

Received November 27, 2018, accepted December 12, 2018, date of publication December 24, 2018, date of current version January 16, 2019.

Digital Object Identifier 10.1109/ACCESS.2018.2889395

# A Survey of Super-Resolution in Iris Biometrics With Evaluation of Dictionary-Learning

FERNANDO ALONSO-FERNANDEZ<sup>1</sup>, (Member, IEEE),  
REUBEN A. FARRUGIA<sup>2</sup>, (Senior Member, IEEE), JOSEF BIGUN<sup>1</sup>, (Fellow, IEEE),  
JULIAN FIERREZ<sup>3</sup>, (Member, IEEE), AND ESTER GONZALEZ-SOSA<sup>4</sup>

<sup>1</sup>School of Information Technology, Halmstad University, SE 301-18 Halmstad, Sweden

<sup>2</sup>Department of Communications and Computer Engineering, University of Malta, MSD 2080 Msida, Malta

<sup>3</sup>Biometrics and Data Pattern Analytics Lab-ATVS, Escuela Politecnica Superior, Universidad Autonoma de Madrid, 28049 Madrid, Spain

<sup>4</sup>Nokia Bell Labs, 28045 Madrid, Spain

Corresponding author: Fernando Alonso-Fernandez (feralo@hh.se)

This work was supported by the EU COST Action under Grant IC1106. The work of F. Alonso-Fernandez and J. Bigun was supported in part by the Swedish Research Council, in part by the Swedish Innovation Agency, and in part by the Swedish Knowledge Foundation through the CAISR/SIDUS-AIR projects. The work of J. Fierrez was supported by the Spanish MINECO/FEDER through the CogniMetrics Project under Grant TEC2015-70627-R. The authors acknowledge the Halmstad University Library for its support with the open access fees.

**ABSTRACT** The lack of resolution has a negative impact on the performance of image-based biometrics. While many generic super-resolution methods have been proposed to restore low-resolution images, they usually aim to enhance their visual appearance. However, an overall visual enhancement of biometric images does not necessarily correlate with a better recognition performance. Reconstruction approaches thus need to incorporate the specific information from the target biometric modality to effectively improve recognition performance. This paper presents a comprehensive survey of iris super-resolution approaches proposed in the literature. We have also adapted an eigen-patches' reconstruction method based on the principal component analysis eigen-transformation of local image patches. The structure of the iris is exploited by building a patch-position-dependent dictionary. In addition, image patches are restored separately, having their own reconstruction weights. This allows the solution to be locally optimized, helping to preserve local information. To evaluate the algorithm, we degraded the high-resolution images from the CASIA Interval V3 database. Different restorations were considered, with  $15 \times 15$  pixels being the smallest resolution evaluated. To the best of our knowledge, this is the smallest resolutions employed in the literature. The experimental framework is complemented with six publicly available iris comparators that were used to carry out biometric verification and identification experiments. The experimental results show that the proposed method significantly outperforms both the bilinear and bicubic interpolations at a very low resolution. The performance of a number of comparators attains an impressive equal error rate as low as 5% and a Top-1 accuracy of 77%–84% when considering the iris images of only  $15 \times 15$  pixels. These results clearly demonstrate the benefit of using trained super-resolution techniques to improve the quality of iris images prior to matching.

**INDEX TERMS** Iris hallucination, iris recognition, eigen-patch, super-resolution, PCA.

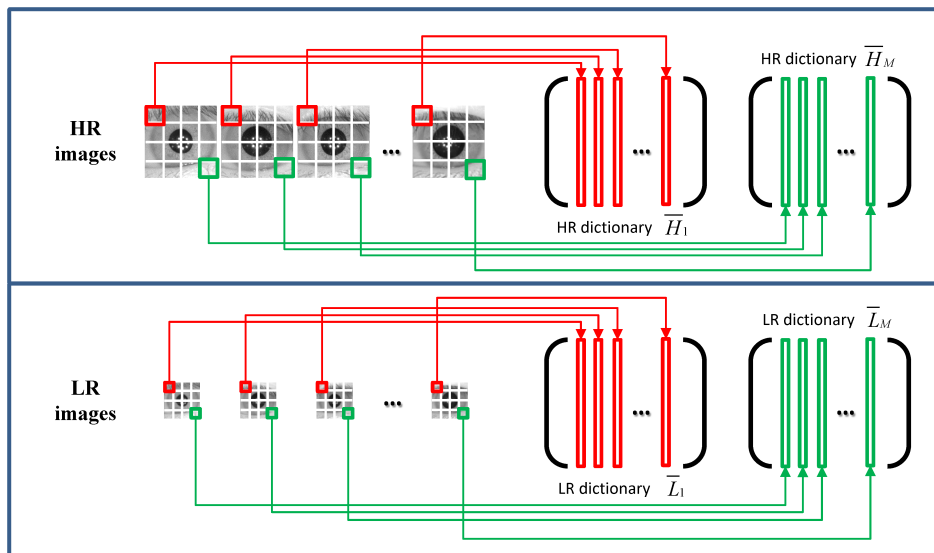
## I. INTRODUCTION

Iris recognition systems are known to achieve very high accuracy when captured in controlled environments and using the near infrared (NIR) spectrum. Nevertheless, recognition in applications such as mobile biometrics, surveillance and recognition at a distance has not reached the same level of maturity [1]. In these environments, the acquisition cannot be controlled, and performance can significantly drop due

to the lack of pixel resolution [2]. Furthermore, smart cards or remote applications may further reduce the quality of the image using JPEG2000 compression [3]. In this context, super-resolution techniques can be used to enhance the quality of low resolution images, in order to improve the recognition performance of biometric systems [4].

Two main categories of super-resolution methods are usually distinguished in the literature [5]: reconstruction-based





**FIGURE 2.** Dictionary construction of coupled low- and high-resolution patches using the position-patch principle.

to our experimental framework [15]–[18] in addition to the two previously employed [19], [20].

Simulation experiments conducted in this paper show that the proposed method achieves accuracies superior to those attained using bilinear and bicubic interpolation schemes. It is observed that recognition rates degrade more rapidly with both bilinear and bicubic interpolation as resolution decreases. At the smallest iris resolution ( $15 \times 15$  pixels), two particular comparators stand out, with EER values of  $\sim 5\%$  and a Top-1 accuracy of 77-84% in this extreme case. It is also shown that recognition performance is not significantly degraded with any given comparator until a resolution of only  $29 \times 29$  pixels when using our proposed PCA iris super-resolution method. This allows to reduce the storage or data transmission requirements, or to increase the distance between the user and the iris sensor, two important requirements for biometric technologies to achieve massive adoption [2]. We have also observed that, despite iris images reconstructed with PCA have better subjective quality, the image fidelity measures employed (PSNR and SSIM) do not have the same sensitivity to reductions in resolution. This is also acknowledged in other studies which have pointed out that image fidelity metrics do not behave equally under the same image degradation [21], [22]. In addition, each comparator has different behavior when resolution is reduced. Most of the comparators show a stable authentication performance until a certain resolution is employed, but the cut-off resolution is different for each one. On the other hand, one particular comparator does not suffer a significant degradation in performance, giving a consistent accuracy across nearly all resolutions. These differences in behavior of image fidelity metrics and biometric comparators among themselves highlight the necessity of adapting super-resolution techniques to cater for the particularities of

a specific biometric modality [10]. It is also crucial not to assess only the fidelity of the reconstruction in the visual sense, but to evaluate the capability of the reconstruction algorithm to improve authentication performance with the particular recognition features to be employed [9], [23].

The rest of the paper is organized as follows. The remaining of this section summarizes the main contributions of this paper. In Section II, we provide a comprehensive overview of the application of super-resolution techniques to iris biometrics. This is followed in Section III by the description of the proposed super-resolution algorithm, which adopts low- and high-resolution coupled dictionaries to learn an optimal up-scaling function for each patch. Then, this super-resolution algorithm is studied for iris biometrics. While this method is employed for iris super-resolution, it is general enough to be applied to other biometric modalities. The experimental framework, including database, protocol, and iris recognition algorithms employed, is given in Section IV, while results are given in Section V. Conclusions are finally drawn in Section VI.

### A. CONTRIBUTIONS

The contributions of this paper are as follows:

- A survey of super-resolution applied to iris biometrics (Table 1). We give a comprehensive overview of references found in the literature (we focus primarily on papers that appeared in IEEE Xplore, ScienceDirect or SpringerLink, as these appear to currently be the major sources of publications in the biometrics field). We provide a basic algorithmic descriptions of each approach, including the database employed for its evaluation, the smallest size of input low-resolution images considered, the features used for recognition experiments, and the reported recognition results (if any). We also provide

**TABLE 1. Overview of existing iris super-resolution works. Smallest low-resolution (LR) size refers to the smallest size of the input data used in the reported experiments (ED=Eye Diameter, refers to the average diameter of the iris in pixels; PI=Polar Image). Simulated downsampling indicates if the images employed in the study have been down-sampled from high-resolution reference images. The reported accuracy corresponds to the best accuracy obtained with the smallest iris size (shown in column 7). In all cases, near-infrared (NIR) data is used, except in the works [25], [34], [42], [45]. All other terms are explained in the text or in referenced papers.**

PIXEL DOMAIN: RECONSTRUCTION-BASED

Ref.	Algorithm	Input Data	Patch Based	Database	Input Images	Smallest LR Size	Simulated Downs.	Recognition Features	Accuracy	
									EER	Rank-1
[24]	Inverse Optimization	Iris	No	Own videos	15	30px ED	no	Iris Code	n/a	n/a
[25]	Auto-regressive models	Iris	No	Own videos	9	n/a	no	Iris Code	n/a	n/a
[26] [27]	Weighted average	Polar	No	MBGC Portal	5	90px ED	no	Iris Code	0.7%	n/a
[28]	Robust-mean average	Polar	No	MBGC Portal	var	90px ED	no	Iris Code	4.1%	n/a
[29]	Mean/median	Polar	No	MBGC Iris	10	220px ED	no	Iris Code	0.7%	n/a
[30]	PCT enhance + average	Polar	No	MBGC Iris	6	20px ED	yes	Log-Gabor	1.76%	n/a
[31]	Iterated back projection	Polar	No	CASIA 3.0	4	128 × 16 PI	yes	Iris Code	8.69%	n/a
[32] [33]	Weighted average	Polar	No	MBGC Portal	n/a	90px ED	no	Iris Code	2.58%	n/a
					6	110px ED	no	Iris Code	n/a	n/a
[34]	Wavefront coding + exp-weighted average	Polar	Yes	Own videos @3m	11	n/a	no	Log-Gabor	0%	n/a
[35]	Gaussian Process Regression	Polar	Yes	CASIA Long Range	n/a	300 × 40 PI	no	GLCM, moments, statistical features	n/a	n/a
[36]	Total Variation	Polar	No	CASIA Long Range	6	300 × 40 PI	no	GLCM	n/a	n/a

