Ear Recognition: More Than a Survey

Žiga Emeršič, Vitomir Štruc, Peter Peer

Abstract—**Automatic identity recognition from ear images represents an active field of research within the biometric community. The ability to capture ear images from a distance and in a covert manner makes the technology an appealing choice for surveillance and security applications as well as other application domains. Significant contributions have been made in the field over recent years, but open research problems still remain and hinder a wider (commercial) deployment of the technology. This paper presents an overview of the field of automatic ear recognition (from 2D images) and focuses specifically on the most recent, descriptor-based methods proposed in this area. Open challenges are discussed and potential research directions are outlined with the goal of providing the reader with a point of reference for issues worth examining in the future. In addition to a comprehensive review on ear recognition technology, the paper also introduces a new, fully unconstrained dataset of ear images gathered from the web and a toolbox implementing several state-of-the-art techniques for ear recognition. The dataset and toolbox are meant to address some of the open issues in the field and are made publicly available to the research community.**

Index Terms—**biometry, dataset, in-the-wild, unconstrained image, descriptor-based method, open-source toolbox, ear recognition**

I. Introduction

Ear images used in automatic ear recognition systems can typically be extracted from profile head shots or video footage. The acquisition procedure is *contactless* and *nonintrusive* and also *does not depend on the cooperativeness* of the person one is trying to recognize. In this regard ear recognition technology shares similarities with other image-based biometric modalities. Another appealing property of ear biometrics is its *distinctiveness* [1]. Recent studies even empirically validated existing conjectures that certain features of the ear are distinct for identical twins [2]. This fact has significant implications for security related applications and puts ear images on par with epigenetic biometric modalities, such as the iris. Ear images can also serve as *supplements* for other biometric modalities in automatic recognition systems and provide identity cues when other information is unreliable or even unavailable. In surveillance applications, for example, where face recognition technology may struggle with profile faces, the ear can serve as a source of information on the identity of

Ž. Emeršič and Peter Peer are with the Faculty of Computer and Information Science, University of Ljubljana, e-mail: {ziga.emersic, peter.peer}@fri.uni-lj.si.

V. Štruc is with the Faculty of Electrical Engineering, University of Ljubljana, e-mail: vitomir.struc@fe.uni-lj.si.

people in the surveillance footage. The importance and potential value of ear recognition technology for multi-modal biometric systems is also evidenced by the number of research studies on this topic, e.g. [3], [4], [5], [6], [7].

Today, ear recognition represents an active research area, for which new techniques are developed on a regular basis and several datasets needed for training and testing of the technology are publicly available, e.g., [8], [9]. Nevertheless, despite the research efforts directed at ear biometrics, to the best of our knowledge, there is only one commercial system currently available on the marked that exploits ear biometrics for recognition, i.e., the Helix from Descartes Biometrics [10]. We conjecture that the limited availability of the commercial ear recognition technology is a consequence of the open challenges that by today have still not been appropriately addressed. This paper is an attempt to meet some of these challenges and provide the community with a point of reference as well as with new research tools that can be used to further advance the field.

A. Contributions and Paper Organization

Prior surveys related to ear recognition, such as [11], [12], [13], [14], [15], [16], provide well written and well structured reviews of the field. In this paper we contribute to these surveys by discussing recent 2D ear recognition techniques proposed until the end of 2015. We pay special attention to descriptor-based approaches that are currently considered state-of-the-art in 2D ear recognition. We present comparative experiments with the new dataset and toolbox to establish an independent ranking of the state-of-the-art techniques and show that there is significant room for improvement and that ear recognition is far from being solved.

We make the following contributions in this paper:

- **Survey:** We present a comprehensive survey on ear recognition, which is meant to provide researchers in this field with a recent and up-to-date overview of the state-of-technology. We introduce a taxonomy of the existing 2D ear recognition approaches, discuss the characteristics of the technology and review the existing state-of-the-art. Most importantly, we elaborate on the open problems and challenges faced by the technology.
- **Dataset:** We make a new dataset of ear images available to the research community. The dataset, named Annotated Web Ears (AWE), contains images collected from the web and is to the best of our knowledge the first dataset for ear recognition gathered "in the wild". The images of the AWE dataset contain a
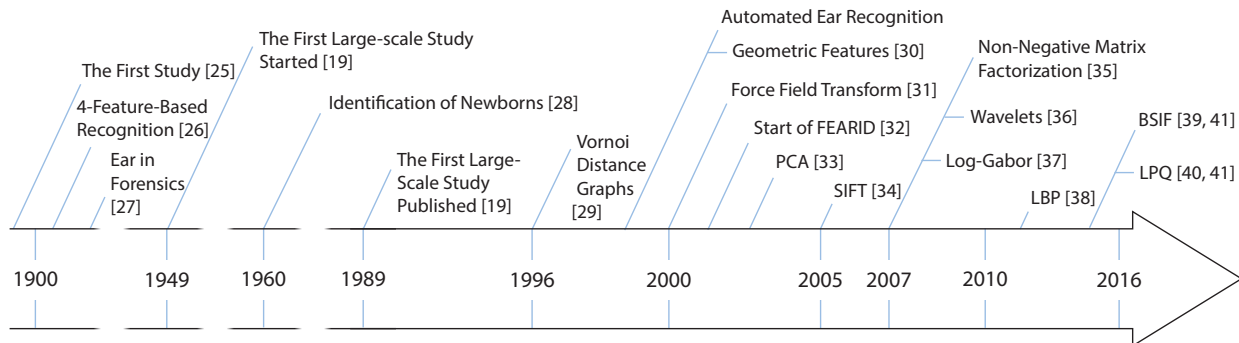
Figure 1: Development of ear recognition approaches through time. Selected milestones are presented.

high degree of variability and present a challenging problem to the existing technology, as shown in the experimental section.

- **Toolbox:** We introduce an open source (Matlab) toolbox, i.e., the AWE toolbox, for research in ear recognition. The toolbox implements a number of state-of-the-art feature extraction techniques as well as other important steps in the processing pipeline of ear recognition systems. It contains tools for generating performance metrics and graphs and allows for transparent and reproducible research in ear recognition. The toolbox is available from: http://awe.fri.uni-lj.si.

- **Reproducible evaluation:** We conduct a comparative evaluation of several state-of-the-art methods on a number of popular ear datasets using consistent experimental protocols, which enables direct comparisons of the state-of-the-art in ear recognition. All experiments are conducted with our AWE toolbox making all presented results reproducible.

The rest of the paper is structured as follows. In Section II we present the background and basic terminology related to ear recognition. In Section III existing ear datasets are discussed and compared on the basis of some common criteria of interest. Section IV introduces our new AWE dataset and the accompanying AWE toolbox. Comparative experiments and results with the new toolbox are presented in Section V. In Section VI open problems and promising future research directions are examined. The paper concludes with some final comments in Section VII.

II. Ear Recognition Essentials

A. Ear Structure

The human ear develops early during pregnancy and is already fully formed by the time of birth. Due to its role as the human hearing organ, the ear has a characteristic structure that is (for the most part) shared across the population. The appearance of the outer ear is defined by the shapes of the tragus, the antitragus, the helix, the antihelix, the incisura, the lope and other important structural parts as shown in Fig. 2. These anatomical cartilage formations differ in shape, appearance and relative
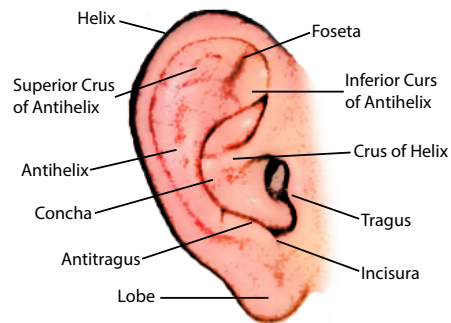

Figure 2: Ear structure.

positions from person to person and can, therefore, be exploited for identity recognition.

In general, the left and right ears of a person are similar to such an extent that makes matching the right ear to the left ear (and vice versa) with automatic techniques perform significantly better than chance. Yan and Bowyer [17], for example, reported a recognition performance of around 90% in cross-ear matching experiments. They observed that for most people the left and right ears are at least close to bilateral symmetric, though the shape of the two ears is different for some [17]. Similar findings were also reported by Abaza and Ross in [18].

The literature suggests that the size of the ear changes through time [19], [15], [16], [20]. Longitudinal studies from India [21] and Europe [20], [22], [23] have found that the length of the ear increases significantly with age for men and women, while the width remains relatively constant. How ear growth affects the performance of automatic recognition system is currently still an open research question. The main problem here is the lack of appropriate datasets captured over a long enough period of time that could help provide final and conclusive answers. Some initial studies appeared recently on this topic, but only featured images captured less than a year apart [24].

B. Chronological Development of Ear Recognition

The chronological development of ear recognition techniques can be divided into a *manual (pre-automatic)* and *automatic era*. During the pre-automatic era several studies and empirical observations were published pointing

to the potential of ears for identity recognition [25], [26], [27], [28]. One of the biggest contributions to the field during this era was made by Iannarelli in 1989 [19], when he published a long-term study on the potential of ear recognition. Iannarelli's seminal work included more than $10,000$ ears and addressed various aspects of recognition, such as ear similarity of siblings, twins and triplets, relations between the appearance of the ears of parents and children as well as racial variations of ear appearance [11].

The 1990s marked the beginning of automatic ear recognition. Various methods were developed during this time and were introduced in the literature. In 1996, for example, Burge and Burger [29] used adjacency graphs computed from Voronoi diagrams of the ears curve segments for ear description and in 1999 Moreno et al. [30] presented the first fully automated ear recognition procedure exploiting geometric characteristics of the ear and a compression network. In 2000 Hurley et al. [31] described an approach for ear recognition that relied on the Force Field Transform, which proved highly successful for this task. A year later, in 2001, the Forensic Ear Identification Project (FEARID) project was launched, marking the first large-scale project in the field of ear recognition [32].

With the beginning of the new millennium, automatic ear recognition techniques started to gain traction with the biometric community and new techniques were introduced more frequently. In 2002 Victor et al. [33] applied principal component analysis (PCA) on ear images and reported promising results. In 2005 the Scale Invariant Feature Transform (SIFT) [34] was used for the first time with ear images, raising the bar for the performance of the existing recognition techniques. In 2006 a method based on non-negative matrix factorization (NMF) was developed by Yuan et al. [35] and applied to occluded and non-occluded ear images with competitive results. In 2007 a method based on the 2D wavelet transform was introduced by Nosrati et al. [36], followed by a technique based on log-Gabor wavelets in the same year [37]. More recently, in 2011, local binary patterns (LBP) were used for ear-image description in [38], while later binarized statistical image features (BSIF) and local phase quantization (LPQ) features also proved successful for this task [39], [40], [41]. A graphical representation of the main milestones[1] in the development of ear recognition technology (briefly discussed above) is shown in Fig. 1.

C. Ear Recognition Approaches

Techniques for automatic identity recognition from ear images can in general be divided into techniques operating on either 2D or 3D ear data. Here, we focus only on 2D approaches and refer the reader to other reviews, such as [16], for a detailed coverage of the field of 3D ear recognition.

