Neural Networks for Face Recognition with Occlusion Maps and Batch Triplet Loss" ad-dresses the problem of unconstrained face recognition with con-volutional neural networks (CNNs). The authors describe two main contributions: *i)* a novel mechanism for identifying dis-criminative facial areas in images and a corresponding proce-dure for incorporating this information into the CNN learning procedure, and *ii)* a novel loss function that extends the standard triplet loss with an extra term and leads to improved recogni-tion performance. Experiments are presented on the AR and Labeled Faces in the Wild (LFW) datasets with highly encouraging results.

The second paper "Patch-based Face Recognition using a Hierarchical Multi-label Matcher" by *Zhang*, *Dou* and *Kakadiaris* introduces an inventive approach for patch-based face recognition. The proposed approach (iteratively) partitions an input face image into a hierarchy of multi-level image patches and then encodes the generated patches using a selected feature extraction technique. Next, a class-label is assigned to each of the patches resulting in a hierarchical label set that is associated with each input image. Finally, a hierarchical multi-label matcher is introduced, which considers patch relationships to refine local (patch-level) matching results and then aggregates the refined scores across all levels to produce the final matching result. The proposed approach is not limited to a specific feature extraction technique and matcher, so experiments are presented for multiple choices. Competitive performance is reported on the Extended Yale B, AR, UHDB31 and IJB-A datasets.

The third paper by *Wu*, *Liu*, *Li* and *Fu* entitled "Improving Face Representation Learning with Center Invariant Loss" focuses on efficient face representation learning with imbalanced training data. The authors propose a new CNN learning objective called center invariant loss that tries to align the centers of all classes in the training data evenly around the global mean - regardless of the amount of images in any given class. Extensive experiments are presented on the Labeled Faces in the Wild (LFW) and YouTube Faces (YTF) benchmarks using a 22-layer ResNet as the basis for the experimental evaluation. Results show consistent improvements over the baseline model when the proposed loss is used as an additional supervision signal next to standard learning objectives.

The fourth special issue paper "Learning CNNs from Weakly Annotated Facial Images" by *Franc* and *Čech* revolves around the problem of learning convolutional neural networks (CNNs) for facial analysis tasks from weakly annotated data. The authors assume that a pair of labels (i.e., an attribute and an identity label) is available for an input image with multiple faces, but the target face for the labels is unknown. To train a CNN-based attribute predictor on such data, they formulate the learning procedure as a maximum-likelihood-estimation problem and solve it by using an expectation-maximization algorithm that learns the predictor in its inner loop. Cross-dataset experiments on age and gender estimation using the IMDB,

APPA-REAL, ChaLearnAge, FG-NET, and LFW demonstrate the efficacy of the proposed solution and its superiority over heuristic approaches that are otherwise used for the studied task.

In the fifth paper "Structured Deep Fisher Pruning for Efficient Facial Trait Classification" *Tian*, *Arbel* and *Clark* present a novel pruning approach designed to reduce the model size and computational complexity of the existing CNNs. The structured filter-level pruning approach described by the authors relies on Fisher's Linear Discriminant Analysis (LDA) to identify less informative filters in the convolutional layers of the CNN models which can be discarded without loss of performance. As shown in facial-trait-classification experiments on the LFWA and CelebA datasets, the proposed method is able to ensure reductions in space of over 95% for the VGG-16 model and 81% for GoogLeNet and computational savings as high as 80% for VGG-16 and 61% for GoogLeNet, respectively.

In the sixth special issue paper by *Boulkenafet*, *Komulainen*, and *Hadid* titled "On the generalization of color texture-based face anti-spoofing" the authors study the problem of presentation attack detection (PAD) in face recognition systems. Specifically, they perform a rigorous in-depth analysis of seven color texture-based methods in challenging cross-database experiments and present important findings about the generalization ability of the tested techniques. Based on the findings, they develop an attack-specific approach to cope with the problem of generalized face PAD and report highly competitive cross-dataset results with three benchmark face spoofing datasets.

The seventh paper entitled "Multi-View 3D Face Reconstruction with Deep Recurrent Neural Networks" by *Dou* and *Kakadiaris* introduces a novel approach for 3D face reconstruction that relies on a subspace representation of the 3D face shape and a combination of a deep convolutional neural network (DCNN) and a recurrent neural network (RNN). The convolutional part of the approach disentangles facial identity from other facial attributes for each input image, while the recurrent part of the approach aggregates the identity information over the entire set of images. The reported results show considerable improvements in terms of accuracy and consistency of the recovered 3D face reconstructions when compared to the state-of-the-art.

The eighth special issue paper titled "Found a good match: should I keep searching? - Accuracy and Performance in Iris Matching Using 1-to-First Search" by *Kuehlkamp* and *Bowyer* analyzes different optimization strategies that enable fast searches in large scale iris databases. Specifically, they look at the characteristics of the 1:First search strategy, where the search is terminated after the first sufficiently good match, and compare it to the traditional 1:N search strategy. Different scenarios are considered and important findings are reported.

In the ninth paper titled "Gait Recognition in the Wild using Shadow Silhouettes" by *Verlekar*, *Soares* and *Correia* presents a gait recognition system that, unlike competing ap-proaches, does not examine the gait patterns of individuals di-rectly, but rather analyzes the shadow silhouettes of people with the goal of extracting suitable gait features for recognition. The authors show that the proposed approach matches the perfor-mance of state-of-the-art methods in constrained settings, but generalizes significantly better to unconstrained settings than competing methods from the literature.

The tenth paper by *Chen*, *Duffner*, *Stoian*, *Dufour*, and *Baskurt* entitled, "Deep and Low-level Feature based Attribute Learning for Person Re-identification," describes a CNN-based attribute assisted person re-identification framework. The framework relies on part-based features learned by a CNN model and hand-crafted Local Maximal Occurrence (LOMO) features to build an efficient attribute classifiers, and then combines the learned attribute representations with identity features of another CNN for person re-identification. In re-identification experiments on the CUHK03 dataset the authors show that the proposed combination of identity and attribute features compares favorably to the state-of-the-art.

The last two special issue papers are surveys on topics of considerable importance for the field of biometrics. In the paper "Modelling of Facial Aging and Kinship: A Survey" *Georgopoulos*, *Panagakis* and *Pantic* review recent advances in modelling of facial aging and kinship. The authors present a survey of the existing computational models used in the two covered areas, discuss popular datasets and analyze existing evaluation protocols and performance metrics commonly used for experimentation. In the last paper "Context Awareness in Biometric Systems and Methods: State of the Art and Future Scenarios", the authors, *Nappi*, *Ricciardi* and *Tistarelli*, survey another interesting aspect of unconstrained biometric recognition, i.e., the use of contextual information. The authors cover the concept of context-awareness across different biometric modalities, introduce a taxonomy of context-aware-biometric systems, and critically analyze and discuss the characteristics of the existing state-of-the-art with respect to the use of contextual information.

We would like to thank everyone involved in this special issue for their valuable contributions. We hope the selection of papers in this issue will inspire new ideas in the area of unconstrained biometrics and help to spark interest in this challenging field of research.

Guest Editors:

Vitomir Štruc
University of Ljubljana, Slovenia

Bir Bhanu
University of California, Riverside, United States

Abdenour Hadid
University of Oulu, Finland

Qiang Ji
Rensselaer Polytechnic Institute, United States

Mark Nixon
University of Southampton, United Kingdom

Advisory Editors:

Rama Chellappa
University of Maryland, United States

Josef Kittler
University of Surrey, United Kingdom