PIXEL DOMAIN: LEARNING-BASED

Ref.	Algorithm	Input Data	Patch Based	Database	Input Images	Smallest LR Size	Simulated Downs.	Recognition Features	Accuracy	
									EER	Rank-1
[37]	High frequency inference	Polar	Yes	CASIA	1	n/a	yes	Circular Filters	15%	89.7%
[38]	Multi-layer Perceptrons	Iris	Yes	CASIA Interval	1	53px ED	yes	Iris Code	1.39%	n/a
[39]	Score-level mapping	Polar	Yes	Q-FIRE	1	110px ED	no	Ordinal Measures	1.6%	n/a
[40]	Bayesian Modelling + Sparse Representation	Polar	Yes	CASIA	1	n/a	yes	n/a	2%	n/a
[11] [12]	PCA modelling	Iris	Yes	CASIA Interval	1	11px ED	yes	Log-Gabor, SIFT	6.44%	n/a
This paper						13px ED	yes	6 comparators	4.79%	84.2%
[41]	Neighbour Embedding	Iris	Yes	CASIA Interval	1	13px ED	yes	Log-Gabor, SIFT	3.58%	n/a
[42]	PCA modelling Neighbour Embedding	Iris	Yes	VSSIRIS	1	13px ED	yes	Log-Gabor, SIFT	4.1%	n/a
[43]	Convolutional Networks Random Forests	Polar	Yes	CASIA Mob 1	1	110px ED	no	Ordinal Measures	3.61%	n/a
				CASIA Mob 2	1	132px ED	no		1.82%	n/a
[44]	Convolutional Networks Stacked Auto-Encoders	Iris	Yes	CASIA Interval	1	13px ED	yes	Log-Gabor, SIFT	6.26%	n/a
[45]	Convolutional Networks	Iris	Yes	CASIA Interval	1	13px ED	yes	CG, QSW, SIFT	27.6%	n/a
				VSSIRIS	1	13px ED	yes		12%	n/a

FEATURE DOMAIN: LEARNING-BASED

Ref.	Algorithm	Input Data	Patch Based	Database	Input Images	Smallest LR Size	Simulated Downs.	Recognition Features	Accuracy	
									EER	Rank-1
[46]	Bayes MAP	Eigen-Iris	No	MBGC Iris	1	50px ED	yes	Eigen-Iris	4.5%	n/a
[47] [48]	Bayes MAP	Gabor	No	MBGC Portal	5	90px ED	no	Iris Code	0.5%	n/a
[49]	Markov Networks	Iris Code	No	Q-FIRE	n/a	110px ED	no	Ordinal Measures	2.6%	n/a

a taxonomy of existing iris super-resolution methods based on different factors (Figure 3), which include the domain of operation (pixel or feature domain), the input data source (iris image, polar image, or feature representations), or the spatial representation (if the method uses complete images or local image patches to carry out the reconstruction).

- A generic super-resolution method which employs low- and high-resolution coupled dictionaries to learn the optimal up-scaling function for each patch. This approach is able to recover important texture detail which is essential for most biometric recognition

systems, including iris. Following the work of Chen and Chien initially developed for face biometrics [13], a PCA Eigen-transformation is computed on local patches of the input low-resolution image. The high-resolution patch is then hallucinated as a linear combination of collocated high-resolution patches contained in the training dictionary. This way, every patch has its own optimal coefficients, which allows to reconstruct patches that are locally optimal and thus, to recover more texture detail. We evaluate our general super-resolution approach using iris images of resolution as low as 15 × 15 pixels (corresponding to an average iris diameter ~13 pixels). To the

## Iris Super-Resolution Reconstruction Methods

Upsampling Method	Domain of Operation	Input Data	Spatial Representation
<ul style="list-style-type: none"> <li>• Reconstruction-based</li> <li>• Learning-based</li> </ul>	<ul style="list-style-type: none"> <li>• Pixel domain</li> <li>• Feature domain</li> </ul>	<ul style="list-style-type: none"> <li>• Iris image</li> <li>• Polar image</li> <li>• Feature representation</li> </ul>	<ul style="list-style-type: none"> <li>• Complete image</li> <li>• Local patches</li> </ul>

**FIGURE 3.** Taxonomy of existing iris super-resolution reconstruction methods.

best of our knowledge, this iris size is much smaller than any other work reported in the literature, apart from ours (see Table 1). Another benefit is that, unlike other methods that restore the normalized polar image, our method is agnostic of the feature extraction method used, since it is applied directly on the iris low-resolution image. This makes the proposed method generic and independent from the iris comparator used. It also allows the use of features which are extracted directly from iris images without conversion to polar coordinates, see [50], [51].

- *Multi-algorithmic evaluation.* In our previous works [11], [12], we used only two iris comparators for the experimental study. Here, we use six different publicly available iris feature extraction methods from popular and state-of-the-art schemes [52] based on 1D log-Gabor filters [19], the SIFT operator [20], [50], local intensity variations in iris textures [15], the Discrete-Cosine Transform [16], cumulative-sum-based grey change analysis [17], and Gabor spatial filters [18]. The SIFT method exploits local features where discrete key-points are extracted directly from the iris region, while the other methods exploit other texture properties from the iris polar image computed according to Daugman’s rubber sheet model [53].
- *Comprehensive evaluation on a database of near-infrared iris images.* We employ in our experiments 1,872 images from the CASIA-Iris Interval v3 database of the Institute of Automation, Chinese Academy of Sciences (CASIA) [54]. High-resolution images, of size  $231 \times 231$  pixels and an average iris diameter of 210 pixels, are sub-sampled to reduce their size by 1/2, 1/4, 1/6, 1/8, 1/10, 1/12, 1/14 and 1/16. The latter corresponds to an image size of just  $15 \times 15$  pixels and an average iris diameter of  $\sim 13$  pixels. The performance of the iris super-resolution algorithm is measured in terms of PSNR and SSIM full reference metrics, which compute the fidelity between the original high-resolution image and the restored ones. Moreover, we carry out verification and identification experiments with the mentioned iris recognition algorithms, being one of the few studies in the literature that reports identification experiments. This is also the most extensive and up-to-date experimental framework in the context of iris super-resolution literature, providing extensive validation experiments.

## II. SUPER-RESOLUTION FOR IRIS BIOMETRICS

Super-resolution (SR) techniques aim to recover the missing high resolution (HR) image  $\bar{Y}$  given a low-resolution (LR) image  $\bar{X}$ . The low-resolution image is considered as a warped, blurred and down-sampled version of the high-resolution image. This can be mathematically expressed using

$$\bar{X} = DBW\bar{Y} + \bar{n} \tag{1}$$

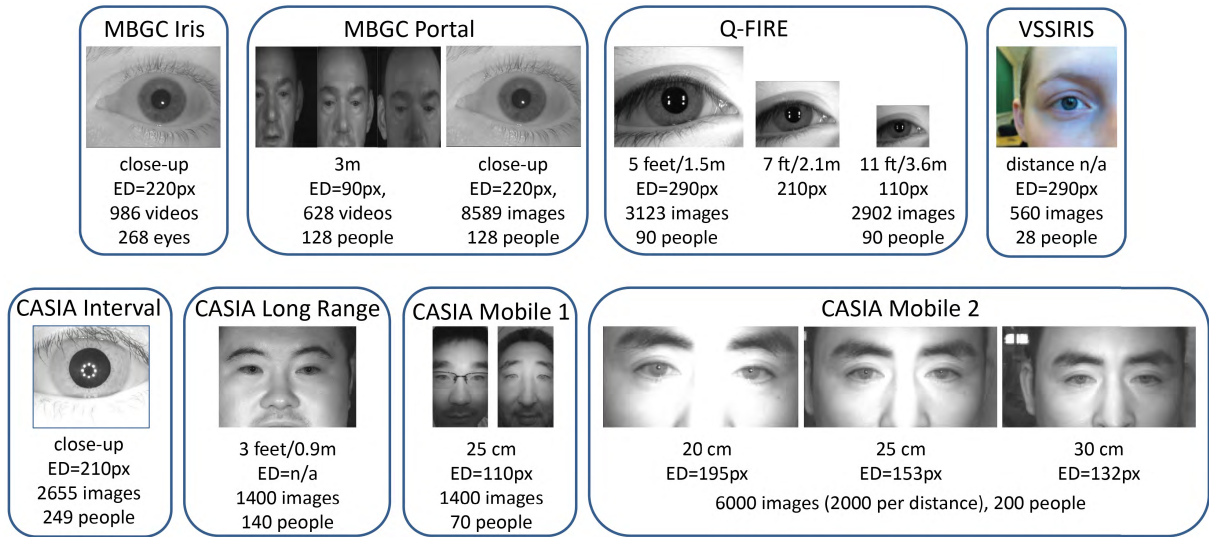
where  $W$  is the warping matrix,  $B$  is the blurring kernel (also called point spread function in some studies),  $D$  is the downsampling matrix, and  $\bar{n}$  represents additive noise. For simplicity, some works omit the warp matrix and noise, leading to

$$\bar{X} = DB\bar{Y} \tag{2}$$

Super-resolution techniques are classically divided into two categories [5]: reconstruction- and learning-based methods. *Reconstruction-based* methods register and fuse a number of low-resolution images to estimate the high-resolution image. Several images are aligned and combined in a pixel-wise manner to obtain a reconstructed image. Given a set of  $N$  images  $\bar{X}_i$ , the super-resolved image  $\tilde{Y}$  is estimated as

$$\tilde{Y}(x, y) = \frac{\sum_{i=1}^N w_i \bar{X}_i(x, y)}{\sum_{i=1}^N w_i} \tag{3}$$

where  $\tilde{Y}(x, y)$  is the intensity value at pixel  $(x, y)$  of the super-resolved image,  $\bar{X}_i(x, y)$  is the intensity value at the same location of the input image  $i$ , and  $w_i$  are the combination weights. While these methods can exploit the correlation and redundancies present in multiple frames to restore the missing detail, they cannot be employed in cases where only one image is available. Moreover, these methods are known to fail to reconstruct dynamic non-rigid objects and can only achieve small magnification factors. On the other hand, *learning-based* methods use coupled low- and high-resolution dictionaries to learn the mapping relations between low- and high-resolution image pairs in order to synthesize a high-resolution image from the observed low-resolution one. Learning-based methods have the advantage of needing only one image as input, and they generally allow to achieve higher magnification factors [5].



**FIGURE 4.** Samples of databases used in iris super-resolution research. Information is also given regarding: the distance between the individual and the acquisition sensor, the number of available images or videos, the number of individuals captured, and the average iris diameter in pixels of the images contained in the database (ED=Eye Diameter).

**A. TAXONOMY OF IRIS SUPER-RESOLUTION ALGORITHMS**

Table 1 gives a summary in chronological order of existing works on iris super-resolution. Apart from the distinction between reconstruction- and learning-based methods, they can be also categorized based on other factors, which are summarized in Figure 3:

- Domain of operation (pixel or feature domain). The majority of studies work in the *pixel domain* (top and middle part of Table 1), estimating pixel intensity values of the enhanced image. As a result, a new image with improved resolution is produced, which translates to a visual related enhancement. A few studies carry out the enhancement in the feature space (bottom part of Table 1), shifting the reconstruction operation from the pixel domain to the *feature domain* employed for recognition [46]–[49]. The latter approaches explicitly aim at improving the recognition performance, instead of the visual appearance. On the other hand, they have the disadvantage of being tied to a particular feature representation.
- Type of data used as input for enhancement (iris image, polar image, or feature representation). This is indicated in column 3 of Table 1. The majority of pixel-based methods super-resolve the *polar* image directly [53], while others reconstruct the entire *iris* image. The latter has the advantage of being usable with feature extraction methods that do not employ polar representation [50], [51]. Feature-based methods, on the other hand, receive as input a *feature* representation of the low-resolution image. Then, instead of producing an enhanced image as output, they estimate a feature representation of the reconstructed image.
- Spatial representation employed (complete image or local patches). This is indicated in column 4 of Table 1.

Some approaches, also called global methods, carry out reconstruction of the *complete image*. Patch-based methods, on the other hand, hallucinate *local patches* separately. The reconstructed high-resolution patches are then stitched together to form the high-resolution image. This allows each patch to have its own optimal reconstruction coefficients, providing better quality reconstructed prototypes with better local detail and lower distortion [8].