A typical (2D) ear recognition system operates on images or video footage captured with commercial of-the-shelf cameras, surveillance systems, CCTV or similar everyday hardware. In a fully automatic setting, the system
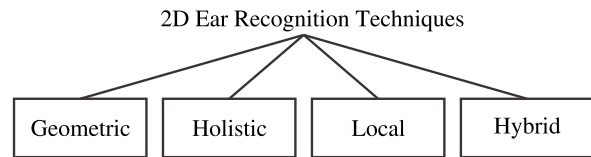
---

[1]In the opinion of the authors.

Figure 3: Taxonomy of 2D ear recognition approaches used in this paper.

first detects a profile face in the input image and segments the ear from the detected region. The ear is then aligned to some predefined canonical form and processed to account for potential variability in illumination. Features are extracted from the pre-processed image and finally identity recognition is conducted based on the computed features using a suitable classification technique.

Depending on the type of feature extraction technique used, 2D ear recognition approaches can be divided into *geometric*, *holistic*, *local* and *hybrid* methods, as illustrated in Fig. 3. Geometric techniques exploit the geometrical characteristics of the ear, such as the shape of the ear, locations of specific ear parts and relations between these parts. Holistic approaches treat the ear as a whole and consider representations that describe the global properties of the ear. Local approaches describe local parts or the local appearance of the ear and exploit these descriptions for recognition. The last category of hybrid approaches covers techniques that do not strictly fit into any of the other categories and usually comprises methods that combine elements from all categories or rely on multiple representations to improve performance.

Table I presents a comparison of the techniques surveyed in this work. The techniques are ordered chronologically and labeled in accordance with our taxonomy. Information about the datasets used to evaluate the techniques as well as the reported performance is also provided. It should be noted that the results in Table I are not directly comparable with each other, as authors typically use their own experimental protocols for most datasets. The label "Own" in the column "Dataset" denotes the fact that the authors used their own dataset, while the label "NA" indicates that the information was not provided by the authors. The performance of the techniques is given in terms of the rank-1 recognition rate (R1), the equal error rate (EER) or the verification rate (VR) [1].

1) Geometric Approaches: Similar to other fields of computer vision, early approaches to automatic ear recognition focused mainly on the extraction and analysis of the geometrical features of ear images. Techniques from this category are in general computationally simple and often rely on edge detection as a pre-processing step. Typically, edge information is exploited to describe the geometric properties of ears or derive geometry-related statistics that can be used for recognition. Only information pertaining to the geometry of the ear is used, which makes it easy to devise methods invariant to geometric distortions, such as rotation, scaling or even small perspective changes. On the down side, most of the texture information that could

Table I: Comparative summary of 2D ear recognition techniques. The column "Method – Short Description" provides a brief description of the technique used by the authors. The column "Type" defines the category of the approach in accordance with our taxonomy. The column "Dataset" presents the dataset that was used for testing, while the column "#Sub." stands for the number of subjects in the datasets. The last column "Perf." denotes the reported performance.

Reference	Method – Short Description	Type	Dataset	#Sub.	Perf. [%]
1996, Burge and Burger [29]	Adjacency Graphs of Voronoi Diagrams	Geometric	Own	NA	NA
1999, Moreno et al. [30]	Geometric Features | Morphological Description	Geometric	Own	48	43 | 83 (R1)
2000, Hurley et al. [31]	Force Field Transform	Holistic	Own	NA	NA
2002, Victor et al. [33]	Principal Component Analysis – PCA	Holistic	Own	294	40 (R1)
2003, Chang et al. [42]	PCA	Holistic	UND E	114	71.6 (R1)
2004, Mu et al. [43]	Geometrical Measures on Edge Images	Geometric	USTB II	77	85 (R1)
2005, Zhang et al. [44]	Independent Component Analysis – ICA	Holistic	CP	17	94.1 (R1)
			Own	60	88.3 (R1)
2006, Abate et al. [45]	Generic Fourier Descriptor	Holistic	Own	70	96 (R1)
2006, Ab.-Mottaleb et al. [46]	Force Field and Contour Extraction	Holistic	Own	29	87.9 (R1)
2006, Choras and Choras [47]	Geometrical Approaches on Longest Ear Contours	Geometric	Own	NA	100 (R1)
2006, Dewi and Yahagi [48]	SIFT	Local	CP	17	78.8 (R1)
2006, Yuan et al. [35]	Non-Negative Matrix Factorization	Holistic	USTB II	77	91 (R1)
2007, Arbab-Zavar et al. [49]	SIFT with Pre-filtered Points	Local	XM2VTS	63	91.5 (R1)
2007, Kumar and Zhang [37]	Log-Gabor Wavelets	Local	UND	113	90 (R1)
2007, Nanni and Lumini[50]	Multi-Matcher	Hybrid	UND E	114	80 (R1)
2007, Nosrati et al. [36]	Wavelet Transformation and PCA	Hybrid	USTB II	77	90.5 (R1)
			CP	17	95.0 (R1)
2007, Rahman et al. [51]	Geometric Features	Geometric	Own	100	88 (R1)
2007, Yuan and Mu [52]	Full-Space LDA	Holistic	USTB II	77	90 (R1)
2008, Arbab-Zavar et al. [53]	Log-Gabor Filters	Local	XM2VTS	63	85.7 (R1)
2008, Choras [54]	Geometry of Ear Outline	Geometric	Own	188	86.2 (R1)
2008, Dong and Mu [55]	Force Field Transform and NKDFA	Holistic	USTB IV	29	75.3 (R1)
2008, Guo and Xu [56]	Local Binary Pattern and Cellular NN	Local	USTB II	77	93.3 (R1)
2008, Naseem et al. [57]	Sparse Representation	Holistic	UND	32	96.9 (R1)
2008, Wang et al. [58]	Haar Wavelet Transformation and LBP	Hybrid	USTB III	79	92.4 (R1)
2008, Xie an Mu [59]	Enhanced Locally Linear Embedding	Holistic	USTB III	79	90 (R1)
2008, Zhang and Liu [60]	Null Kernel Discriminant Analysis	Holistic	USTB I	60	97.7 (R1)
2009, H.-Long and Z.-Chun [61]	Wavelet Transformation	Local	USTB II	77	85.7 (R1)
			USTB III	79	97.2 (R1)
2009, Nanni and Lumini [62]	Gabor Filters	Local	UND E	114	84 (R1)
2009, Xiaoyun and Weiqi [63]	Block Partitioning and Gabor Transform	Local	USTB II	60	100 (R1)
2010, Alaraj et al. [64]	PCA and NN	Holistic	UND	17	96 (R1)
2010, Bustard and Nixon [65]	SIFT	Local	XM2VTS	63	96 (R1)
2010, De Marsico et al. [66]	Partitioned Iterated Function System	Local	UND E	114	59 (R1)
			FERET	100	93 (R1)
2010, Guiterrez et al. [67]	Wavelet Transform and NN	Local	USTB II	77	97.5 (R1)
2012, Chan and Kumar [68]	2D Quadrature Filter	Local	IITD I	125	96.5 (R1)
			IITD II	221	96.1 (R1)
2012, Kumar and Wu [9]	Orthogonal Log-Gabor Filter Pairs	Local	IITD II	221	95.9 (R1)
2013, Baoqing et al. [69]	Sparse Representation Classification – SRC	Holistic	USTB III	79	>90 (R1)
2013, Kumar and Chan [70]	SRC of Local Gray-Level Orientations	Hybrid	IITD I	125	97.1 (R1)
			IITD II	221	97.7 (R1)
2013, Lakshmanan [71]	Multi-Level Fusion	Hybrid	USTB II	77	99.2 (VR)
2013, Prakash and Gupta [72]	Enhanced SURF with NN	Local	IITK I	190	2.8 (EER)
2014, Basit and Shoaib [73]	Non-Linear Curvelet Features	Local	IITD II	221	96.2 (R1)
2014, Benzaoui et al. [39]	BSIF	Local	IITD II	221	97.3 (R1)
2014, Galdamez et al. [74]	Hybrid – based on SURF, LDA and NN	Hybrid	Own	NA	97 (R1)
2014, Jacob and Raju [75]	Gray Level Co-Occurrence + LBP + Gabor Filter	Hybrid	IITD II	221	94.1 (R1)
2014, Pflug et al. [40]	LPQ	Local	Several	555	93.1 (R1)
2014, Pflug et al. [41]	LPQ, BSIF, LBP and HOG – all with LDA	Hybrid	UND J2	158	98.7 (R1)
			AMI	100	100 (R1)
			IITK	72	99.4 (R1)
2014, Ying et al.[76]	Weighted Wavelet Transform and DCT	Hybrid	Own	75	98.1 (R1)
2015, Benzaoui et al.[77]	LBP and Haar Wavelet Transform	Hybrid	IITD	121	94.5 (R1)
2015, Benzaoui et al.[78]	BSIF	Local	IITD I	125	96.7 (R1)
			IITD II	221	97.3 (R1)
2015, Bourouba et al.[79]	Multi-Bags-Of-Features Histogram	Local	IITD I	125	96.3 (R1)
2015, Meraoumia et al.[80]	Gabor Filters	Local	IITD II	221	92.4 (R1)
2015, Morales et al.[81]	Global and Local Features on Ear-prints	Hybrid	FEARID	1200	91.3 (R1)

provide valuable discriminative information is discarded and not used for recognition.

Many geometric approaches were presented in the literature and exploite: Voronoi diagrams of curve segments extracted from ear images [29], coordinates of character-

istic ear points [30], morphology descriptors of the ear shape and the structure of wrinkles [30], relative relationships between selected points on the ear contour [43], [51], contour tracing techniques [47], angle-based contour representations [47], geometric parameterizations of the longest contour found in the ear image [47] or geometric properties of subparts of ears [54].

One limitation of geometric approaches is their dependency on edge detectors, which are known to be susceptible to illumination variations and noise. Many techniques also require exact locations of specific ear points that may be difficult to detect in poor quality images or when an ear is occluded [16], [82].

2) Holistic Approaches: Holistic approaches differ conceptually from geometric techniques discussed in the previous section. Instead of geometric properties, holistic approaches exploit the global appearance of the ear and compute representations from input images that encode the ear structure as a whole. As the appearance of an ear changes significantly with pose or illumination changes, care needs to be taken before computing holistic features from images and normalization techniques need to be applied to correct for these changes prior to feature extraction. Holistic techniques were reported to provide competitive performance on well aligned and well preprocessed images, but may experience difficulties if this is not the case.

One popular (holistic) approach to ear recognition was presented in [31], where the so-called Force Field Transform was introduced. The technique computes a force field from the input ear image by treating the pixels as sources of a Gaussian force field. The force field is then analyzed and its properties are exploited to compute similarities between ear images. Force-field-transform-based techniques proved very successful and were extended several times, e.g., in [34], [46], [55], [83].

Linear and nonlinear subspace projection techniques have also been introduced to the field of ear recognition. With subspace projection techniques, ear images are represented in the form of weights of a linear combination of some basis vectors in either the input (pixel) space or some high-dimensional Hilbert space. Existing examples of techniques from this group include: principal component analysis (PCA) [33], [42], [64], independent component analysis (ICA) [44], non-negative matrix factorization (NMF) [35], full space linear discriminant analysis (LDA) [52], enhanced locally linear embedding (ELLE) [59], or null kernel discriminant analysis (NKDA) [60]. A similar idea is also exploited with sparse representations of ear images, where images are represented in the form of a sparse linear combination of the training images (or some other dictionary) [57], [69].