In addition to the aforementioned properties, we also provide information in Table 1 regarding:

- The database employed (indicated in column 5). As mentioned in the caption, nearly all studies make use of near-infrared (NIR) data. Some sample images from each database are also shown in Figure 4, together with their most representative information. In particular, we indicate: the distance to the acquisition sensor, the number of images or videos available, the number of individuals, and the average iris diameter of the images contained in the database. The following is a short description of each database, highlighting its most important features not contained in Figure 4. The Multiple Biometric Grand Challenge Portal database, or *MBGC Portal* [55], contains face video sequences of people walking naturally through a portal located 3m from a fixed-focal-length NIR camera (Pulnix TM-4000CL). It also contains iris images of good quality captured with a close-up NIR iris sensor from the same individuals. The *MBGC Iris* video database contains videos of irises collected using a close-up NIR iris camera (Iridian LG EOU 2200). The *CASIA Long Range* database contains face images captured at 3 feet (0.9 m) with a high resolution NIR camera. The *CASIA*

*Iris Interval* database has iris images captured with a close-up NIR sensor. The *CASIA Iris Mobile* v1.0 and v2.0 databases contain face images at various distances captured with a NIR imaging module connected to a smart-phone by USB. The *Q-FIRE* database [56] has iris videos captured at 5, 7, and 11 feet away with a Dalsa 4M30 infrared camera and a Tamron AF 70-300mm telephoto lens. And finally, the *VSSIRIS* database [57]. This is the only database in visible range, containing images captured using the rear camera of two different smart-phones (Apple iPhone 5S and Nokia Lumia 1020).

- The number of low-resolution images used by the reconstruction algorithm to generate a high-resolution representation (indicated in column 6). Existing reconstruction-based methods employ a variable number which goes from four to fifteen, while learning-based algorithms in the pixel domain only employ one. It is also found that many learning-based algorithms working in the feature domain employ several images as input. However, since the mapping between low- and high-resolution manifolds is learned, we classify these methods as learning-based. Indeed, nothing prevents learning-based methods to employ more than one image, although one of their main advantages is that they can generate a high-resolution representation from only one low-resolution sample.
- The smallest size of the input low-resolution image (indicated in column 7). Some databases are naturally captured at a certain distance, as it can be observed in Figure 4. For example: *CASIA Iris Mobile* v2.0 (with images having an average iris diameter of 132 pixels), *Q-FIRE* (110 pixels), or the *MBGC Portal* database (90 pixels). To achieve a smaller image size, a number of studies perform sub-sampling of high-resolution images (indicated in column 8). Simulated downsampling is a common practice in the super-resolution literature [14], [58], due to the lack of databases with very low-resolution images. For example, in the work [46], the authors down-sampled images from the *MBGC Iris* video database to an average iris diameter of 50 pixels. Images of the same database were reduced to an average diameter of 20 pixels in the work [30]. The *CASIA Interval* database has been also used for the same purpose in a number of studies [11], [12], [38], [44]. The average iris diameter of sub-sampled images in these studies ranged from 11 to 53 pixels. Finally, in our studies with the *VSSIRIS* database [42], [45], we down-sampled iris images to an average diameter of 13 pixels.
- The features used to carry out recognition experiments (indicated in column 9). Most studies only employ the popular *Iris Code* representation [53]. Very few works compare two or more feature extraction methods. Among those, the present paper stands out as the only one employing six different comparators.

- The biometric authentication results reported (indicated in columns 10-11) when images of the smallest size are used for recognition purposes. The present paper is the only one, together with [37], which reports identification experiments. We further analyze these authentication results by plotting in Figure 5 the verification accuracy given in Table 1 against the iris size. Although the results are not directly comparable due to different databases and feature extraction methods being used, there is an inverse proportion between the EER and the diameter of the iris. Among the methods making use of very small iris images, our works with PCA [11], [12], [42] are among the most competitive. Recent studies adapting deep-learning frameworks [44], [45] still report an accuracy significantly worse in some cases. There is one reconstruction-based method using PCT enhancement [30] which also stands out for its excellent performance with an iris diameter of only 20 pixels. It is also worth noting that other works employing images with a higher diameter do not provide a significantly better performance, see for example the work [38], based on Multi-layer Perceptrons, or the work [46], based on Bayes MAP probability estimation. The same appreciation can be done with the works employing the *MBGC Portal*, *Q-FIRE* or *CASIA Mobile* databases. Although these employ images with a iris diameter (in the range of 90-130 pixels), their performance is not much better in some cases. These results suggest that there is still room for improving the performance of super-resolution methods in iris biometrics.

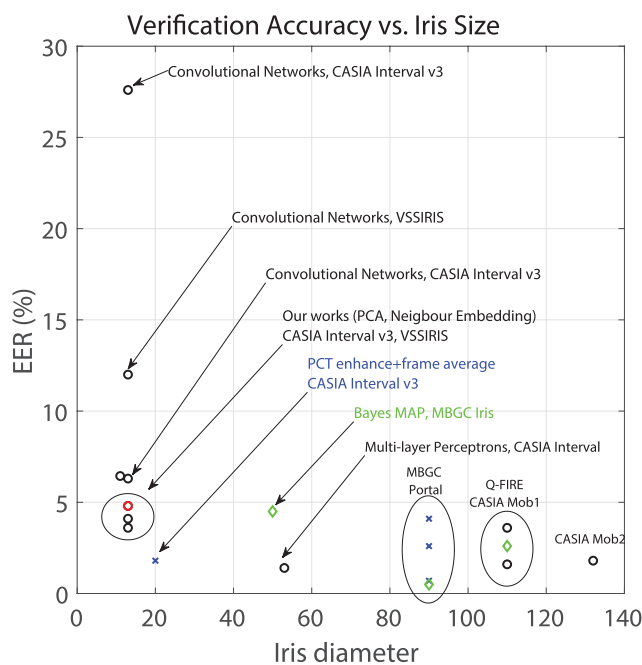


FIGURE 5. Verification accuracy reported by existing iris super-resolution studies vs. employed iris size (in pixels). The algorithm and/or database used is also mentioned.

In the remaining of this section, we provide a brief description of the works summarized in Table 1, categorized by the domain of operation (pixel or feature domain), and by the upsampling approach (reconstruction- or learning-based).

### B. RECONSTRUCTION-BASED METHODS IN THE PIXEL DOMAIN

Reconstruction-based methods for iris images started in 2006 with the work of Barnard *et al.* [24], where they employed a multi-lens imaging hardware system to capture multiple iris images. They carried out the reconstruction by modelling the least square inverse problem associated with Equation 1. For this purpose, the blurring kernel,  $B$ , the warp function,  $W$ , and the downsampling function,  $D$ , were estimated (the noise term was omitted). To minimize the reconstruction error, they used the conjugate gradient method (CGLS). In their experiments, they employed up to 15 low-resolution images as input. They measured the quality of the reconstruction by computing the Hamming Distance between the Iris Code of a reference high-resolution image, and the corresponding reconstructed image. Experiments showed a reduction in the Hamming Distance in comparison with the distance to a low-resolution image.

Later in 2007, Fahmy [25] proposed an algorithm where high-resolution iris images were estimated using an auto-regressive model to fuse a sequence of low-resolution iris images. He first applied a cross-correlation model to register iris images from consecutive frames of videos captured 3 feet away of the subject. Then, an iris image which is 4-times higher in resolution was constructed from every 9 low resolution images. This process was iterated until an image which is 16-times higher was obtained. Two drawbacks of this method are that registration was done using the whole eye image, and that they employed only focused images. These assumptions can be problematic in unconstrained conditions, where off-angle or out-of-focus images may be present.

Video data from the Multiple Biometric Grand Challenge database [55] has been used in several studies [26]–[29], where several polar images were aligned and combined pixel-wise to obtain a reconstructed image according to Equation 3. Hollingsworth *et al.* [29] created a single average polar image from 10 multiple frames of a frontal iris video. Data employed was from the MBGC Iris NIR database, having an average iris diameter of 220 pixels. The ten best focused images from each video were selected by employing the filter kernel proposed in [59]. They tested both the mean and median fusion, concluding that the mean is better, since it employs statistics from all available pixels. The median, on the other hand, only employs statistics of one or two pixels. They achieved an EER of 0.7% by fusion of 10 frames, compared to an EER of 1.56% obtained by employing only one-gallery and one-probe frame. Nguyen *et al.* [28] employed the robust mean, which consist of fusing intensity values of an individual pixel over multiple frames by taking the mean of values within two standard deviation from the mean. Data employed was from the MBGC Portal NIR database, with an average

iris diameter of 90 pixels. Authentication experiments were done by comparing super-resolved images to high resolution still iris images captured with a close-up NIR camera, which are also provided with the MBGC Portal dataset (Figure 4). The authors argued that in unconstrained environments like the MBGC Portal database, extreme pixel values can appear in different locations in different frames due to reflections, shadows, eyelids, or eyelashes. For this reason, the proposed approach is more robust against unexpected extreme pixel values. The obtained EER is of 4.1%, contrasted to 4.6% when the approach of [29] was applied.

Considering that frames in an iris video sequence can have different quality when captured in adverse conditions, Nguyen *et al.* [26], [27] employed quality measures to compute the weights of Equation 3. In the paper [26], they employed the focus level of each frame, which was measured by evaluating the high frequency total energy of the image. A high-resolution image was then estimated using the focus-score weighted average of the available frames. Heavily de-focused frames were discarded from further processing, while the others were fused to super-resolve the iris image. With the proposed approach, they achieved an EER of 2.1% using the MBGC Portal database. In the paper [27], the authors combined four quality factors (focus, off-angle, motion blur, and illumination variation) into a unified quality score for each iris frame. For this purpose, they employed the Dempster-Shafer theory proposed in [60]. Another novelty of this work was that instead of using the conventional weighted average of Equation 3, frames were fused by using an exponential weighted average. Experiments were also carried out to determine the optimum number of frames for fusion. The authors concluded that 5 is the optimal number of frames to fuse, although they acknowledged that a different dataset may lead to a different number if the acquisition conditions are different. The reported EER in this case is of 0.7%. When the number of frames increases beyond 5, they observed that the poor quality of the additional frames counteracted the introduction of extra information. Following a similar vein, Othman *et al.* [32], [33] expanded this idea by computing the quality of local image patches. For this purpose, they estimated a Gaussian Mixture Model (GMM) of clean iris texture distribution. Then, during the fusion, each pixel was weighted individually with the quality value of the associated local patch, instead of employing a single quality score for the whole image. This way, regions with better quality contribute more in the reconstruction of the fused image than regions with poorer quality. Using the MBGC Portal database, they reported an EER of 2.58%, compared to an EER of 4.9% obtained by simple pixel intensity average. They also employed the Q-FIRE database, which contains videos captured at 5, 7, and 11 feet away with a telephoto lens. The number of input images employed for reconstruction varied from 2 to 10, with the frames ranked according to their quality. In the experiments, the authors observed that performance with Q-FIRE improved until the best 6 frames were chosen, then the error increased when extra frames were



added. The reported FRR @ FAR=0.1% with this database is of 2.54% (images captured 5 feet away), 4.37% (7 feet), and 2.04% (11 feet).

Some works have included specific preprocessing for image enhancement purposes. For example, Jillela *et al.* [30] applied Principal Component Transform (PCT, variation of PCA) to polar iris images of the MBGC Iris database in order to highlight the variance information among the pixel intensity. Then, the enhanced images were fused by image averaging. The optimum number of frames per video was empirically chosen as 6, with the best quality frames selected manually. Low-resolution data was generated by sub-sampling the original images by a factor of 1/2, 1/4, and 1/8. This resulted in low-resolution image sets with an average iris diameter of 110, 50, and 20 pixels, respectively. Authentication experiments were done by comparing super-resolved images to a separate set of original high-resolution images set aside as gallery set. The reported EER using images with the lowest resolution is of 1.76%. When no reconstruction is carried out (i.e. each low-resolution frame is compared separately against gallery frames), the reported EER is of 6.09%, highlighting the benefit of the proposed approach. In the work by Hsieh *et al.* [34] the authors incorporated optical wavefront coding techniques [61] to increase the depth of field (DoF) in long-range iris portal acquisition. This was achieved by optimizing the optical architecture of the acquisition hardware. An extended depth of field (EDoF) allows a higher capture volume as the subjects walks to the camera. They also exploited image quality measures to weight the contribution of low-resolution images by exponential weighted average. They employed their own video data from 16 subjects captured with a telephoto lens. Enrolment images were captured with subjects standing at 3m, and test images were captured at 11 defocus positions (from -15 to +15 cm in 3 cm intervals). For each defocus position, two images were captured, one without the EDoF hardware, and one with the EDoF hardware. Quality of each frame was assessed by calculating the Hamming distance between all possible pairs of the 11 test images, and then computing the average distance of all pair-wise comparisons associated to each test image. A high quality image is expected to have a lower average value, and vice versa. A novelty of this approach with respect to the already mentioned approaches of this section is that reconstruction is made in local patches. When the 11 test images captured with the EDoF hardware are fused following the proposed method, the reported recognition results are EER=0%, and FFR=0% at a FAR of 0.1%. On the contrary, when the images are captured without the EDoF hardware, the paper reports an EER of 11.5%, and a FRR of 37% at a FAR of 0.1%.

The CASIA Long Range database has also been used in a number of studies [35], [36]. Deshpande and Patavardhan [35] employed Gaussian Process Regression (GPR) and Enhanced Iterated Back Projection (EIBP) to super-resolve iris images. The best frame was selected

as reference for alignment purposes by using the Discrete Cosine Transform. To account for local image deformations, they carried out reconstruction in local patches of the polar image. A threshold is applied to the intensity variance of each patch. If the variance is higher than the threshold, the patch is reconstructed with GPR, otherwise it is reconstructed with EIBP. The GPR is a time consuming process, therefore patches with less amount of information (measured by their variance of intensity) are processed with the faster EIBP algorithm. Performance was evaluated in [35] by downsampling iris images, and then super-resolving them. The authors reported several image fidelity measures in the pixel domain between original high-quality images and reconstructed images. Recognition results were also reported, with a FAR of 3.86% and a FRR of 4.21%. although no information is given regarding the size of the low-resolution images involved in the authentication experiments. The same authors [36] employed an enhanced Total Variation regularization algorithm [62] to super-resolve iris images. Low resolution input images were first de-blurred in order to remove motion blur and then, motion estimation between consecutive image frames was computed. In the regularization process, the estimated blur kernel and motion vectors were taken into account to iteratively generate a high resolution reconstructed image. The authors employed six polar images of 300×40 pixels to estimate one super-resolved image of twice the input resolution (600×80 pixels). The authors evaluated the algorithms by reporting several image fidelity and textural measures between original high-quality images and reconstructed images. However, no person authentication experiments were reported. Iterative back projection was also employed by Ren *et al.* [31] in a previous work. A frame was randomly selected as reference for alignment, which was done in the Fourier domain. The output high-resolution image was initialized by upsampling the first low-resolution frame using nearest neighbor interpolation. Then, the estimated high-resolution image was iteratively updated with the gradient of the total square error in the pixel domain between the reference image and the low-resolution frames. When the total square error achieved a threshold value, the iterative process was finished. The authors employed iris images from the CASIA 3.0 database in their experiments. Low-resolution data was generated by sub-sampling polar images of 512×64 pixels to reduce their size by a factor of 1/2, and 1/4. The resulting polar images were of size 256×32, and 128×16 pixels, respectively. The EER reported by the authors without reconstruction is of 12.75% (with polar images of 256×32 pixels) and 13.7% (polar images of 128×16). Sub-sampled polar images were then reconstructed to their original size of 512×64. The number of input images evaluated for reconstruction was 2, 4 and 6, concluding that 4 images is a good compromise between performance and processing time. With the proposed algorithm, the reported EER is of 6.87% and 8.69% (using polar images of 256×32 and 128×16 pixels, respectively).

### C. LEARNING-BASED METHODS IN THE PIXEL DOMAIN

Learning-based iris reconstruction was first proposed in 2003 by Huang *et al.* [37]. In this method, the probabilistic relation between different frequency bands is learned, in order to predict the missing high-frequency information of low-resolution images. It is based on the general purpose method by Freeman *et al.* [63]. The training set of high-resolution polar images is first pre-processed as follows. Each high-resolution image is separated in three bands: low-frequency, by downsampling and upsampling the high-resolution image; medium-frequency, by applying a Circular Symmetric Filter (CSF); and high-frequency, by subtracting the high-resolution and the low-frequency images. Images are then divided into patches, which each position having associated a set of low-, medium-, and high-frequency patches. Given an input low-resolution image to be reconstructed, it is first up-sampled by cubic interpolation. Then, a feature image is constructed by applying a Circular Symmetric Filter (CSF) to extract the medium-frequency information. The image is then divided into patches. For each patch, the 200 patch sets of the training set whose feature vectors are closest to the input patch are selected using the L1 distance. The best matching set from this sub-set is then computed based on spatial constraints at adjacent patch borders, and the corresponding high-frequency patch is selected. A high-frequency image is then obtained by stitching together the resulting high-frequency patches. The output reconstructed image is finally obtained by adding the high-frequency image to the test input image. In the experiments reported in this paper, low resolution data was generated by sub-sampling images from the CASIA dataset to three different low-resolutions (not specified). The experiments report a significant improvement in rank-1 and EER metrics in comparison with cubic interpolation, and with the original method described in [63].

Shin *et al.* [38] employed multiple Multi-layer Perceptrons (MLP) to restore local iris patches. Each block of the input image is classified into one of 3 types (vertical, horizontal, and non-edge) based on difference of pixel intensities. Three MLPs are trained, one per type of block, to estimate selected pixel values of the high-resolution patch. An advantage of this method is that it does not require accurate image registration. Reconstructed blocks are then assembled together, and missing pixels are filled by bilinear interpolation. This is because the MLPs are not trained to predict all pixels of the high-resolution patch, but only a part of them. Low-resolution data was generated in [38] by sub-sampling images from the CASIA Interval v3 database to 1/3 and 1/4 of the original image size. This resulted in image sets with an average iris diameter of 70 and 53 pixels. The MLPs were trained using 12 randomly selected images from different eyes. The reported EER using the smallest images is of 1.39%, compared with 1.49% by bilinear interpolation, or 0.89% with original high-resolution images.

Sparse representation in over-complete dictionaries was used in the work of Aljadaany *et al.* [40]. Traditional approaches in this regard, such as K-Singular Value Decomposition (K-SVD), have the limitation that the number of dictionary items and the number of sparse coefficients has to be predefined. To overcome this limitation, the authors used a non-parametric Bayesian approach, named Beta Process (BP), to build the discriminative over-complete dictionary and discover the necessary parameters automatically. During the training phase, high-resolution polar images are down-sampled to create their low-resolution counterparts, which are then used to learn the relationship between high- and low-resolution patches. The dictionaries in both manifolds are assumed to have the same sparse weights, which simplifies the reconstruction process. During the testing phase, sparse weights of the low-resolution image are first computed. Then, the weights are transferred to the high-resolution manifold, which are then used to calculate the conditional expectation of a high resolution iris image. This approach was evaluated in a subset of the CASIA database. Low-resolution images were generated by sub-sampling images to 25% of their original size and adding Gaussian noise. The authors reported and improved recognition performance with the proposed method (EER  $\sim 2\%$ ), in comparison with linear interpolation (EER  $\sim 3\%$ ).

Instead of synthesizing high-resolution iris images, Liu *et al.* [39] learned a non-linear mapping function in which homogeneous (high resolution-high resolution) and heterogeneous (high resolution-low resolution) comparison scores are non-linearly mapped into a common high dimensional space. In this space, separation between inter- and intra-class distributions are maximized regardless whether they originate from homogeneous or heterogeneous samples. This converts the multi-class problem (discriminating between different identities) into a two-class problem (discriminating between inter- and intra-class comparisons). Each class is composed of two sets: homogeneous comparisons (high-to-high resolution) and heterogeneous comparisons (high-to-low resolution). During the training phase, all possible inter/intra-class comparisons of training data are carried out. Polar images are divided into patches, and comparisons are done for each patch separately, so given two images, their comparison results in a set of patch-scores. During the testing phase, a low-resolution test sample is compared to the gallery of high-resolution templates, and the resulting scores are mapped to the learned common space, where a rejection/acceptance decision can be made. The authors employed the Q-FIRE database. Data captured at 5 and 11 feet away were selected respectively as high- and low-resolution sets. Eye detection and iris segmentation were carried out in the input iris videos, resulting in multiple images available for each subject. Heavily de-focused or occluded images were discarded for further use. A total of 1400 images (700 low- and 700 high-resolution images) from 100 different classes were used as training data. The reported EER with

the proposed method is of  $1.6 \pm 0.92\%$  (10-fold validation), compared with  $\sim 4.6\%$  without score mapping.

In the work [11], we presented an iris reconstruction method based on Principal Component Analysis (PCA) that is the basis of the present paper. The technique is inspired by the system of [13] for face images. Given a test low-resolution image, a PCA Eigen-transformation is conducted for each patch using a set of low-resolution training images. After the reconstruction weights are computed in the low-resolution manifold, they are transferred to the high-resolution manifold, where high-resolution patches are then reconstructed using a linear combination of collocated high-resolution patches from training images. The method was evaluated using the CASIA Interval v3 database, and Log-Gabor wavelets [19] as feature extraction method. Low-resolution data was generated by sub-sampling high-resolution images to reduce their size by a factor from 1/2 to 1/18 (the latter corresponding to an average iris diameter of 11 pixels). The paper reported an EER of 6.44% with images of the lowest resolution, compared with 12.23% when bi-cubic reconstruction is used.

The current paper extends our previous studies with the PCA method, including additional comparators and experiments. A limitation of this method is that it assumes that low- and high-resolution manifolds have similar local geometrical structure. Reconstruction weights are estimated on the low-resolution manifold, and they are simply transferred to the high-resolution manifold. However, the geometrical structure of the low-resolution manifold is distorted by the one-to-many relationship between low- and high-resolution patches [64]. Therefore, the reconstruction weights estimated on the low-resolution manifold do not necessarily correlate with the actual weights needed to reconstruct the unknown high-resolution patch. To cope with this limitation, we later considered to use iterative neighbor embedding of local patches (LINE) [41], where the geometry of the low- and high-reconstruction manifolds are jointly taken into account during the reconstruction. During the testing phase, the reconstruction weights are computed by minimizing a regularization function that considers the distance of the input patch to the training dictionary both in the low- and in the high-resolution manifolds. The lowest resolution evaluated consisted of images with an average iris diameter of 13 pixels. The LINE method compared well with the PCA method, showing additional performance improvements at very low resolutions. The PCA and LINE methods were also evaluated in [42] using smart-phone images from the VSSIRIS database [57]. In this work, high-resolution images were down-scaled by a factor of 1/22 (corresponding to an average iris diameter of  $\sim 13$  pixels). The experiments showed a superior performance of the trained reconstruction approaches in comparison to bilinear or bicubic methods, with the LINE approach showing better performance than PCA. The best recognition rates reported were 4.64% (PCA) and 4.1% (LINE) with the fusion of log-Gabor wavelets and the SIFT operator.