Holistic techniques operating in the frequency domain have been used for ear recognition as well. One such example is the technique of Abate et al. [45], which computes a global image descriptor in the frequency domain to represent ear images and is able to achieve rotation invariance by representing the ear image in polar coordinates prior to descriptor computation.

3) Local Approaches: Local approaches for ear recognition extract features from local areas of an image and use these features for recognition. Different from geometric approaches, local approaches do not rely on the locations of specific points or relations between them, but on the description of the local neighborhood (or area) of some points in the image. The points of interest to local approaches must not necessary correspond to structurally meaningful parts of the ear, but can in general represent any point in the image. Two types of techniques can be considered local in this context: *i)* techniques that first detect keypoint locations in the image and then compute separate descriptors for each of the detected keypoints [84], and *ii)* techniques that compute local descriptors densely over the entire image.

Local techniques (together with hybrid approaches) currently achieve the most competitive results in the field of ear recognition. Unlike geometric or holistic approaches, local techniques mostly encode texture information [84]. Many of the techniques from this category exhibit some level of robustness to small image rotations and alignment errors owing to the fact that histograms are computed over smaller image areas for most descriptors. The computational complexity of many low-level descriptors, such as local binary patterns (LBPs), is low and the descriptors can be computed extremely efficiently. However, others, such as Gabor features, for example, are computationally more demanding.

Examples of techniques that first detect keypoint locations in ear images and then extract SIFT (scale invariant feature transform) descriptors around these keypoints were presented in [48], [53] and [65]. An approach similar in nature albeit based on speeded up robust features (SURF) was presented in [72]. A desirable characteristic of this group of techniques is the independent description of keypoint locations, which makes partial matching of the extracted descriptors possible and allows designing methods that exhibit robustness to partial occlusions of the ear. On the other hand, information about the global ear structure that could help with the recognition process is typically lost with this group of methods.

Techniques from the second group compute local descriptors in a dense manner without first detecting keypoints in an image. The resulting image representation is comprised of local descriptors that still encode the global structure of the image, but at the expense of loosing robustness to partial occlusions (in most cases). Typical examples of techniques from this group rely on: Gabor or log-Gabor filters and other types of wavelet representations [37], [9], [53], [61], [62], [63], [67], [68], [80], curvelet representations [73], dense SIFT (DSIFT) [85], local binary patterns (LBP) [56], local phase quantization features (LPQ) [40] and rotation invariant LPQs (RILPQ) [86], histograms of oriented gradients (HOG) [41], [87], patterns of oriented edge magnitudes POEM) [88] or learned local descriptors such as the binarized statistical image features (BSIF) [39], [78]. Other methods that also fall into this

group are based on fractals [66] or the bag-of-features model [79].

4) Hybrid Approaches: The last category of methods in our taxonomy covers methods that combine elements from other categories or use multiple representations to increase the recognition performance.

One such approach was introduced by Nosrati et al. [36] and combined wavelets and PCA. Another technique, presented in [58], used LBPs on images transformed with the Haar wavelet transform. Kumar and Chan [70] adopted the sparse representation classification algorithm and applied it to local gray-level orientation features, while Galdamez et al. [74] first computed SURF features from the images and then reduced the dimensionality using LDA. A combination of the wavelet transform and the discrete cosine transform (DCT) was presented in [76] and a hybrid method based on LBPs and the Haar transform was introduced in [77]. Various combinations of local descriptors and subspace projection techniques were assessed in [41].

A hybrid technique trying to improve performance by using multiple matchers was presented by Nanni and Lumini in [50]. Lakshmanan described an approach based on multi-level fusion in [71]. Jacob and Raju [75] combined gray-level co-occurrence Matrices, LBPs and Gabor filters for ear recognition, while Morales et al. [81] tried to exploit global and local features to reliably recognize ear-prints.

Hybrid ear recognition techniques are often quite competitive, but depending on their building blocks may also be computationally more complex than simpler holistic or local techniques.

III. Existing Datasets

In this section we present an overview of the existing ear datasets that can be used for training and evaluation of 2D ear recognition approaches. We specifically focus on datasets of cropped ear images suitable for studying ear recognition rather than ear detection approaches. A comparative summary of the most popular ear datasets is provided in Table II. The datasets have different characteristics and the corresponding ear images exhibit different levels of variability, as also illustrated in Fig. 4.

A. CP Ear Dataset

The Carreira-Perpinan (CP) ear dataset [89] represents one of the earliest publicly available datasets for ear recognition. Introduced in 1995, the dataset contains 102 ears of 17 distinct subjects. All images have been captured in controlled conditions, so most of the image variability is due to minor pose variations and, of course, subject identity. A few sample images from the dataset are also shown at the top of Fig. 4.

B. IITD Ear Dataset

The ear dataset of the Indian Institute of Technology Delhi consists of two sub-datasets [9]. The first contains 493 gray-scale images of 125 subjects and the second 793 images of 221 subjects. All images were taken at different
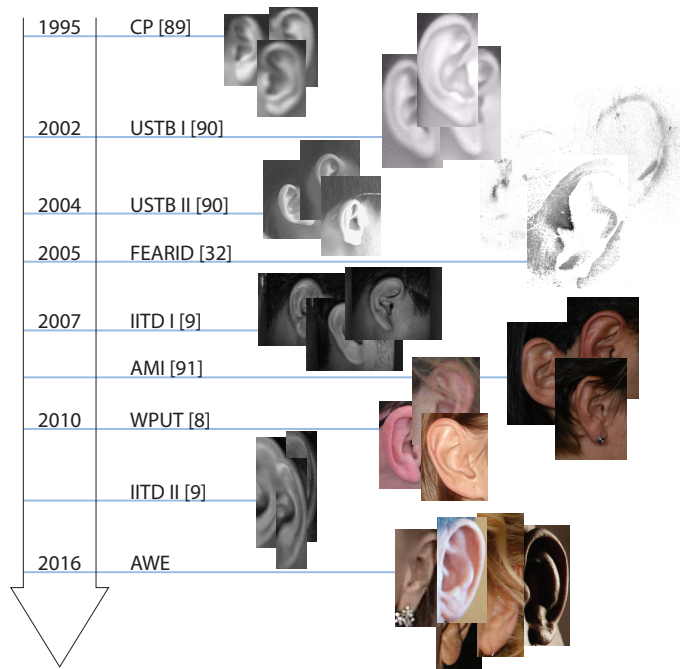

Figure 4: Sample images of available ear datasets. The figure also shows the development of the datasets through time. Note how the datasets have progressed towards more realistic imaging conditions.

indoor lightning conditions and from approximately the same profile angle. The first dataset (IITD I) is available in raw and pre-processed form, while the second (IITD II) is available only in the pre-processed form. With pre-processing the authors of the datasets ensured that: *i)* all ears are tightly cropped, *ii)* all images are of equal dimensions, and *iii)* all ears are centered and mutually aligned. Additionally, all images of the left ear are mirrored, so the entire dataset appears to consist of images of the right ear. The number of images per subject ranges from 3 to 6. No major occlusions are present. The dataset is available per request.

C. USTB Ear Datasets

The University of Science & Technology Beijing introduced four distinct ear datasets enjoying a high level of popularity within the community [90]. The first two datasets, USTB I and USTB II, contain only cropped ear images, while datasets III and IV feature complete head profile shots. In the first dataset the authors captured 3 to 4 images of 60 volunteers resulting in a total of 185 images. For the second dataset the authors acquired 4 images of 77 volunteers, altogether 308 images. In both datasets indoor lightning was used, but the second dataset contains more loosely cropped images, a higher level of illumination-induced variability and pitch angles between ±30° making it more challenging. The USTB III and USTB IV datasets contain full face-profile images captured under specific angles and occlusions. USTB III contains 1600 images of 79 subjects and USTB IV contains images

Table II: Comparative summary of the existing ear datasets and their characteristics. The category "Sides" evaluates whether images from the left or right side of the head are present in the dataset, "Yaw" and "Pitch" provide information about the ear rotation, "Occlusions" and "Accessories" assess whether occlusion and accessories are visible in the images and the last two categories "Gender" and "Ethnicities" indicate whether both sexes are present and what kind of racial variation is accounted for in the datasets.

DB	Year	#Subjects	#Images	Sides	Pitch	Yaw	Occlusions	Accessories	Gender	Ethnicities
CP [89]	1995	17	102	left	none	none	none	none	both	white
USTB I [90]	2002	60	185	right	mild	none	none	yes	both	asian
USTB II [90]	2004	77	308	right	mild	mild	none	yes	both	asian
IITD I [9]	2007	125	493	right	severe	none	none	yes	both	asian
AMI [91]	NA	100	700	both	severe	mild	mild	none	both	white
WPUT [8]	2010	501	2071	both	severe	mild	severe	yes	both	white
IITD II [9]	NA	221	793	right	none	none	none	none	both	asian
AWE	2016	100	1000	both	severe	severe	severe	yes	both	various

of 500 subjects with everyone photographed at 17 different angles with 15° steps in between. The USTB datasets are available per request and are free for research purposes.

D. AMI Ear Dataset

The AMI (translated from Spanish as Mathematical Analysis of Images) ear dataset [91] was collected at the University of Las Palmas and consists of 700 images of a total of 100 distinct subjects in the age group of 19–65 years. All images were taken under the same illumination and with a fixed camera position. For each subject 1 image of the left ear and 6 images of the right ear were taken. All ears in the dataset are loosely cropped, so that the ear area covers approximately one half of the image. The poses of subjects vary slightly in yaw (all images are still from profile nonetheless) and severely in pitch (subject looking up at 45° angle) angles. The dataset is publicly available.

E. WPUT Ear Dataset

The ear dataset of the West Pommeranian University of Technology (WPUT) [8] was introduced in 2010 with 2071 ear images belonging to 501 subjects. However, the dataset currently available for download contains 3348 images of 471 subjects, with 1388 duplicates. Images belonging to subject IDs from 337 to 363 are missing. The dataset contains annotation according to the following categories: gender, age group, skin color, head side, two types of rotation angles, lightning conditions, background (cluttered or heterogeneous), and occlusion type (earrings, hat, tattoos etc.). There are between 4 and 10 images per subject in the dataset. Images were taken under different indoor lightning conditions and (head-)rotation angles ranging from approximately 90° (profile) to 75°. The dataset is publicly available free of charge.

F. Other Datasets

In addition to the datasets reviewed above, datasets of ear-prints and profile-face images are also suitable for studying ear recognition technology. While there are many such datasets publicly available, we discuss (in this section) only datasets that were used for ear recognition in the literature.

The FEARID dataset [32] was collected as part of the FEARID project and differs from other datasets in the type of data it contains. Unlike other datasets, FEARID does not contain ear images, but ear-prints, which were acquired from different research groups using specific scanning hardware. Ear-print data differs significantly from regular images, as there are no occlusions, no variable angles and no illumination variations. However, other sources of variability influence the appearance of the ear-prints, such as the force with which the ear was pressed against the scanner, scanning-surface cleanliness and other similar factors. The acquisition protocol for the dataset was designed to simulate the appearance of ear-prints that would normally be found at crime scenes and resulted in a dataset comprising 7364 images of 1229 subjects. The FEARID dataset is interesting because it was collected with forensic applications in mind and with a somewhat different purpose than other ear datasets popular within the biometric community.