Recent studies have also adapted deep-learning frameworks to the task of iris super-resolution [43]–[45]. Zhang *et al.* [43] adapted Super-Resolution Convolutional Neural Networks (SRCNN) [65] and Super-Resolution Forests (SRF) [66] to reconstruct iris images in the polar domain. The SRCNN employed learns the non-linear mapping function between low- and high-resolution images with 3 layers: the first one extract feature maps of low-resolution patches, the second one maps these feature vectors into feature maps of corresponding high-resolution patches, and the last one aggregates high-resolution patches to generate the output image. The loss function employed is the mean squared error between the reconstructed images and the corresponding ground-truth high-resolution image, which in turn corresponds to a fidelity measure between images, and not to a performance metric. This is common to all iris super-resolution studies that employ deep-learning frameworks, which may explain that their performance is still behind methods not based on deep-learning. In the SRF method, Random Forest are used to directly map low-resolution patches to high-resolution patches. During tree growing, a regularized objective function that operates on both output and input domains is used, so higher quality results can be achieved. Training of SRCNN and SRF was done with 91 non-iris images, of use in other super-resolution studies. The two methods were tested using the CASIA Mobile v1.0 and v2.0 databases, which contain near-infrared (NIR) images captured with smart-phones by using a NIR imaging module. The EER achieved with CASIA Mobile v1.0 is of 3.65% (SRCNN) and 3.61% (SRF), compared with 3.89% obtained using iris images without enhancement. These results corresponds to the score fusion of left and right iris. With CASIA Mobile v2.0, the authors reported single-eye experiments comparing images captured at various distances, namely 20-20 cm, 20-25 cm, and 20-30 cm. Reported experiments show that applying the SRCNN and SRF methods results in better performance in comparison with employing images without enhancement, specially at low FAR.

In the works by Ribeiro *et al.* [44], [45], the authors employed several deep learning methods to reconstruct iris images. In these works, images were divided into patches, which were then reconstructed separately using each respective network to obtain high-resolution patches. The work [44] employed Super-Resolution Convolutional Neural Networks (SRCNN) [65] and Stacked Auto-Encoders (SAE) [67]. The experimental framework used the CASIA Interval v3 database as in the work [11]. High-resolution images were down-scaled by a factor from 1/2 to 1/16 (the latter corresponding to an average iris diameter of 13 pixels). Using Log-Gabor wavelets, the PCA method of [11] still shows better performance at the lowest resolution (EER of 4.79% vs. 6.26%). However, the paper evaluated a second comparator based on local SIFT key-points [20], with which the deep-learning methods showed better recognition performance (EER of 19.5% with PCA vs. 17.26% with SRCNN).

In the work [45], the authors evaluated Super-Resolution Convolutional Neural Networks (SRCNN) [65], Very Deep Convolutional Neural Networks (VDCNN) inspired by VGG-net [68], and Super-Resolution Generative Adversarial Networks (SRGAN) [69]. Besides the CASIA Interval v3 database, the authors also employed smart-phone images from the VSSIRIS database in their experiments. As in [44], images of both databases were down-sampled up to an average iris diameter of 13 pixels. They also explored the use of different databases to train the networks for the iris reconstruction task. The databases employed include texture databases, natural image databases, and iris databases. The iris databases contain both near-infrared and visible wavelength images, while the other two categories only include visible data. An interesting finding of this paper is that using texture databases or iris images from other databases for training provides good reconstruction results, even if captured with different lightning. At the lowest resolution, the best recognition performance with CASIA images was given by SRCNN, with a reported EER of 27.6%. With the VSSIRIS database, the best performance at the lowest resolution was given by SRGAN, with an EER of 12%. It should be mentioned nevertheless that the recognition features used in [45] are different than those in previous studies with the same databases [42], [44], which might explain the differences in performance.

#### D. LEARNING-BASED METHODS IN THE FEATURE DOMAIN

Instead of super-resolving pixel intensity values, the methods in this section super-resolve images in the feature space. A commonality found in most of them is that they employ several input images, so they could be considered to be reconstruction-based. However, we categorize them as learning-based since the mapping relation between low- and high-resolution images is learned using training dictionaries.

In the work published in 2011, Nguyen *et al.* [46] carried out the reconstruction by modelling the inverse problem associated with Equation 1. However, they replaced the low- and high-resolution images  $\bar{X}$  and  $\bar{Y}$  with their PCA feature representations [70] estimated from gallery images in the polar domain. Given a low-resolution test image, it is first projected onto the PCA space. Then, a high-resolution iris image is estimated by Bayes MAP probability estimation, which is computed using iterative steepest descent. The method was tested with data from the MBGC Iris NIR database. Two high quality iris images were selected per subject, one used as gallery, while the other was degraded by Gaussian blurring, random warping and downsampling by a factor of 1/4 to create a series of 16 low resolution images. The average iris diameter of low-resolution images in this case was of 50 pixels. In addition, the method was compared with the quality-weighted pixel average technique of the same authors, published in [26]. The effectiveness of the reconstruction was measured by computing the distance of the reconstructed feature vector to the true feature vector

of the original high-resolution image. Reported experiments showed that the proposed feature-domain approach produces closer features than the method in [26], as well as better recognition performance. The proposed feature-domain approach achieved an EER of 4.5%, compared to an EER of 10% obtained when using the pixel average method.

Nguyen *et al.* identified as a drawback of the previous approach that linear features such as PCA are not optimal for recognition, when compared to nonlinear ones such as iris codes extracted from Gabor phase-quadrant encoding [53]. Since Gabor-based features are shown to be one of the most discriminant features for face [71] and iris [72], they proposed to super-resolve iris images in the Gabor domain [47], [48]. A challenge of this approach comes from the difficulty of modelling the relationship between low- and high-resolution images in non-linear feature domains. However, they observed that the response to a Gabor wavelet is linear, whilst the nonlinearity of the process comes from the phase-quadrant encoding applied to compute the iris code. Thus, following the same strategy of [46], they replaced the low- and high-resolution images  $\bar{X}$  and  $\bar{Y}$  of Equation 1 with the complex-valued responses of polar images to Gabor wavelets. In this occasion, the method was tested with data from the MBGC Portal NIR database. Four video sequences of each identity were matched against still high-resolution images. For each video, the frames with the best quality were selected according to the Depmster-Shafer quality assessment of [60]. A threshold to remove poor quality frames was selected through experimentation. With the proposed approach, they achieved an EER of 0.5%. The method outperformed other approaches evaluated in the paper, including bicubic (EER=1.8%), pixel average [29] (EER=1.4%), quality-weighted pixel average [26] (EER=0.9%), PCA feature representation [46] (EER=4.8%), or LDA feature representation (EER=2.1%).

Lastly, Liu *et al.* [49] learnt the statistical relationship between a set of binary codes from low-resolution images, and the binary code of the corresponding high-resolution image. For this purpose, they employed Markov networks. The co-occurrence of neighboring bits in the high-resolution iris code was also modeled with the network. Besides the non-linear relationship between feature codes of low- and high-resolution iris images, the Markov model is also able to produce a weight mask which measures the reliability of each bit in the enhanced iris code. This weight mask can be used in the computation of the Hamming Distance between two iris codes to further enhance recognition accuracy. The authors employed the Q-FIRE database. Images captured at 5 and 11 feet away were selected respectively as high- and low-resolution sets. Eye detection and iris segmentation was carried out on the input iris videos, resulting in multiple iris for each subject. Heavily de-focused or occluded images were discarded for further use. A total of 1000 images (500 low- and 500 high-resolution images) were used as training data. The reported EER with the proposed method is of 2.6%. Reported experiments also show that it outperforms other

existing algorithms, including [27], [29], and [47], specially at low FAR. In a later work with the same database [39] (presented in Section II-C), the same authors improved the EER further to 1.6±0.92%.

### III. EIGEN-PATCH IRIS SUPER-RESOLUTION

The block diagram of the proposed iris super-resolution method is shown in Figure 1, which is based on the eigen-patch hallucination method for face images proposed by Chen and Chien [13]. Each input image  $\bar{X}$  is first separated into overlapping patches. An eigen-transformation is then conducted on each patch using collocated patches of low-resolution iris images from a training set  $\{\bar{L}\}$ , in order to obtain the optimal reconstruction weights of each patch. The reconstruction weights are then transferred to the high-resolution manifold, where the high-resolution patch is rendered using collocated patches of the high-resolution images in the training set  $\{\bar{H}\}$ . A preliminary high-resolution image  $\tilde{Y}'$  is then formed by stitching the reconstructed high-resolution patches. Finally, a reprojection operation is applied to further reduce artifacts and make the output high-resolution image  $\tilde{Y}$  more similar to the input low-resolution image. The methods is described in more detail in the following sub-sections.

#### A. EIGEN-PATCHES

Without loss of generality, suppose that our image recognition problem is iris recognition. Given an input low-resolution iris image  $\bar{X}$ , it is first separated into  $N = N_v \times N_h$  overlapping patches  $\{\bar{x}\} = \{\bar{x}_1, \bar{x}_2, \dots, \bar{x}_N\}$ , where  $N_v$  and  $N_h$  are the vertical and horizontal number of patches, respectively. Since we will consider square images in our experiments, we can assume from the remainder of this paper that  $N_v = N_h$ .

Two super sets of basis patches are computed for each patch position  $i$  from collocated patches of a training database of high resolution images  $\{\bar{H}\}$ . One of the super sets is obtained from collocated high resolution patches as follows. For each patch position  $i$ , we stack patches at the same position from the set of  $M$  high-resolution training iris images as shown in Figure 2, which we denote as  $\bar{H}_i = \{\bar{h}_i^1, \bar{h}_i^2, \dots, \bar{h}_i^M\}$ . By degradation (low-pass filtering and downsampling) using the acquisition model defined in Equation 2, a low-resolution database  $\{\bar{L}\}$  is generated from  $\{\bar{H}\}$ , and the corresponding low-resolution super set  $\bar{L}_i = \{\bar{l}_i^1, \bar{l}_i^2, \dots, \bar{l}_i^M\}$  is obtained for each patch position  $i$ . This way, the structure of the iris image is exploited by the construction of position-dependent dictionaries. In the example in Figure 2, the high-resolution dictionary  $\bar{H}_1$  (marked in red) is composed using the top-left position patches of all  $M$  high-resolution images, while the corresponding low-resolution dictionary  $\bar{L}_1$  is composed using the corresponding top-left patches of the  $M$  low-resolution images.

During testing, a PCA Eigen-transformation is conducted for each low-resolution input patch  $\bar{x}_i$  using the collocated

patches of the low-resolution dictionary  $\bar{L}_i$  to compute the optimal linear reconstruction weights  $\bar{c}_i = \{c_i^1, c_i^2, \dots, c_i^M\}$ . More specifically, we first compute the mean-patch of the  $i$ -th low-resolution dictionary using

$$\bar{m}_i^L = \frac{1}{M} \sum_{j=1}^M \bar{l}_i^j \tag{4}$$

so that the low-resolution dictionary  $\bar{L}_i$  for the  $i$ -th patch is centered by removing the mean-patch  $\bar{m}_i^L$ :

$$\bar{L}'_i = \bar{L}_i - \bar{m}_i^L \tag{5}$$

For an input patch  $\bar{x}_i$ , the weight vector can be computed by projecting it onto the eigen-space using

$$\bar{w}_i = E_i^T (\bar{x}_i - \bar{m}_i^L) \tag{6}$$

where the eigen-patches  $E_i$  are derived by applying PCA to the covariance matrix  $\bar{C}_i$  of the centered dictionary  $\bar{L}'_i$ :

$$\bar{C}_i = \bar{L}'_i{}^T \bar{L}'_i \tag{7}$$

The matrix  $E_i$  can be decomposed using

$$E_i = \bar{L}'_i V_i \Lambda_i^{-\frac{1}{2}} \tag{8}$$

where  $V_i$  and  $\Lambda_i$  are the eigenvectors and eigenvalues provided using PCA. In this work, we retain the eigenvectors of  $\bar{C}_i$  which contains at least 99% of the variance. The reconstruction weights are then derived using

$$\bar{c}_i = V_i \Lambda_i^{-\frac{1}{2}} \bar{w}_i \tag{9}$$

The  $i$ -th high-resolution patch  $\tilde{y}_i$  is then reconstructed from the collocated patches of the high-resolution dictionary  $\bar{H}_i$  using

$$\tilde{y}_i = \sum_{j=1}^M c_i^j \bar{h}_i^j + \bar{m}_i^H \tag{10}$$

where  $\bar{m}_i^H$  is the high-resolution counterpart of Equation 4:

$$\bar{m}_i^H = \frac{1}{M} \sum_{j=1}^M \bar{h}_i^j \tag{11}$$

The recovered patches are then stitched together by averaging overlapping pixels to synthesize the preliminary high-resolution iris  $\tilde{Y}'$ . It is important to mention here that every patch, which represents a particular spatial region within the iris, is optimized using dictionaries of irises at the same position. Therefore, the different reconstruction weights are optimized for every region of the iris. This ensures that iris images of higher quality, which are locally optimized, are reconstructed.