The first dataset from the group of profile-face-image datasets is the dataset of the Indian Institute of Technology Kanpur (IITK) [92]. During data collection, several face images with multiple fixed head poses were captured and organized into three subsets: *i)* the first subset contains 801 profile-only images of 190 subjects, *ii)* the second subset contains 801 images of 89 subjects with pitch angle variations, and *iii)* the third subset contains 1070 images of the same 89 subjects, but with yaw angle variations.

The UBEAR dataset [93] comprises 4429 profile images from 126 subjects taken from both the left and the right side. The images were captured under different illumination conditions. Partial occlusions of the ear are also present on some of the images. This dataset is interesting because images were captured while the subjects were moving. This characteristic is useful for studying techniques for video-based ear recognition where blurring and shearing effects typically appear.

The datasets of the University of Notre Dame (UND)[2] contain multiple subsets with 2D and 3D ear data. Subset E contains 2D images of the left side of the face of 114 subjects. Subsets F, G, J2 and NDOff-2007 all contain images

[2]https://sites.google.com/a/nd.edu/public-cvrl/data-sets

with clearly visible ears and corresponding 3D and 2D data with 942, 738, 1800 and 7398 images, respectively. These subsets are useful not only for 2D or 3D ear detection and recognition, but also enable direct comparison of 2D and 3D approaches.

The FERET [94] dataset in its final version contains 14, 051 images. The dataset is interesting because images contain annotations about the subjects' facial expression (neutral or not) and pose angles, which can be interesting covariates to study in ear recognition.

The Pose, Illumination and Expression (PIE) [95] dataset from the Carnegie Mellon University contains images of 68 subjects. The dataset comprises images captured in 13 different poses, 43 illumination conditions and 4 different facial expressions for each of the 68 subjects resulting in a total of 40.000 images for the entire dataset.

The XM2VTS [96] dataset features frontal and profile face images as well as video recordings of 295 subjects. The dataset was captured in controlled conditions during eight recording sessions and over a period of around five months.

Another dataset with video recordings is the ear dataset of the West Virginia University (WVU). This dataset is interesting because the authors used advanced capturing procedures (described in [97]) that allowed them to precisely record ears from different angles. The dataset contains images/videos of 137 subjects.

IV. AWE Dataset and Toolbox

Most datasets reviewed in the previous section were captured in controlled, laboratory conditions. In this section we introduce the Annotated Web Ears (AWE) dataset, which differs significantly in this regard from other existing datasets. We also present the AWE Matlab toolbox for ear recognition. The dataset and toolbox are both publicly available from http://awe.fri.uni-lj.si.

A. The AWE Dataset

Motivation: Most existing datasets for evaluating ear recognition technology contain images captured in controlled laboratory settings or profile shots of people posing outdoors in front of a camera. The variability of these images is usually limited and the sources of variability are pre-selected by the authors of the datasets. This is in stark contrast to the field of face recognition, where in the last decade the research focus shifted to datasets of completely uncontrolled images, which helped to advance the state-of-technology significantly. A cornerstone in this development were datasets gathered directly from the web, such as the LFW [98], PubFig [99], FaceScrub [100], Casia WebFaces [101], IJB-A [102] and others. These datasets introduced the notion of *images captured in the wild*, with the wordplay indicating that the images were acquired in uncontrolled environments. The difficulty of these datasets provided significant room for improvements in the technology and spurred the development of the field over recent years.
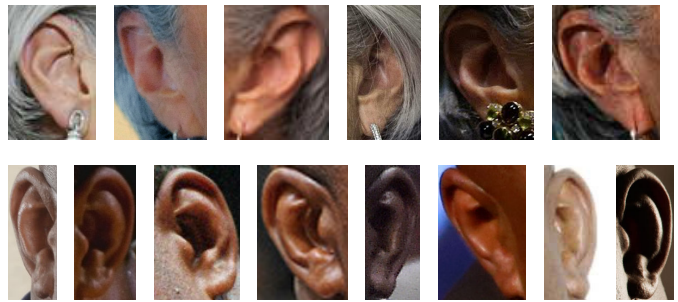

Figure 5: Sample images from the Annotated Web Ears (AWE) dataset. Images in each row correspond to one subject from the dataset. Note the extend of variability present in the images.

By today, the reported performance of ear-recognition techniques has surpassed the rank-1 recognition rate of 90% on most available datasets. This fact suggests that the technology has reached a level of maturity that easily handles images captured in laboratory-like settings and more challenging datasets are needed to identify open problems and provide room for further advancements. Inspired by the success of datasets collected *in the wild* for face recognition, we gather a new dataset of ear images from the web and make it available to the research community.

Dataset collection: All images of the AWE dataset were collected from the web by first compiling a list of people, for whom it was reasonably to assume that a large amount of images could be found online. Similar to other datasets collected in the wild this list mostly featured actors, musicians, politicians and the like. A web crawler was then used to gather large amounts of images for each person from the list by sending appropriate search queries to Google's image search[3]. To avoid pre-filtering of the images with automatic detection techniques [102] and, thus, limit the image variability, all images were manually inspected and 10 images per subject were selected for inclusion in the dataset.

Structure and annotations: In total, the dataset contains 1000 ear images of 100 subjects. Each subject has 10 images of a different quality and size. The smallest image in the dataset has 15×29 pixels, the biggest 473×1022 pixels and the average image size is 83×160 pixels. The ear images are tightly cropped and do not contain profile faces as shown in Fig. 5.

Each image in the dataset was annotated by a trained researcher according to the following categories: gender, ethnicity, accessories, occlusion, head pitch, head roll, head yaw, and head side. The pitch, roll and yaw angles were estimated from the head pose. The available labels for each category are shown in Table III and the distribution of the labels for the entire dataset in Fig. 6. Additionally, the location of the tragus was also marked in each image.

Experimental protocols: We propose two types of ex-

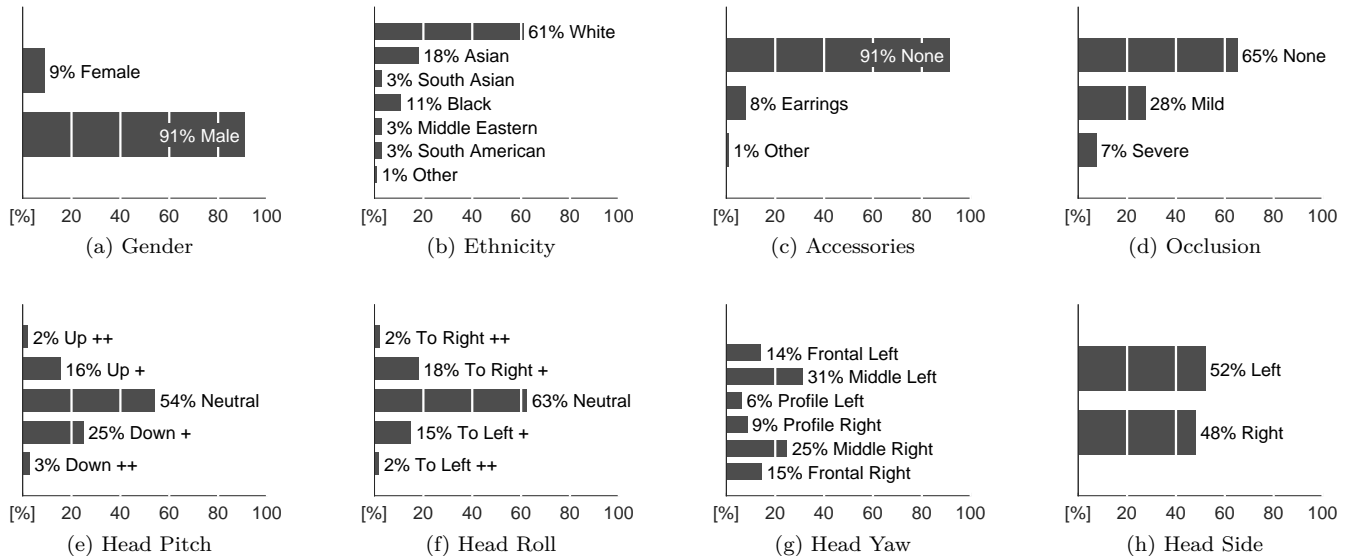[3]In most cases the name and surname generated enough hits.

Figure 6: Distribution of image labels in the Annotated Web Ears (AWE) dataset.

Table III: Annotation categories and corresponding labels for Annotated Web Ears (AWE) dataset.

Category	Available labels
Gender	Male, Female
Ethnicity	White, Asian, South Asian, Black Middle Eastern, South American, Other
Accessory	None, Earrings, Other
Occlusion*	None, Mild, Severe
Head Pitch**	Up ++, Up +, Neutral, Down +, Down ++
Head Roll**	To Right ++, To Right +, Neutral To Left +, To Left ++
Head Yaw	Frontal Left, Middle Left, Profile Left Profile Right, Middle Right, Frontal Right
Head Side	Left, Right

* *mild* denotes occlusions, where small parts of the ear are covered by earrings or hair strands, *severe* denotes occlusions where larger parts of the ear are not visible and covered by hair or other objects
** ++ means larger than 10°, + means between 2° and 10°

perimental protocols for the dataset to enable comparisons of techniques evaluated on our dataset. The first corresponds to protocols for identification (i.e., one-vs-many) experiments and the second to protocols for verification (i.e., one-vs-one) experiments.

For both types of protocols we first partition the data into a *development set* that contains 60% of all images and a *test set* that contains the remaining 40% of images. We propose to use a 5-fold cross validation procedure on the development set to train potential background models, subspaces or classifiers and then apply the trained models on the (hold-out) test set for evaluation. Results should be reported on the development set by presenting all performance metrics in the form of means and standard deviations computed over all 5-folds. Similarly, results on the test set should again be presented in the form of mean values with corresponding standard deviations of all

performance metrics, but this time the statistics should be generated through bootstrapping. Files containing lists for the experiments are distributed with our dataset.

The following performance curves and metrics should be generated when reporting performance [1]:

- *Identification Experiments*:
  - Cumulative Match-score Curves (CMCs),
  - Rank-1 recognition (R1) rates,
- *Verification Experiments*:
  - Receiver Operating Characteristics (ROC) curves,
  - Equal Error Rates (EERs).

B. The AWE Toolbox

Motivation: Reproducing experimental results is one of the most important steps in many research fields related to computer vision and machine learning. We typically observe a gap between methods described in the literature and available open source implementations [103], which makes implementing prior work a standard task accompanying every research project. This process is not only time consuming but also affects the reproducibility of published results. To address this problem for the field of ear recognition, we introduce in this section the AWE toolbox, an open source Matlab toolbox designed for reproducible research in ear recognition (and beyond).

Software requirements: The AWE toolbox is written entirely in Matlab and has been fully tested on the 2015a and 2015b editions, but it should run with older versions as well. To access all of the functionality of the toolbox, Matlab's Computer Vision System Toolbox needs to be installed, as some of the feature extraction techniques supported by the AWE toolbox rely on Matlab's built-in implementations. The toolbox itself exploits a few external libraries but these ship with the toolbox so there is no need for manual installation.

Structure and Capabilities: The toolbox consists of four major parts devoted to: *i)* data handling, *ii)* feature extraction, *iii)* distance calculation or/and classification, and *iv)* result visualization.

The first part of the toolbox (devoted to data handling) takes an arbitrary dataset as an input. The dataset needs to be organized into a directory structure with each folder containing samples of only one class. In this regard, the toolbox is not limited solely to ear images, but can also be used with other biometric modalities. Using functions from the first toolbox part, images are read, normalized and cached for faster reuse at a later time. Currently supported normalization techniques include more than 20 state-of-the-art photometric normalization techniques, which are provided by the INFace toolbox[4] that is integrated with our toolbox.

The second part of our toolbox is oriented towards feature extraction. For the intended work-flow of the toolbox preselected feature extraction techniques are applied to all normalized images of the selected dataset. The computed feature vectors are (by request) stored in comma-separated-values (CSV) files so that users can either use them in subsequent experiments or export them to other processing or visualization tools. The toolbox currently implements 8 descriptor-based feature extraction techniques, but users can easily plug in their own methods or add functionality to the existing code. The toolbox currently supports extraction of LBP, LPQ, RILPQ, BSIF, DSIFT, POEM, Gabor and HOG features.

The third part of the toolbox, intended for matching and classification of the computed feature vectors, supports two options: *i)* feature vectors can either be matched directly using various distance measures, or *ii)* classifiers may be trained on the training data of the datasets to classify samples. While the toolbox currently supports several distance measures for the first option, we have not included any specific classifiers in the toolbox yet as many ship with every Matlab installation.

The last part of the toolbox generates all relevant performance curves and metrics, such as ROC plots, CMC plots, histograms of client and impostor score distributions, rank-1 recognition rates, equal error rates, and similar metrics useful for comparisons.

The entire processing pipeline implemented in the four parts of the toolbox may be easily configured in any combination and automatically executed for any dataset provided to the toolbox.

One of the biggest advantages of the AWE toolbox is speeding up of the time needed to develop new methods for ear recognition. Several state-of-the-art techniques are already implemented, which makes comparisons straightforward and the results reproducible. The toolbox is open source so it can be modified, shared and used freely without limitations.

V. EXPERIMENTS AND RESULTS

To highlight the current capabilities of the AWE toolbox and provide an independent comparative analysis of several state-of-the-art descriptor-based ear recognition techniques we conduct in this section recognition experiments with two popular datasets as well as our newly introduced AWE dataset. All presented experiments are fully reproducible using the AWE toolbox and rely on the same experimental protocol.

A. Comparative Assessment on Existing Datasets

Many existing ear datasets do not define an experimental protocol that can be used during experimentation. Consequently, authors set up their own protocols, which makes comparisons of published results difficult. To establish an objective ranking of the state-of-the-art (descriptor-based) methods for ear recognition, we conduct verification as well as identification experiments with 8 different techniques from the AWE toolbox on two popular ear datasets. We follow a similar processing pipeline for all techniques, which: *i)* rescales the images to a fixed size of 100×100 pixels, *ii)* corrects for illumination-induced variability by applying histogram equalization to the resized images, *iii)* subjects the images to the selected feature (or descriptor) extraction techniques, and *iv)* produces a similarity score for each probe-to-gallery comparison by computing the distance between the corresponding feature vectors.

A brief summary of the feature extraction techniques and selected hyper-parameters used in the experiments is given in Table IV. We implement feature extraction techniques based on LBP, LPQ, BSIF, POEM, HOG, DSIFT, RILPQ and Gabor features and select the open hyper-parameters through cross-validation. We use the chi-square distance to measure the similarity between the feature vectors of the probe and gallery images for all histogram-based descriptors and the cosine similarity measure for the Gabor features.

We perform the comparative assessment on images from the IITD II and USTB II datasets[5]. These datasets are among the most popular datasets for ear recognition and have, therefore, also been selected for our experiments. The reader is referred to Section III for a more in-depth description of these datasets.

Two types of experiments are conducted. The first uses the cropped and pre-aligned images that ship with the IITD II dataset, while the second employs cropped but non-aligned images from the USTB II dataset. As already suggested above, the goal of these experiments is to study the behavior of the state-of-the-art descriptor-based methods on images of varying quality and to demonstrate the current capabilities of our AWE toolbox.

We conduct 5-fold cross validation experiments on the two datasets and generate average ROC and CMC plots

---

[4]Available from: http://luks.fe.uni-lj.si/sl/osebje/vitomir/face_tools/INFace/

[5]We also performed experiments on the IITD I and USTB I datasets with comparable results but, for the sake of brevity, do not report the results here.

Table IV: Summary of implemented techniques used for the comparison. The symbol "pix" stands for "pixels".

Method	Detailed information
LBP [41], [56], [77], [104]	uniform LBP, radius: 2 pix, neighborhood size: 8, block size: 8×8 pix, no block overlap
LPQ [105]	window size: 5×5 pix, block size: 18×18 pix, no block overlap
BSIF [41], [106]	#filters: 8, filter size: 11×11 block size: 18×18 pix, no block overlap
POEM [88]	#orientations: 3, uniform LBP, radius: 2 pix, neighborhood size: 8, block size: 12×12 pix, no block overlap
HOG [41], [87], [107]	cell: 8×8 pix, block: 16×16 pix, block overlap: 8 pix
DSIFT [48], [53], [65], [85], [108], [109]	#grid points: 100, patch size: 16×16 pix, bin size: 8
RILPQ [86]	window radius: 6 pix, #angles: 12, block size: 16×16 pix, no block overlap
Gabor [37], [9], [62], [63], [80], [110], [111], [112], [113]	#orientations: 8, #scales: 5, #filters 40, down-sampling factor: 64

(a) IITD II CMCs (b) USTB II CMCs

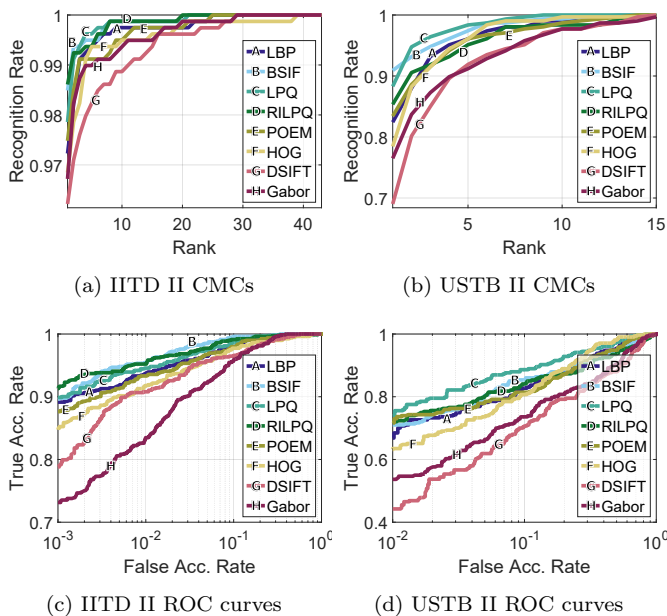(c) IITD II ROC curves (d) USTB II ROC curves

Figure 7: Comparative assessment of 8 state-of-the-art techniques. The upper row shows CMC curves and the lower row ROC plots for the two datasets. Experiments on the IITD II dataset were conducted with pre-aligned images, while the results on the USTB II datasets were generated with non-aligned images. Note that the performance of all assessed techniques is better on the aligned than on the nonaligned dataset and not much room is left for further improvements.

over all 5 folds as shown in Fig. 7. We adjust the scale on the *y*-axis for each dataset to better highlight the differences in the performance of the assessed techniques. To avoid clutter in the graphs, we do not provide error bars

for the performance curves but instead include standard deviations for the performance metrics in Table V. Here, we report the mean rank-1 recognition rate (R1) for the identification experiments and the equal error rate (EER) for the verification experiments..

Table V: Comparative assessment of 8 state-of-the-art techniques on the pre-aligned images from the IITD II dataset and the non-aligned images from the USTB II dataset. For each performance metric the results are provided in the form of the mean and the standard deviation over 5-folds. The performance of the assessed techniques on both datasets is comparable to the existing state-of-the-art from the literature.

Method	IITD II		USTB II	
	R1	EER	R1	EER
LBP	97.2 ± 0.9	3.9 ± 1.7	82.4 ± 7.5	15.0 ± 7.5
BSIF	98.5 ± 1.0	2.4 ± 1.4	90.9 ± 6.5	13.4 ± 7.2
LPQ	97.9 ± 0.5	3.3 ± 1.5	88.3 ± 7.6	11.1 ± 5.6
RILPQ	98.6 ± 0.9	3.1 ± 1.3	85.4 ± 8.7	13.7 ± 8.0
POEM	97.5 ± 1.1	3.9 ± 1.4	83.4 ± 7.2	15.3 ± 8.3
HOG	96.9 ± 0.7	4.5 ± 1.7	78.5 ± 7.0	14.6 ± 7.2
DSIFT	96.2 ± 1.9	4.5 ± 2.0	69.0 ± 9.1	20.5 ± 8.5
Gabor	96.7 ± 1.6	6.6 ± 1.6	76.5 ± 4.4	19.3 ± 7.6

The results show that the performance of all assessed techniques is significantly better for aligned than for non-aligned images. This is expected as pose variability is generally considered one of the major factors influencing ear recognition performance.

On the aligned images all evaluated techniques perform similarly in the identification experiments (see Fig. 7 – (a)). The difference in the recognition performance is in fact minimal for techniques that rely on any of the binary pattern variants to encode the ear texture. For the verification experiments (Fig. 7 – (c)) the differences among the techniques are somehow larger. Here, HOG,

DSIFT and Gabor features result in lower performance, while the binary-pattern-based techniques again perform better and more or less on the same level.

On the non-aligned images the ranking of the techniques does not change significantly compared to the aligned images. Techniques based on binary patterns perform very similarly, while HOG, DSIFT and Gabor-based techniques again result in worse performance (Fig. 7 – (b), (d)). The main conclusion that we can draw from this series of experiments is that techniques based on binary patterns should be preferred over other local methods, while the choice of a specific method within this group should be made based on other non-performance related criteria.

The presented results are comparable to what was reported by other authors for similar techniques on the same datasets, e.g., [39], [41], [56], [78], [80], which suggests that the AWE toolbox provides competitive implementations of all evaluated descriptor-based techniques.
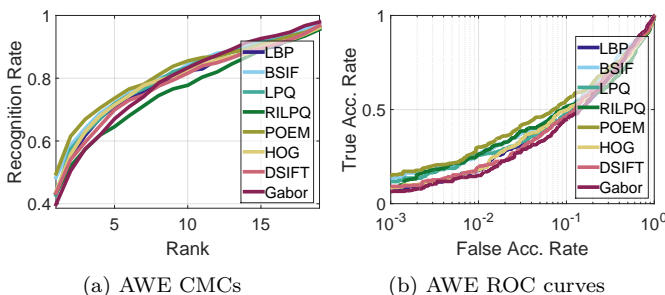

(a) AWE CMCs   (b) AWE ROC curves

Figure 8: Performance curves generated during 5-fold cross-validation experiments on the development part of the Annotated Web Ears (AWE) dataset. The left graph depicts CMC and the right graph ROC curves of the experiments. The results show the usefulness and the difficulty of the AWE dataset and indicate that there is significant room for improvement of ear recognition technology. Here, the performance curves are not marked as in Fig. 7 due to the almost equal performance of all descriptors.