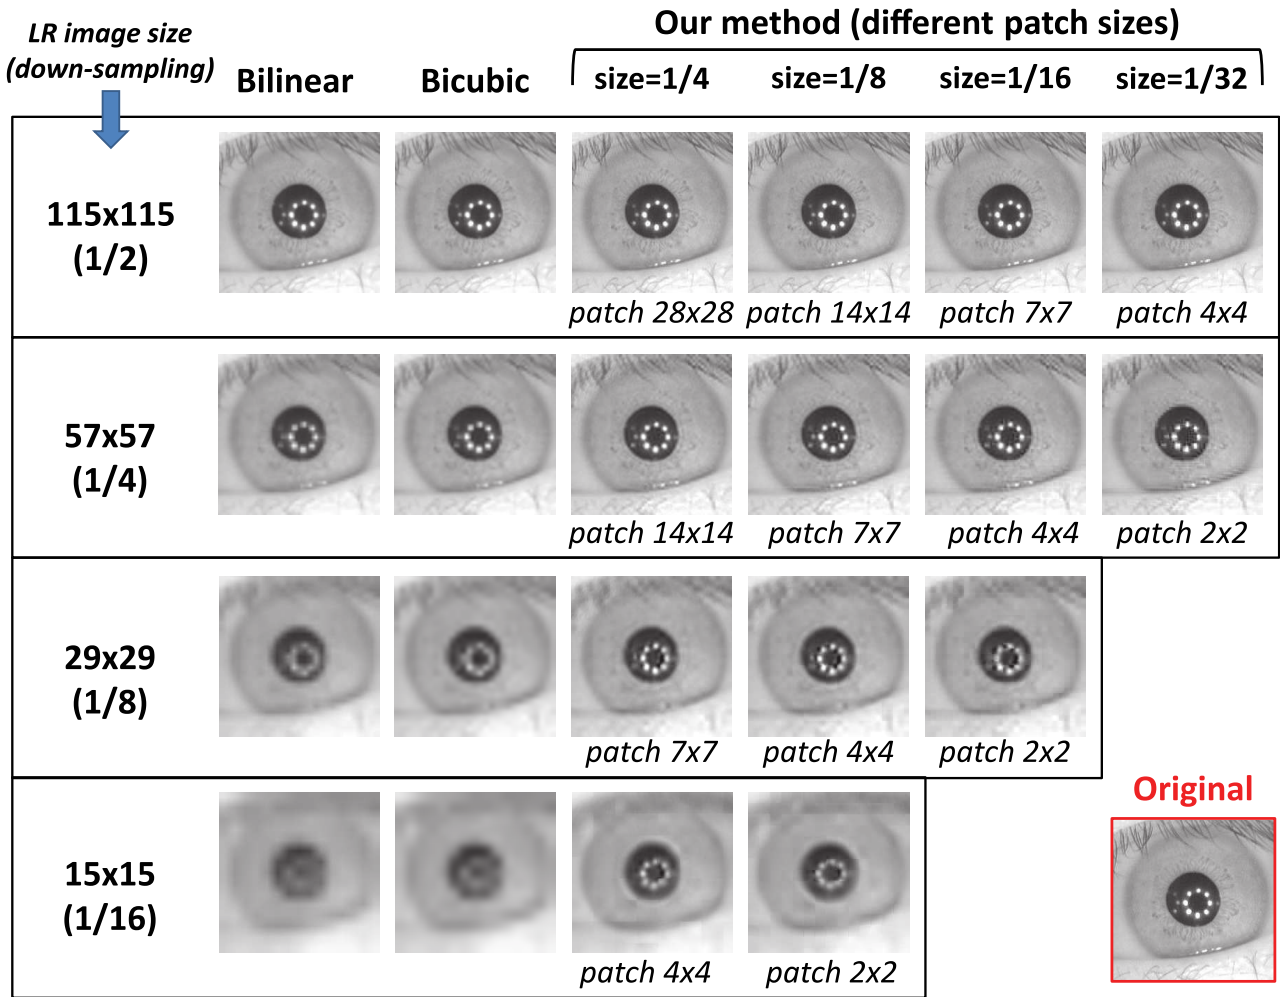


FIGURE 6. Resulting high-resolution hallucinated images for different downsampling factors and patch sizes.

**B. IMAGE REPROJECTION**

A re-projection step is further applied to  $\tilde{Y}^t$  to reduce artifacts and make the output image  $\tilde{Y}$  more similar to the input image  $\bar{X}$ . The image  $\tilde{Y}^t$  is re-projected to  $\bar{X}$  using the model of Equation 2 via:

$$\tilde{Y}^{t+1} = \tilde{Y}^t - \tau U \left( B \left( DB\tilde{Y}^t - \bar{X} \right) \right) \quad (12)$$

where  $U$  is the upsampling matrix. The process stops when  $|\tilde{Y}^{t+1} - \tilde{Y}^t|$  is smaller than a threshold. For our experiments in iris biometrics, we use  $\tau = 0.02$  and  $10^{-5}$  as the difference threshold.

**IV. EVALUATION FRAMEWORK**

**A. DATASET**

For our experiments, we used the CASIA Interval v3 iris database [54]. It consists of 2,655 NIR images from 249 contributors, captured in 2 sessions (the number of images per contributor and per session is not constant). A close-up iris camera was used to capture the images, with a resolution

of  $280 \times 320$  pixels. Manual annotation of the database is available [73], [74]. All images were resized via bicubic interpolation to have the same sclera radius ( $R = 105$ , average sclera radius of the whole database according to the ground-truth). Then, images were aligned by extracting a square region of  $231 \times 231$  pixels around the pupil center, which corresponds to about  $1.1 \times R$ . In case that such extraction was not possible (for example if the eye is close to an image side), the image was discarded. After this procedure, 1,872 images remained, which were used for our experiments. The dataset of aligned images was further divided into two sets: a training set comprised of images from the first 116 contributors ( $M = 925$  images), used as dictionary images to train the eigen-patch hallucination method, and a test set comprised of the remaining 133 contributors (947 images), which was used for validation.

The test and dictionary images were down-sampled via bicubic interpolation using MATLAB's `imresize` function by  $1/n$ , with  $n \in \{2, 4, 6, 8, 10, 12, 14, 16\}$ . This resulted in down-sampled images of  $115 \times 115$ ,  $57 \times 57$ ,  $39 \times 39$ ,  $29 \times 29$ ,

$23 \times 23$ ,  $19 \times 19$ ,  $17 \times 17$ , and  $15 \times 15$  pixels respectively. The dictionaries for each patch and for each downsampling factor were constructed using the position-patch method described in Section III-A. Down-sampled test images were then used as input low-resolution images, from which hallucinated high-resolution images were computed using the proposed algorithm. This simulated downsampling is the approach followed in most previous studies [14], due to the lack of databases with low-resolution and corresponding high-resolution reference images. The proposed method was compared with bilinear and bicubic interpolation. The method were implemented in MATLAB. All simulations were run using a machine with Intel (R) Core (TM) i7-4600U CPU at 2.10GHz running Windows 64-bit Operating system.

### B. IRIS AUTHENTICATION EXPERIMENTS

We conducted verification and identification experiments in the test set with several iris recognition algorithms (Section IV-C). We considered two scenarios, shown in Figure 7: scenario 1), where enrolment samples were taken from original high-resolution images of the test set, and query samples from hallucinated images; and scenario 2), where both enrolment and query samples were hallucinated images. The first case simulates a controlled enrolment scenario, where one of the samples is of high-resolution. The second case, on the other hand, simulates a totally uncontrolled scenario, where both samples are of low-resolution (albeit for simplicity, both have similar resolution).




FIGURE 7. Scenarios considered in our identity recognition experiments.

In our authentication experiments, each eye were considered as a different user of the system. *Verification experiments* were done as follows. Genuine trials were obtained by comparing each image of a user to the remaining images of the same user, avoiding symmetric comparisons. Impostor trials were obtained by comparing the 1<sup>st</sup> image of a user to the 2<sup>nd</sup> image of the remaining users. With this procedure, we obtained 2,607 genuine and 19,537 impostor scores per comparator and per scenario. To carry out *identification experiments*, we used the first image of each user as enrolment sample, and the remaining images for evaluation. Given an evaluation sample, the user was recognized by searching the enrolment samples of all the  $K$  subjects in the database for a match (one-to-many). As a result, the system returned a ranked list of candidates. For identification experiments, only users (eyes) with two or more samples were considered. This resulted in  $K = 162$  available users, and 764 different one-to-many trials (totaling 124,532 comparisons) per comparator and per scenario.

### C. IRIS RECOGNITION ALGORITHMS

Iris recognition experiments were conducted using six different algorithms according to the state of the art [75], namely 1D log-Gabor filters (LG) [19], the SIFT operator (SIFT) [20], local intensity variations in iris textures (CR) [15], the Discrete-Cosine Transform (DCT) [16], cumulative-sum-based grey change analysis (KO) [17], and Gabor spatial filters (QSW) [18].

All algorithms (except SIFT) extract features from a normalized rectangle image which is computed from the iris image using the Daugman's rubber sheet model [53]. This normalization produces a 2D polar array of  $20 \times 240$  pixels, height  $\times$  width, (LG system) and  $64 \times 512$  pixels (others), with horizontal dimensions of angular resolution, and vertical dimensions of radial resolution. Feature encoding is implemented according to the different feature extraction methods, leading to fixed-length templates with are matched using distance measures (details are given in the respective papers). Rotation is accounted for by shifting the 2D polar array of the query image in counter- and clock-wise direction and selecting the lowest distance, which corresponds to the best match between the two templates. In the SIFT comparator, SIFT key points are directly extracted from the iris image (without normalization), and the recognition metric is the number of matched key points, normalized by the average number of detected key-points in the two images under comparison.

We used open source code implementations of these algorithms. The LG implementation is from the Libor Masek code [19]. SIFT feature extraction and matching was carried out using a free toolkit<sup>1</sup> with the adaptations described in [50] (particularly, it includes a post-processing step to remove spurious matching points using geometric constraints). The remaining algorithms used are from the University of Salzburg Iris Toolkit software package (USIT) [75].