B. Comparative Assessment on the AWE Dataset

We repeat the experiments from the previous section with the same 8 feature extraction techniques, i.e., based on LBP, LPQ, RILPQ, BSIF, DSIFT, POEM, Gabor, and HOG features, on the AWE dataset and follow the experimental protocols outlined in Section IV. We use the same hyper-parameters as in the previous section for all techniques. The goal of these experiments is to demonstrate the difficulty of recognizing ear images captured in completely uncontrolled conditions and to establish the performance of the state-of-the-art descriptor-based methods on images captured in the wild. The images of the AWE dataset are not aligned; the results presented in this section should, therefore, be compared with the results reported on the USTB II dataset when assessing the difficulty of the new dataset.

Following the predefined experimental protocols, we split the data into two parts and conduct experiments on the development part of the dataset using a 5-fold cross-validation procedure. As suggested in Section IV, files containing lists of the required experiments for each fold are distributed with the dataset. We report the results in the form of CMC and ROC curves and with the predefined performance metrics. When reporting quantitative results, we provide mean and standard deviations for each required performance metric computed over all 5-folds. We conduct a total of 600 identification experiments and $1,856$ legitimate vs. $34,308$ illegitimate authentication trials for the verification experiments.

Note that the dataset partitioning into separate development and test sets is meant to facilitate training of classification techniques (such as [114]). In this case the development set is intended for model learning and parameter tuning, while the test set is reserved for the final evaluation. However, because we do not use classification models in our experiments, we do not report the results on the test set in this paper, but only state that they are similar to the results obtained on the development set.

For the development set the results are presented in Fig. 8 and Table VI. As can be seen, the performance of all assessed techniques on the AWE dataset is significantly lower than on the datasets used in the previous section for both the identification and verification experiments. The mean rank-1 recognition rates range from 39.8% achieved with Gabor features to 49.6% achieved by the POEM descriptor. The remaining recognition rates are close and all above 40%. For the verification experiments the performance of all evaluated methods is also close but, considering the entire ROC curve, the POEM descriptor is again the top performer. Here, the best performance with respect to EER is around 30%. These results highlight the difficulty of the AWE dataset and clearly show that more research is needed to make the ear-recognition technology applicable to completely uncontrolled images.

Table VI: Comparative assessment of 8 state-of-the-art techniques on the development part of the Annotated Web Ears (AWE) dataset. For each performance metric the results are provided in the form of the mean and the standard deviation over 5-folds generated during cross-validation experiments. The results demonstrate the usefulness and the difficulty of the dataset.

Method	R1	EER
LBP	43.5 ± 7.1	32.0 ± 7.4
BSIF	48.4 ± 6.8	30.0 ± 9.6
LPQ	42.8 ± 7.1	30.0 ± 7.4
RILPQ	43.3 ± 9.4	34.0 ± 6.4
POEM	49.6 ± 6.8	29.1 ± 9.1
HOG	43.9 ± 7.9	31.9 ± 7.8
DSIFT	43.4 ± 8.6	30.9 ± 8.4
Gabor	39.8 ± 7.1	32.1 ± 8.1

VI. Open Questions and Research Directions

Ear recognition still represents a relatively young research area compared to other biometric modalities. Problems and research questions that are well studied in other fields require more research and provide room for exploration [115], [116], [117]. In this section we briefly describe what we feel are the most important open problems and most promising research directions in the field. Our goal here is to establish a list of topics worth investigating in the years to come.

A. Ear Detection and Alignment

Fully automatic ear recognition systems require reliable ear detection techniques capable of determining the location of the ear in the input images (or video) in an efficient and possibly real-time manner. While there has been some research effort in this direction (see, for example, [15], [16] for recent surveys related to detection), ear detection is still an unsolved problem and no widely adopted solution has been proposed yet in the literature. In our opinion, two major issues need to be addressed to foster advancements in this area. The first is the introduction of large-scale ear datasets captured outside laboratory environments that would allow training competitive ear detectors to be capable of operating on real-world imagery and video footage. As a shortcut towards addressing this issue, the existing face datasets with suitable profile images could be used for this purpose. The second issue is the implementation of publicly available open-source tools for ear detection that would provide baselines for experimentation and development. To the best of our knowledge no competitive open-source implementation of ear detection techniques can currently be found on-line.

Another major problem in the automatic ear recognition systems is alignment. Pose variations are generally considered to be one of the biggest issues plaguing the technology [15], [16]. In-plane rotations as well as moderate out of plane variations can effectively be addressed with alignment techniques, which makes alignment a key factor of robustness for ear recognition. Manual ear alignment is common, but existing automatic alignment techniques include active shape models [52] or combinations of active contours models and ovoids over ear contours [118]. However, recent advancements in boosted/cascaded landmark detection and alignment, e.g., [119], [120], can also be used for ear images provided that appropriately annotated data for training becomes available. Some initial attempts in this direction have been presented in [41].

B. Handling Ear Variability

Appearance variability encountered with ear images is induced by different factors, such as occlusions by hair, clothing or accessories, self-occlusions due to view-point changes, or pose variations (among others). These factors represent the main sources of errors in ear recognition systems and therefore need to be addressed appropriately.

Various techniques have been proposed in the literature capable of handling moderate occlusions (by external objects or due to view point changes), e.g., [53], [65], though it is not clear where the breaking points of these techniques are or how much occlusion can be tolerated by the existing technology. Ear occlusion does not only affect recognition schemes but likely has an equally large effect on ear detection and alignment techniques. Errors from these steps are propagated through the processing pipelines and degrade the recognition performance of automatic ear recognition systems. Potential solutions here could exploit contextual information and address occlusions by modeling the entire facial profile instead of the ear region alone. A similar idea was, for example, successfully exploited in [121].

Pose variations can to a large extent be addressed with suitable alignment techniques, as already emphasized in the previous section. 3D data and generic 3D models that could be exploited to normalize the geometry of the ear may play a bigger role in the future. Especially appealing in this context are techniques based on the recently popularized concept of analysis-by-synthesis [122] where images are not aligned into predefined canonical forms. Instead the poses of probe and gallery images are matched through a reposing procedure that renders one of the images in the pose of the second. This pair-wise approach may represent a viable research direction for pose-invariant ear recognition in the future.

C. Feature Learning

Ear recognition approaches mainly followed the development of other computer vision fields and progressed from early geometric techniques and holistic methods to more recent local and hybrid approaches. The next step in this development is the use of feature learning techniques based on convolutional neural networks (CNNs) for ear image description or even end-to-end system design. While techniques based on neural networks are becoming increasingly popular in other areas of computer vision and machine learning, ear recognition has not yet benefited from recent advances in this area mainly due to the lack of large-scale datasets.

D. Large-scale Deployment

To date, all available datasets for evaluating ear recognition technology are still relatively modest in size, featuring images of at most a few hundred subjects. It is still not clear how the existing technology scales with the number of subjects, especially for large-scale identification tasks. Experimental studies on this topic are needed as well as mathematical models of ear individuality that would lay the theoretical basis for the distinctiveness of ears [15]. Some research on this topic can be found in the forensic literature but more work is required to provide empirical evidence for the suitability of ear images for large-scale identification tasks.

E. Imaging Modalities

Most research work related to ear recognition focuses on 2D images captured with common off-the-shelf (COTS) cameras or CCTV systems. However, other modalities, such as 3D ear data or ear-prints, are equally important and offer similar research challenges as 2D images. Detection, segmentation, feature extraction and modeling techniques need to be developed for these modalities and we believe that research focused on these problems will become more important in the future.

Another interesting problem that we believe will garner more interest in the long term is heterogeneous ear recognition where ear images captured with different imaging modalities will be matched against each other.

F. Understanding Ear Recognition

As pointed out by [15], [16], there are several factors pertaining to ear recognition that are not yet well understood and offer room for further exploration. These factors include: *i) Ear symmetry:* some studies have found that the left and right ear are close to symmetric for many subjects but more research is needed to find ways of exploiting this fact in automatic recognition systems; *ii) Aging:* it is known that the size of the ear increases with age, e.g., [20], [22], [23], but how this effect influences the recognition performance of automatic techniques is still unclear; the main problem here is the lack of suitable datasets that would be needed for a detailed analysis of the effects of template aging; *iii) Inheritance:* as pointed out by [16], the work of Iannarelli [19] suggested that some characteristics of the outer ear can be inherited; whether these characteristics could be identified by automatic techniques and exploited to determine if two persons are related or whether they could present a source of error for automatic recognition systems is still an open question.

VII. Conclusion

In this paper we surveyed the field of ear recognition including the most popular datasets used in the ear recognition research, and introduced a new, publicly available ear dataset gathered from the web. The dataset is the first of its kind and represents a considerable challenge for the existing technology. We demonstrated the difficulty of the dataset in recognition experiments with a toolbox of the state-of-the-art ear recognition techniques that was also introduced in this paper. Finally, we identified several open problems in the field and provided an outlook into the future of ear recognition.

We hope that the survey, toolbox and dataset will help the research community to improve the state-of-technology in ear recognition and inspire new research in this area. The toolbox is open to the wider biometric community and useful beyond ear recognition and we plan to include more techniques in the future, such as: ear detection and alignment approaches, normalization techniques, new descriptors and more complex procedures for feature learning.

References

[1] A. Jain, A. Ross, K. Nandakumar, Introduction to biometrics, Springer Science & Business Media, 2011.

[2] H. Nejati, L. Zhang, T. Sim, E. Martinez-Marroquin, G. Dong, Wonder Ears: Identification of Identical Twins from Ear Images, in: Proceedings of the International Conference on Pattern Recognition, IEEE, 2012, pp. 1201–1204.

[3] X. Xu, Z. Mu, L. Yuan, Feature-Level fusion method based on kfda for multimodal recognition fusing ear and profile face, in: Proceedings of the International Conference on Wavelet Analysis and Pattern Recognition, IEEE, 2007, pp. 1306–1310.

[4] X. Pan, Y. Cao, X. Xu, Y. Lu, Y. Zhao, Ear and face based multimodal recognition based on KFDA, in: Proceedings of the International Conference on Audio, Language and Image Processing, IEEE, 2008, pp. 965–969.

[5] T. Theoharis, G. Passalis, G. Toderici, I. A. Kakadiaris, Unified 3D face and ear recognition using wavelets on geometry images, Pattern Recognition 41 (3) (2008) 796–804.

[6] S. Islam, M. Bennamoun, A. Mian, R. Davies, Score level fusion of ear and face local 3D features for fast and expression-invariant human recognition, in: Proceedings of the International Conference on Image Analysis and Recognition, 2009, p. 387–396.

[7] D. R. Kisku, P. Gupta, H. Mehrotra, J. K. Sing, et al., Multimodal belief fusion for face and ear biometrics, Intelligent Information Management 1 (03) (2009) 166.

[8] D. Frejlichowski, N. Tyszkiewicz, Image Analysis and Recognition: 7th International Conference, ICIAR 2010, Póvoa de Varzin, Portugal, June 21–23, 2010, Proceedings, Part II, Springer Berlin Heidelberg, 2010, Ch. The West Pomeranian University of Technology Ear Database – A Tool for Testing Biometric Algorithms, pp. 227–234.