## V. EXPERIMENTAL RESULTS


### A. IMAGE FIDELITY

This section reports the performance of the hallucination algorithms by measuring the Peak Signal-to-Noise ratio (PSNR) and the Structural Similarity (SSIM) index between the hallucinated image and its corresponding high-resolution reference image of the test set. We used MATLAB's `psnr` and `ssim` functions for this purpose. These are the metrics usually employed in the super-resolution literature [6]. The PSNR is a measure of the ratio (in dBs) between the maximum possible power of a signal and the power of corrupting noise that affects the fidelity of its representation. The noise in this case is defined as the difference between the reference high-resolution image  $\bar{Y}$ , and its estimation  $\tilde{Y}$  by the reconstruction algorithm. A higher PSNR generally indicates that the reconstruction is of higher quality. If the two images are identical,  $\bar{Y} = \tilde{Y}$ , then  $\text{PSNR}(\bar{Y}, \tilde{Y}) = \infty$ . The SSIM index, on the other hand, is a perception-based model that considers image degradation as

<sup>1</sup><http://vision.ucla.edu/~vedaldi/code/sift/assets/sift/index.html>

**TABLE 2.** Hallucination results (PSNR and SSIM values) with different downsampling factors and patch sizes (average values on the test dataset). Patch size is indicated in proportion to the size of the iris image. The best results for each downsampling factor are marked in bold. ‘Diff’ refers to the difference between the best PCA case and the bicubic method.

LR size (scaling)		Full image							Unwrapped iris region						
		bilinear	bicubic	Our method (patch size)				diff	bilinear	bicubic	Our method (patch size)				diff
				1/4	1/8	1/16	1/32				1/4	1/8	1/16	1/32	
115×115 (1/2)	psnr	33	34.04	34.23	<b>34.65</b>	34.62	34.11	+0.61	36.94	38.22	38.77	<b>39.15</b>	39.1	38.59	+0.93
	ssim	0.91	0.93	0.92	0.93	<b>0.94</b>	0.93	+0.01	0.96	<b>0.97</b>	<b>0.97</b>	<b>0.97</b>	<b>0.97</b>	<b>0.97</b>	+0.00
57×57 (1/4)	psnr	28.36	29.18	<b>29.91</b>	29.9	29.53	28.78	+0.73	31.64	32.35	32.69	<b>32.72</b>	32.38	31.83	+0.37
	ssim	0.79	0.8	0.8	<b>0.81</b>	0.8	0.78	+0.01	0.85	0.87	<b>0.88</b>	<b>0.88</b>	<b>0.88</b>	0.86	+0.01
39×39 (1/6)	psnr	26.21	26.8	27.98	<b>28.05</b>	27.86	27.61	+1.25	29.61	30.21	30.71	<b>30.85</b>	30.73	30.64	+0.64
	ssim	0.73	0.74	0.74	<b>0.75</b>	<b>0.75</b>	<b>0.75</b>	+0.01	0.79	0.81	0.81	<b>0.82</b>	<b>0.82</b>	<b>0.82</b>	+0.01
29×29 (1/8)	psnr	24.86	25.33	<b>26.73</b>	26.55	26.16	-	+1.4	28.18	28.74	<b>29.57</b>	29.49	29.2	-	+0.83
	ssim	0.69	0.7	<b>0.71</b>	<b>0.71</b>	0.7	-	+0.01	0.74	0.75	<b>0.77</b>	<b>0.77</b>	0.76	-	+0.02
23×23 (1/10)	psnr	23.94	24.41	<b>25.76</b>	25.62	25.24	-	+1.35	27.09	27.63	<b>28.69</b>	28.47	28.11	-	+1.06
	ssim	0.67	0.68	<b>0.69</b>	<b>0.69</b>	0.67	-	+0.01	0.7	0.71	<b>0.73</b>	0.72	0.71	-	+0.02
19×19 (1/12)	psnr	23.24	23.71	<b>25.12</b>	24.95	24.57	-	+1.41	26.21	26.71	<b>27.88</b>	27.53	27.18	-	+1.17
	ssim	0.65	0.66	<b>0.68</b>	0.67	0.65	-	+0.02	0.66	0.68	<b>0.69</b>	0.68	0.67	-	+0.01
17×17 (1/14)	psnr	22.85	23.32	<b>24.69</b>	24.05	-	-	+1.37	25.72	26.22	<b>27.39</b>	26.53	-	-	+1.17
	ssim	0.64	0.65	<b>0.66</b>	0.64	-	-	+0.01	0.65	0.66	<b>0.67</b>	0.63	-	-	+0.01
15×15 (1/16)	psnr	22.39	22.86	<b>24.32</b>	24.17	-	-	+1.46	25.16	25.64	<b>26.99</b>	26.79	-	-	+1.35
	ssim	0.64	0.64	<b>0.66</b>	0.65	-	-	+0.02	0.63	0.64	<b>0.66</b>	0.65	-	-	+0.02



**FIGURE 8.** Hallucination results with different downsampling factors (left: PSNR values, right: SSIM values). The best PCA case for each downsampling factor is shown (marked in bold in Table 2).

a perceived change in structural information (luminance and contrast). This is achieved by using first and second order statistics of grey values on local image windows. The SSIM index is a decimal value between  $-1$  and  $1$ , and value  $1$  is only reachable in the case of two identical images.

Table 2 reports the average PSNR and SSIM values of all images in the test set. We report PSNR and SSIM metrics in two cases: *i*) between the full reconstructed image and its reference high-resolution image; and *ii*) between the normalized polar image versions of these two images (with size  $20 \times 240$  pixels), computed according to the Daugman’s rubber sheet model [53]. Examples of hallucinated iris images can also be seen in Figure 6 (only for a selection of downsampling factors for the sake of space). In the described PCA method, the size of local image patches is an important parameter. In order to test the influence of this parameter, we evaluated the performance of the PCA algorithm with different patch sizes. In particular, we tested patches of size equal to  $1/4$ ,  $1/8$ ,  $1/16$  and  $1/32$  of the iris image. We defined the patch size in proportion to the dimensions of the iris image

to ensure that they cover the same relative size across different scaling factors. Overlapping between patches was set to  $1/3$ .

Table 2 shows in bold the best result for each downsampling factor. As it can be seen, the eigen-patch hallucination method outperforms the bilinear and bicubic interpolations at all resolutions. Also, among bilinear and bicubic interpolations, the latter gives the best results. We can also observe that, as resolution decreases, the gain of the PCA method becomes larger, at least as measured by the PSNR. This can be better assessed in Figure 8, where we plot the results of Table 2 (the best PCA case for each downsampling factor is selected). At a resolution of  $115 \times 115$  pixels, the PSNR gain of PCA is of  $0.61\text{dB}$  over bicubic (full image) and  $0.93\text{dB}$  (iris region), and it becomes larger at lower resolutions. At a resolution of  $15 \times 15$  pixels, the gain reaches  $1.46\text{dB}$  (full image) and  $1.35\text{dB}$  (iris region), showing the advantage of the utilized eigen-patch hallucination method at very low resolutions.

It is also worth noting that, although the average SSIM values are smaller as resolution decreases, the gain of PCA



**TABLE 3.** Verification results (EER values) of the two scenarios considered for different downsampling factors. The best case for each comparator and for each downsampling factor is marked in bold. The relative EER variation of PCA with respect to bilinear/bicubic (best of the two) is also given.

Comparator	Method	Scenario 1 (original vs. down-sampled)				Scenario 2 (down-sampled vs. down-sampled)			
		Downsampling				Downsampling			
		115×115 (1/2)	57×57 (1/4)	29×29 (1/8)	15×15 (1/16)	115×115 (1/2)	57×57 (1/4)	29×29 (1/8)	15×15 (1/16)
LG	Bilinear	<b>0.69%</b>	0.69%	1.61%	10.39%	<b>0.61%</b>	0.76%	2.38%	11.03%
	Bicubic	<b>0.69%</b>	<b>0.68%</b>	1.42%	9.59%	0.73%	<b>0.65%</b>	1.88%	11.25%
	Our (var)	0.76% (+10.1%)	0.8% (+17.6%)	<b>1.11%</b> (-21.8%)	<b>7.29%</b> (-24%)	0.73% (+19.7%)	0.69% (+6.2%)	<b>1.18%</b> (-37.2%)	<b>4.79%</b> (-56.6%)
SIFT	Bilinear	4.05%	10.42%	28.23%	50.52%	<b>3.01%</b>	4.26%	14.82%	41.66%
	Bicubic	<b>3.51%</b>	7.41%	24.99%	47.33%	3.13%	<b>3.08%</b>	11.6%	36.37%
	Our (var)	4.17% (+18.8%)	<b>4.74%</b> (-36%)	<b>15.65%</b> (-37.4%)	<b>36.67%</b> (-22.5%)	3.9% (+29.6%)	3.11% (+1%)	<b>7.46%</b> (-35.7%)	<b>18.98%</b> (-47.8%)
CR	Bilinear	10.12%	13.06%	20.96%	33.74%	8.86%	9.24%	13%	15.76%
	Bicubic	<b>9.85%</b>	11.98%	19.45%	34.1%	<b>8.84%</b>	<b>7.83%</b>	10.98%	17.49%
	Our (var)	10.41% (+5.7%)	<b>11.33%</b> (-5.4%)	<b>16.88%</b> (-13.2%)	<b>28.1%</b> (-16.7%)	10.1% (+14.3%)	9.04% (+15.5%)	<b>9.98%</b> (-9.1%)	<b>14.44%</b> (-8.4%)
DCT	Bilinear	<b>2.15%</b>	3.59%	16.64%	44.51%	<b>1.77%</b>	2.55%	9.25%	17.77%
	Bicubic	2.16%	2.95%	14.05%	44.03%	2.06%	2.14%	7.63%	18.11%
	Our (var)	2.19% (+1.9%)	<b>2.73%</b> (-7.5%)	<b>10.39%</b> (-26%)	<b>39.09%</b> (-11.2%)	2.11% (+19.2%)	<b>2.13%</b> (-0.5%)	<b>6.75%</b> (-11.5%)	<b>11.72%</b> (-34%)
KO	Bilinear	<b>12.18%</b>	12.68%	14.39%	18.41%	12.48%	12.92%	14.08%	15%
	Bicubic	12.55%	12.74%	13.95%	17.49%	<b>12.37%</b>	12.49%	13.75%	14.92%
	Our (var)	12.49% (+2.5%)	<b>12.45%</b> (-1.8%)	<b>13.44%</b> (-3.7%)	<b>16.63%</b> (-4.9%)	12.54% (+1.4%)	<b>12.29%</b> (-1.6%)	<b>12.8%</b> (-6.9%)	<b>13.41%</b> (-10.1%)
QSW	Bilinear	<b>0.7%</b>	0.76%	2.09%	14.47%	<b>0.76%</b>	0.84%	3.06%	12.49%
	Bicubic	0.74%	0.77%	1.8%	<b>12.13%</b>	0.8%	0.8%	2.93%	12.95%
	Our (var)	0.76% (+8.6%)	<b>0.75%</b> (-1.3%)	<b>1.67%</b> (-7.2%)	12.31% (+1.5%)	<b>0.76%</b> (0%)	<b>0.77%</b> (-3.8%)	<b>1.91%</b> (-34.8%)	<b>5.55%</b> (-55.6%)

over bicubic remains between 0.01 and 0.02. In Figure 6, it can be observed that images reconstructed with bilinear or bicubic methods have much more blur than those reconstructed with PCA when the resolution decreases. While this is reflected in a higher gain in PSNR, it is not the case with the SSIM. As it has been pointed out in previous studies, image quality metrics like PSNR or SSIM do not have the same sensitivity to image degradations [21], [22]. Although iris images reconstructed with PCA have better subjective quality, this is not captured well by the SSIM index.

Regarding the appropriate patch size, it can be observed in Table 2 that the best results are obtained with a bigger patch. In particular, 1/4 is the best size at very low-resolutions. Using a bigger patch results in less artifacts due to overlapping patches being stitched together. This can be seen for example in Figure 6 (image of size 57×57 pixels), where more artifacts appear for patch sizes of 1/16 or 1/32. It also has computational implications, because there are less patches to process per image, although they have a bigger size.