[9] A. Kumar, C. Wu, Automated human identification using ear imaging, Pattern Recognition 45 (3) (2012) 956–968.

[10] Homepage of Descartes Biometrics, http://www.descartesbiometrics.com/, accessed: 2016-02-26.

[11] R. Purkait, Role of External Ear in Establishing Personal Identity - A Short Review, Austin Journal of Forensic Science and Criminology 2 (2) (2015) 1–5.

[12] S. Islam, M. Bennamoun, R. Owens, R. Davies, Biometric approaches of 2D - 3D ear and face: a survey, in: T. Sobh (Ed.), Advances in computer and information sciences and engineering, Springer, 2008, pp. 509–514.

[13] M. Choraś, Image feature extraction methods for ear biometrics–a survey, in: Proceedings of the International Conference on Computer Information Systems and Industrial Management Applications, IEEE, 2007, pp. 261–265.

[14] K. Pun, Y. Moon, Recent advances in ear biometrics, in: Proceedings of the International Conference on Automatic Face and Gesture Recognition, IEEE, 2004, pp. 164–169.

[15] A. Abaza, A. Ross, C. Hebert, M. A. F. Harrison, M. Nixon, A Survey on Ear Biometrics, ACM Computing Surveys 45 (2) (2013) 1–22.

[16] A. Pflug, C. Busch, Ear biometrics: a survey of detection, feature extraction and recognition methods, Biometrics 1 (2) (2012) 114–129.

[17] P. Yan, K. W. Bowyer, Empirical Evaluation of Advanced Ear Biometrics, in: Proceedings of the Conference on Computer Vision and Pattern Recognition, IEEE, 2005, p. 41.

[18] A. Abaza, A. Ross, Towards Understanding the Symmetry of Human Ears: A Biometric Perspective, in: Proceedings of the Conference on Biometrics: Theory, Applications and Systems, IEEE, 2010, pp. 387—-396.

[19] A. Iannarelli, Ear Identification, Forensic Identification Series, Paramont Publishing Company, 1989.

[20] C. Sforza, Grandi, M. G., Binelli, D. Tommasi, R. Rosati, V. Ferrario, Age- and sex-related changes in the normal human ear, Forensic Anthropology Population Data 187 (1–3) (2009) 110.e1–110.e7.

[21] R. Purkait, P. Singh, Anthropometry of the normal human auricle: a study of adult Indian men, Aesthetic Plastic Surgery (31) (2007) 372–379.

[22] R. Gualdi-Russo, Longitudinal study of anthropometric changes with aging in an urban Italian population, Homo - Journal of Comparative Human Biology 49 (3) (1998) 241–259.

[23] V. Ferrario, C. Sforza, V. Ciusa, G. Serrao, G. Tartaglia, Morphometry of the normal human ear: a cross-sectional study from adolescence to midadulthood, Journal of Craniofacial Genetics and Development Biology 19 (4) (1999) 226–233.

[24] M. Ibrahim, M. Nixon, S. Mahmoodi, The effect of time on ear biometrics, in: Proceedings of the International Joint Conference on Biometrics, IEEE, 2011, pp. 1–6.

[25] A. Bertillon, Signaletic instructions including the theory and practice of anthropometrical identification, Werner Co., 1896.

[26] R. Imhofer, Die Bedeutung der Ohrmuschel für die Feststellung der Identität, Archiv für die Kriminologie 26 (150–163) (1906) 3.

[27] E. Locard, L'identification des récidivistes, A. Maloine, 1909.

[28] C. Fields, H. Falls, C. Warren, M. Zimberoff, The ear of the newborn as an identification constant, Obstet Gynecol. 16 (1960) 98–102.

[29] M. Burge, W. Burger, Biometrics: Personal Identification in Networked Society, Springer US, Boston, MA, 1996, Ch. Ear Biometrics, pp. 273–285.

[30] B. Moreno, A. Sánchez, J. F. Vélez, On the use of outer ear images for personal identification in security applications, in: Proceedings of the International Carnahan Conference on Security Technology, IEEE, 1999, pp. 469–476.

[31] D. J. Hurley, M. S. Nixon, J. N. Carter, Automatic ear recognition by force field transformations, in: Proceedings of the Colloquium on Visual Biometrics, IET, 2000, pp. 7–1.

[32] I. Alberink, A. Ruifrok, Performance of the FearID earprint identification system, Forensic science international 166 (2) (2007) 145–154.

[33] B. Victor, K. Bowyer, S. Sarkar, An evaluation of face and ear biometrics, in: Proceedings of the International Conference on Pattern Recognition, Vol. 1, IEEE, 2002, pp. 429–432.

[34] D. J. Hurley, M. S. Nixon, J. N. Carter, Ear biometrics by force field convergence, in: Proceedings of the Audio-and Video-Based Biometric Person Authentication, Springer, 2005, pp. 386–394.

[35] L. Yuan, Z.-c. Mu, Y. Zhang, K. Liu, Ear recognition using improved non-negative matrix factorization, in: Proceedings of the International Conference on Pattern Recognition, Vol. 4, IEEE, 2006, pp. 501–504.

[36] M. S. Nosrati, K. Faez, F. Faradji, Using 2D wavelet and principal component analysis for personal identification based on 2D ear structure, in: Proceedings of the International Conference on Intelligent and Advanced Systems, IEEE, 2007, pp. 616–620.

[37] A. Kumar, D. Zhang, Ear authentication using log-gabor wavelets, in: Proceedings of the Symposium onDefense and Security, International Society for Optics and Photonics, 2007, p. 65390A.

[38] W. Zhi-Qin, Y. Xiao-dong, Multi-scale feature extraction algorithm of ear image, in: Proceedings of the International Conference on Electric Information and Control Engineering, IEEE, 2011, pp. 528–531.

[39] A. Benzaoui, A. Hadid, A. Boukrouche, Ear biometric recognition using local texture descriptors, Journal of Electronic Imaging 23 (5) (2014) 053008.

[40] A. Pflug, C. Busch, A. Ross, 2D ear classification based on unsupervised clustering, in: Proceedings of the International Joint Conference on Biometrics, IEEE, 2014, pp. 1–8.

[41] A. Pflug, P. N. Paul, C. Busch, A comparative study on texture and surface descriptors for ear biometrics, in: Proceedings of the International Carnahan Conference on Security Technology, IEEE, 2014, pp. 1–6.

[42] K. Chang, K. W. Bowyer, S. Sarkar, B. Victor, Comparison and combination of ear and face images in appearance-based biometrics, Transactions on Pattern Analysis and Machine Intelligence 25 (9) (2003) 1160–1165.

[43] Z. Mu, L. Yuan, Z. Xu, D. Xi, S. Qi, Shape and structural feature based ear recognition, in: Advances in biometric person authentication, Springer, 2004, pp. 663–670.

[44] H.-J. Zhang, Z.-C. Mu, W. Qu, L.-M. Liu, C.-Y. Zhang, A novel approach for ear recognition based on ICA and RBF network, in: Proceedings of the International Conference on Machine Learning and Cybernetics, Vol. 7, IEEE, 2005, pp. 4511–4515.

[45] A. F. Abate, M. Nappi, D. Riccio, S. Ricciardi, Ear recognition by means of a rotation invariant descriptor, in: Proceedings of the International Conference on Pattern Recognition, Vol. 4, IEEE, 2006, pp. 437–440.

[46] M. Abdel-Mottaleb, J. Zhou, Human ear recognition from face profile images, in: Advances in biometrics, Springer, 2006, pp. 786–792.

[47] M. Choras, R. S. Choras, Geometrical algorithms of ear contour shape representation and feature extraction, in: Proceedings of the International Conference on Intelligent Systems Design and Applications, IEEE, 2006, pp. 451–456.

[48] K. Dewi, T. Yahagi, Ear photo recognition using scale invariant keypoints., in: Proceedings of the Computational Intelligence, 2006, pp. 253–258.

[49] B. Arbab-Zavar, M. S. Nixon, D. J. Hurley, On model-based analysis of ear biometrics, in: Proceedings of the Conference on Biometrics: Theory, Applications and Systems, IEEE, 2007, pp. 1–5.

[50] L. Nanni, A. Lumini, A multi-matcher for ear authentication, Pattern Recognition Letters 28 (16) (2007) 2219–2226.

[51] M. Rahman, M. R. Islam, N. I. Bhuiyan, B. Ahmed, A. Islam, Person identification using ear biometrics, International Journal of The Computer, the Internet and Management 15 (2) (2007) 1–8.

[52] L. Yuan, Z.-C. Mu, Ear recognition based on 2D images, in: Proceedings of the Conference on Biometrics: Theory, Applications and Systems, IEEE, 2007, pp. 1–5.

[53] B. Arbab-Zavar, M. S. Nixon, Robust log-gabor filter for ear biometrics, in: Proceedings of the International Conference on Pattern Recognition, IEEE, 2008, pp. 1–4.

[54] M. Choraś, Perspective methods of human identification: ear biometrics, Opto-electronics review 16 (1) (2008) 85–96.

[55] J. Dong, Z. Mu, Multi-pose ear recognition based on force field transformation, in: Proceedings of the International Symposium on Intelligent Information Technology Application, Vol. 3, IEEE, 2008, pp. 771–775.

[56] Y. Guo, Z. Xu, Ear recognition using a new local matching approach, in: Proceedings of the International Conference on Image Processing, IEEE, 2008, pp. 289–292.

[57] I. Naseem, R. Togneri, M. Bennamoun, Sparse representation for ear biometrics, in: Advances in visual computing, Springer, 2008, pp. 336–345.

[58] Y. Wang, Z.-C. Mu, H. Zeng, Block-based and multi-resolution methods for ear recognition using wavelet transform and uniform local binary patterns, in: Proceedings of the International Conference on Pattern Recognition, IEEE, 2008, pp. 1–4.

[59] Z. Xie, Z. Mu, Ear recognition using lle and idlle algorithm, in: Proceedings of the International Conference on Pattern Recognition, IEEE, 2008, pp. 1–4.

[60] Z. Zhang, H. Liu, Multi-View ear recognition based on b-spline pose manifold construction, in: Proceedings of the World Congress on Intelligent Control and Automation, 2008.

[61] Z. Hai-Long, M. Zhi-Chun, Combining wavelet transform and orthogonal centroid algorithm for ear recognition, in: Proceedings of the International Conference on Computer Science and Information Technology, IEEE, 2009, pp. 228–231.

[62] L. Nanni, A. Lumini, Fusion of color spaces for ear authentication, Pattern Recognition 42 (9) (2009) 1906–1913.

[63] W. Xiaoyun, Y. Weiqi, Human ear recognition based on block segmentation, in: Proceedings of the International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery, IEEE, 2009, pp. 262–266.

[64] M. Alaraj, J. Hou, T. Fukami, A neural network based human identification framework using ear images, in: Proceedings of the International technical conference of IEEE Region 10, IEEE, 2010, pp. 1595–1600.

[65] J. D. Bustard, M. S. Nixon, Toward unconstrained ear recognition from two-dimensional images, Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans 40 (3) (2010) 486–494.

[66] M. De Marsico, N. Michele, D. Riccio, HERO: human ear recognition against occlusions, in: Proceedings of the Computer Vision and Pattern Recognition Workshops, 2010, pp. 178–183.

[67] L. Gutierrez, P. Melin, M. Lopez, Modular neural network integrator for human recognition from ear images, in: Proceedings of the International Joint Conference on Neural Networks (IJCNN), IEEE, 2010, pp. 1–5.

[68] T.-S. Chan, A. Kumar, Reliable ear identification using 2-D quadrature filters, Pattern Recognition Letters 33 (14) (2012) 1870–1881.

[69] Z. Baoqing, M. Zhichun, J. Chen, D. Jiyuan, A robust algorithm for ear recognition under partial occlusion, in: Proceedings of the Chinese Control Conference, 2013, pp. 3800–3804.

[70] A. Kumar, T.-S. T. Chan, Robust ear identification using sparse representation of local texture descriptors, Pattern recognition 46 (1) (2013) 73–85.

[71] L. Lakshmanan, Efficient person authentication based on multi-level fusion of ear scores, Biometrics 2 (3) (2013) 97–106.

[72] S. Prakash, P. Gupta, An efficient ear recognition technique invariant to illumination and pose, Telecommunication Systems 52 (3) (2013) 1435–1448.

[73] A. Basit, M. Shoaib, A human ear recognition method using nonlinear curvelet feature subspace, International Journal of Computer Mathematics 91 (3) (2014) 616–624.

[74] P. Galdamez, A. Gonzalez Arrieta, M. Ramon, Ear recognition using a hybrid approach based on neural networks, in: Proceedings of the International Conference on Information Fusion, 2014, pp. 1–6.

[75] L. Jacob, G. Raju, Advances in Signal Processing and Intelligent Recognition Systems, Springer International Publishing, Cham, 2014, Ch. Ear Recognition Using Texture Features - A Novel Approach, pp. 1–12.

[76] T. Ying, Z. Debin, Z. Baihuan, Ear recognition based on weighted wavelet transform and DCT, in: Proceedings of the Chinese Conference on Control and Decision, IEEE, 2014, pp. 4410–4414.

[77] A. Benzaoui, A. Kheider, A. Boukrouche, Ear description and recognition using ELBP and wavelets, in: Proceedings of the International Conference on Applied Research in Computer Science and Engineering, 2015, pp. 1–6.

[78] A. Benzaoui, N. Hezil, A. Boukrouche, Identity recognition based on the external shape of the human ear, in: Proceedings of the International Conference on Applied Research in Computer Science and Engineering, IEEE, 2015, pp. 1–5.

[79] H. Bourouba, H. Doghmane, A. Benzaoui, A. H. Boukrouche, Ear recognition based on Multi-bags-of-features histogram, in: Proceedings of the International Conference on Control, Engineering Information Technology, 2015, pp. 1–6.

[80] A. Meraoumia, S. Chitroub, A. Bouridane, An automated ear identification system using Gabor filter responses, in: Proceedings of the International Conference on New Circuits and Systems, IEEE, 2015, pp. 1–4.

[81] A. Morales, M. Diaz, G. Llinas-Sanchez, M. Ferrer, Earprint recognition based on an ensemble of global and local features, in: Proceedings of the International Carnahan Conference on Security Technology, IEEE, 2015, pp. 253–258.

[82] P. Yan, K. W. Bowyer, Biometric recognition using 3D ear shape, Transactions on Pattern Analysis and Machine Intelligence 29 (8) (2007) 1297–1308.

[83] D. J. Hurley, M. S. Nixon, J. N. Carter, Force field energy functionals for image feature extraction, Image and Vision computing 20 (5) (2002) 311–317.

[84] J. Li, N. M. Allinson, A comprehensive review of current local features for computer vision, Neurocomputing 71 (10) (2008) 1771–1787.

[85] J. Krizaj, V. Struc, N. Pavesic, Adaptation of SIFT features for robust face recognition, in: Proceedings of the Image Analysis and Recognition, Springer, 2010, pp. 394–404.

[86] V. Ojansivu, E. Rahtu, J. Heikkilä, Rotation invariant local phase quantization for blur insensitive texture analysis, in: Proceedings of the International Conference on Pattern Recognition, IEEE, 2008, pp. 1–4.

[87] N. Damar, B. Fuhrer, Ear recognition using multi-scale histogram of oriented gradients, in: Proceedings of the Conference on Intelligent Information Hiding and Multimedia Signal Processing, 2012, pp. 21–24.

[88] N.-S. Vu, A. Caplier, Face recognition with patterns of oriented edge magnitudes, Computer Vision (2010) 313–326.

[89] M. A. Carreira-Perpinan, Compression neural networks for feature extraction: Application to human recognition from ear images, Master's thesis, Faculty of Informatics, Technical University of Madrid, Spain (1995).

[90] E. R. L. at University of Science & Technology Beijing, Introduction to USTB ear image databases, Online (2002).

[91] E. Gonzalez-Sanchez, Biometria de la oreja, Ph.D. thesis, Universidad de Las Palmas de Gran Canaria, Spain (2008).

[92] S. Prakash, P. Gupta, An efficient ear localization technique, Image and Vision Computing 30 (1) (2012) 38–50.

[93] R. Raposo, E. Hoyle, A. Peixinho, H. Proença, UBEAR: A dataset of ear images captured on-the-move in uncontrolled conditions, in: Proceedings of the Workshop on Computational Intelligence in Biometrics and Identity Management, IEEE, 2011, pp. 84–90.

[94] P. J. Phillips, H. Moon, P. J. Rauss, S. Rizvi, The FERET evaluation methodology for face recognition algorithms, Transactions on Pattern Analysis and Machine Intelligence 22 (10).

[95] T. Sim, S. Baker, M. Bsat, The CMU Pose, Illumination, and Expression (PIE) database, in: Proceedings of the International Conference on Automatic Face and Gesture Recognition, IEEE, 2002, pp. 46–51.

[96] K. Messer, J. Matas, J. Kittler, J. Luettin, G. Maitre, XM2VTSDB: The extended M2VTS database, in: Proceedings of the International conference on audio and video-based biometric person authentication, Vol. 964, Citeseer, 1999, pp. 965–966.

[97] G. Fahmy, A. El-Sherbeeny, S. Mandala, M. Abdel-Mottaleb, H. Ammar, The effect of lighting direction/condition on the performance of face recognition algorithms, in: Defense and Security Symposium, Proceedings of the International Society for Optics and Photonics, 2006, p. 62020J.

[98] G. B. Huang, M. Ramesh, T. Berg, E. Learned-Miller, Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments, Tech. Rep. 7–49, University of Massachusetts, Amherst (Oct. 2007).

[99] N. Kumar, A. C. Berg, P. N. Belhumeur, S. K. Nayar, Attribute and Simile Classifiers for Face Verification, in: Proceedings of the International Conference on Computer Vision, IEEE, 2009.

[100] H. Ng, S. Winkler, A data-driven approach to cleaning large face datasets, in: Proceedings of the International Conference on Image Processing, IEEE, 2014.

[101] D. Yi, Z. Lei, S. Liao, S. Z. Li, Learning Face Representation from Scratch, Tech. Rep. 1411.7923, preprint (arXiv 2014).

[102] B. F. Klare, B. Klein, E. Taborsky, A. Blanton, J. Cheney, K. Allen, P. Grother, A. Mah, M. Burge, A. K. Jain, Pushing the frontiers of unconstrained face detection and recognition: IARPA Janus Benchmark A, in: Proceedings of the Conference on Computer Vision and Pattern Recognition, IEEE, 2015.

[103] J. Klontz, B. Klare, S. Klum, A. Jain, M. Burge, Open Source Biometric Recognition, in: Proceedings of the Conference on Biometrics: Theory, Applications and Systems, IEEE, 2013.

[104] M. Pietikäinen, A. Hadid, G. Zhao, T. Ahonen, Computer Vision Using Local Binary Patterns, Computational Imaging and Vision, Springer, 2011.

[105] V. Ojansivu, J. Heikkilä, Blur insensitive texture classification using local phase quantization, in: Image and signal processing, Springer, 2008, pp. 236–243.

[106] J. Kannala, E. Rahtu, BSIF: Binarized statistical image features, in: Proceedings of the International Conference on Pattern Recognition, IEEE, 2012, pp. 1363–1366.

[107] N. Dalal, B. Triggs, Histograms of oriented gradients for human detection, in: Proceedings of the International Conference on Computer Vision and Patten Recognition, IEEE, 2005, pp. 886–893.

[108] D. Lowe, Distinctive Image Features from Scale-Invariant Keypoints, International Journal of Computer Vision 60 (2) (2004) 91–110.

[109] A. Morales, M. Ferrer, M. Diaz-Cabrera, E. Gonzalez, Analysis of local descriptors features and its robustness applied to ear recognition, in: Proceedings of the International Carnahan Conference on Security Technology, IEEE, 2013, pp. 1–5.

[110] J. Daugman, Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters, Journal of the Optical Society of America A 2 (7) (1985) 1160–1169.

[111] V. Struc, R. Gajsek, N. Pavesic, Principal Gabor filters for face recognition, in: Proceedings of the Conference on Biometrics: Theory, Applications and Systems, IEEE, 2009, pp. 1–6.

[112] V. Struc, N. Pavesic, Gabor-based kernel partial-least-squares discrimination features for face recognition, EURASIP Journal on Advances in Signal Processing 20 (1) (2009) 115–138.

[113] V. Struc, N. Pavesic, The complete gabor-fisher classifier for robust face recognition, EURASIP Journal on Advances in Signal Processing 2010 (2010) 1–26.

[114] Y. Peng, S. Wang, X. Long, B.-L. Lu, Discriminative graph regularized extreme learning machine and its application to face recognition, Neurocomputing 149 (2015) 340–353.

[115] Q. Gao, Y. Huang, X. Gao, W. Shen, H. Zhang, A novel semi-supervised learning for face recognition, Neurocomputing 152 (2015) 69–76.

[116] J. Zhang, J. Yang, J. Qian, J. Xu, Nearest orthogonal matrix representation for face recognition, Neurocomputing 151 (2015) 471–480.

[117] B. Yang, S. Chen, A comparative study on local binary pattern (LBP) based face recognition: LBP histogram versus LBP image, Neurocomputing 120 (2013) 365–379.

[118] E. Gonzalez, L. Alvarez, L. Mazorra, Normalization and feature extraction on ear images, in: Proceedings of the International Carnahan Conference on Security Technology, IEEE, 2012, pp. 97–104.

[119] X. Xiong, F. De la Torre, Supervised Descent Method and its Applications to Face Alignment, in: Proceedings of the Conference on Computer Vision and Pattern Recognition, IEEE, 2013.

[120] N. C. Camgoz, V. Struc, B. Gokberk, L. Akarun, A. A. Kindiroglu, Facial landmark localization in depth images using supervised ridge descent, in: Proceedings of the International Conference on Computer Vision Workshops, IEEE, 2015, pp. 136–141.

[121] J. Bustard, M. Nixon, 3D Morphable Model Construction for Robust Ear and Face Recognition, in: Proceedings of the Conference on Computer Vision and Pattern Recognition, IEEE, 2010, pp. 2582 – 2589.

[122] T. D. Kulkarni, P. Kohli, J. B. Tenenbaum, V. Mansinghka, Picture: A Probabilistic Programming Language for Scene Perception, in: Proceedings of the Conference on Computer Vision and Pattern Recognition, IEEE, 2015, pp. 4390–4399.