Given that a patch size of 1/4 consistently produces top results in the experiments of this section, we will employ this patch size with PCA in the remainder of this paper. This patch size has also been observed to produce top results in previously published verification experiments using the LG comparator [11].

## B. BIOMETRIC VERIFICATION

Next, we report results of the verification experiments using hallucinated iris images. Verification results of the proposed PCA method and the bilinear/bicubic interpolations are given

in Table 3 and Figure 9. We adopt the Equal Error Rate (EER) as the measure of accuracy. Due to space constraints, we only report a selection of downsampling factors (those shown in Figure 6). In Table 3, we also mark in bold the best EER value for each comparator and for each downsampling factor. In addition, the relative EER variation of PCA with respect to bilinear/bicubic (best of the two) is given in brackets.

### 1) INTERPOLATION METHOD ANALYSIS

It can be seen that the PCA method results in better verification performance than bilinear/bicubic interpolations when the resolution decreases. This highlights the benefits of using trained methods to enhance very low-resolution iris images. At an image size of 29×29 pixels (downsampling of 1/8), any given comparator produces better performance with PCA than with bilinear or bicubic interpolations. In scenario 1, for example, the EER reduction of PCA is between 3.7% (KO comparator) and 37.4% (SIFT); whereas in scenario 2, it is between 6.9% (KO) and 37.2% (LG). For some comparators, the improvement with PCA is evident even at a resolution of 57×57 pixels (downsampling of 1/4). For example, the EER reduction in this case is of 36% with the SIFT comparator, and of 7.5% with the DCT system (both in scenario 1). But the biggest benefits of PCA occur at very low-resolutions (15×15 pixels, or downsampling of 1/16). Here, EER improvements are significantly higher, as reflected by the wider gap between curves in Figure 9. With such small image size, an impressive EER reduction of 56-57% is achieved with LG and QSW, followed by 47.8% with SIFT, and 34% with DCT (all in scenario 2).

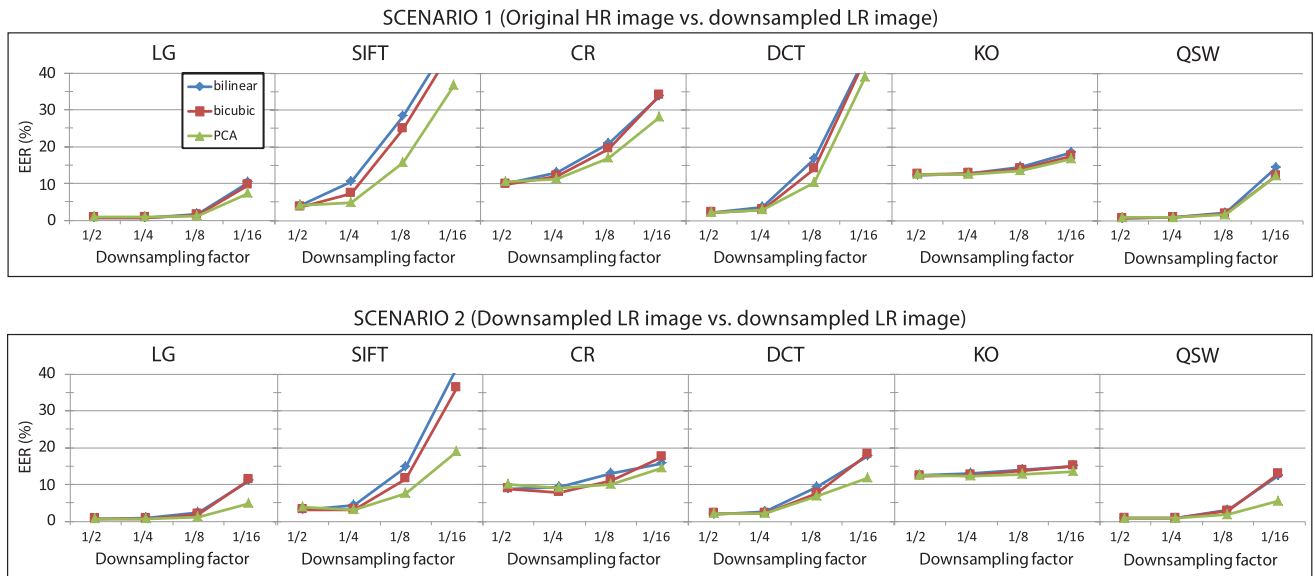


FIGURE 9. Verification results (EER values) with different downsampling factors.

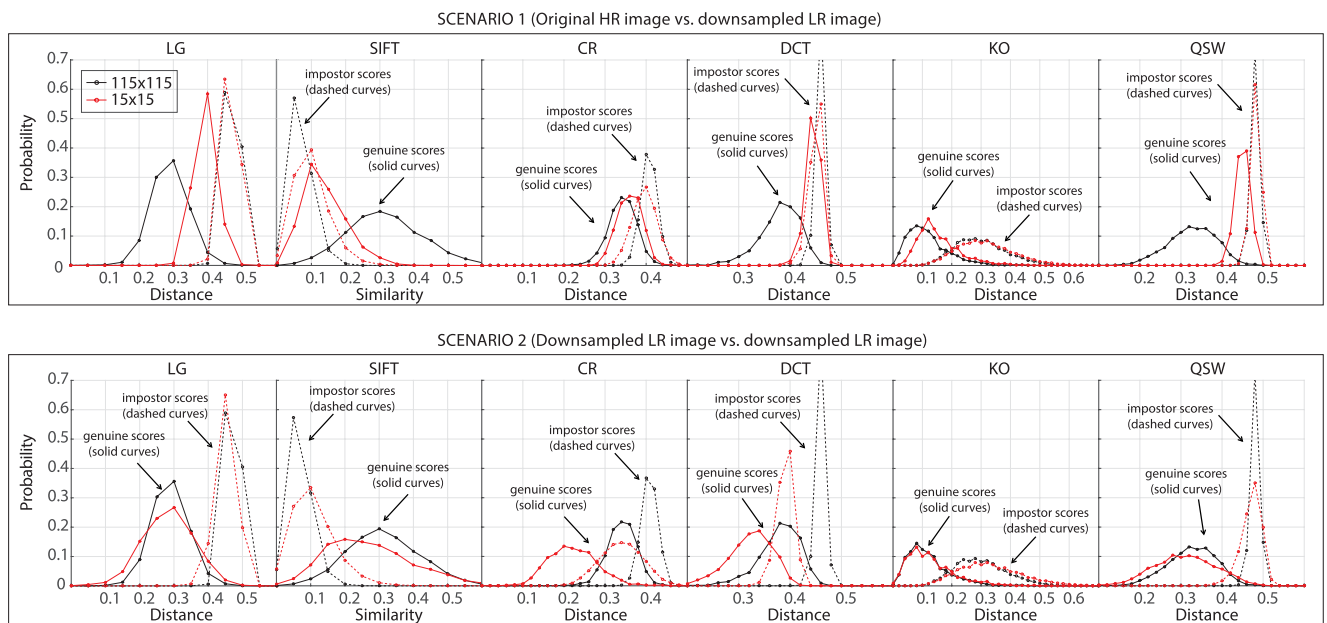



FIGURE 10. Score distributions with PCA enhancement for a downsampling factor of 1/2 (image size 115×115 pixels, black curves) and 1/16 (image size 15×15 pixels, red curves).

## 2) SCENARIO ANALYSIS

In this sub-section, we study the difference between the two scenarios of operation analyzed. It can be observed in Figure 9 that the performance of scenario 1 and 2 is approximately equal up to a certain resolution. Then, if resolution is further reduced, scenario 2 performs much better than scenario 1. This is observed with all comparators, being the only difference the cut-off point. For example, with LG, KO, or QSW, the difference between both scenarios only becomes substantial at a factor of 1/16; but with DCT the difference starts at 1/8, and even earlier with SIFT and CR.

We further analyse this effect by plotting in Figure 10 the score distributions with PCA of the following two cases: moderate downsampling (factor 1/2, image size 115×115 pixels), and extreme downsampling (factor 1/16, image size 15×15 pixels). Note that most comparators employ a distance measure, meaning that genuine distributions are on the left side of the plot, but one comparator (SIFT) employ a similarity measure, so its genuine distributions are on the right side.

In scenario 1, we observe that in most comparators, the distribution of genuine scores is shifted towards the impostor



**FIGURE 11.** Verification results (DET curves) of the iris comparators. Results are given for scenario 2 and a downsampling factor of 1/16 (image size 15×15 pixels).

distribution as resolution decreases. This suggests that the PCA reconstruction algorithm is not able to fully recover the information found in the original high resolution image, at least measured by the features employed by our comparators. As a result, the similarity between high-resolution and reconstructed images of the same user is reduced (recall that in scenario 1, enrolment samples were taken from high-resolution images, whereas query samples were taken from reconstructed images). This is consistent with the results of Figure 8, which show that the PSNR and SSIM values between a high-resolution image and its reconstructed counterpart decreases with the resolution. On the other hand, impostor distributions lie in a similar range in most cases, regardless of the resolution.

In scenario 2, on the contrary, the distribution of genuine scores is not shifted significantly towards the impostor distribution as the resolution changes. Instead, the relative difference between the two distributions is maintained. In this scenario, enrolment and query images undergo the same downsampling and upsampling procedure, regardless whether they come from genuine or impostor trials. Therefore, it has sense that the relative difference between genuine users and impostors is maintained. However, a collateral effect, given by the loss of information when images are down-sampled and reconstructed, is that genuine and score distributions are more spread. This explains the worse EER performance at low resolution. An exception to these observations is SIFT. With this comparator, the genuine and impostor distributions become significantly closer to each other in scenario 2. For this reason, this is the comparator whose performance is most significantly degraded with the resolution (see Figure 9). Another exception, but in the opposite direction, is the KO comparator. This comparator shows a high resilience to changes in resolution, as it can be seen by its nearly ‘flat’ behavior in Figure 9. As a result, its score distributions are in the same range. It should be remarked, however, that its EER at high-resolution is already above 12%, while other comparators start below 1%.

### 3) COMPARATOR ANALYSIS

We further look into the differences between the individual comparators considered in this paper. As it can be seen in

Figure 9, the performance of any comparator is not degraded significantly until a downsampling factor of at least 1/8 (image size of 29×29 pixels). This implies that the size of both gallery and probe images could be kept low without sacrificing performance, and then simply up-sampled with the bicubic or even bilinear method. This has positive implications when lower storage or data transmission capabilities are required. Some comparators (LG, QSW) even show a praiseworthy performance with a downsampling factor of 1/8, having in this case an EER of 1.11%/1.67% in scenario 1, and 1.18%/1.91% in scenario 2. It should be noted as well that these two comparators are based on Gabor wavelets, so it is reasonable that they behave in a similar fashion. It is also remarkable the impressive low EER figures obtained with these comparators when considering very low-resolution images with a size of only 15×15 pixels. For example, an EER of 4.78% is obtained with the LG comparator, and 5.55% with QSW in scenario 2. Also, as mentioned above, the KO comparator shows a high resilience to changes in resolution, with the EER degrading only from 12.54% to 13.41% in scenario 2.

In terms of absolute performance, the best comparator at any resolution is LG, followed closely by QSW (recall that both are based on Gabor wavelets). The third one is DCT, followed by SIFT, although performance of these two degrades significantly at very low-resolutions. For example, DCT has an EER of ~2% at a resolution of 1/2, but it goes up to 39.09% at a resolution of 1/16 in scenario 1, and to 11.72% in scenario 2. CR and KO are the worst performing comparators. They both start with an EER of 10% or higher already at a resolution of 1/2, although they show higher resiliency to reductions in resolution, as seen earlier.

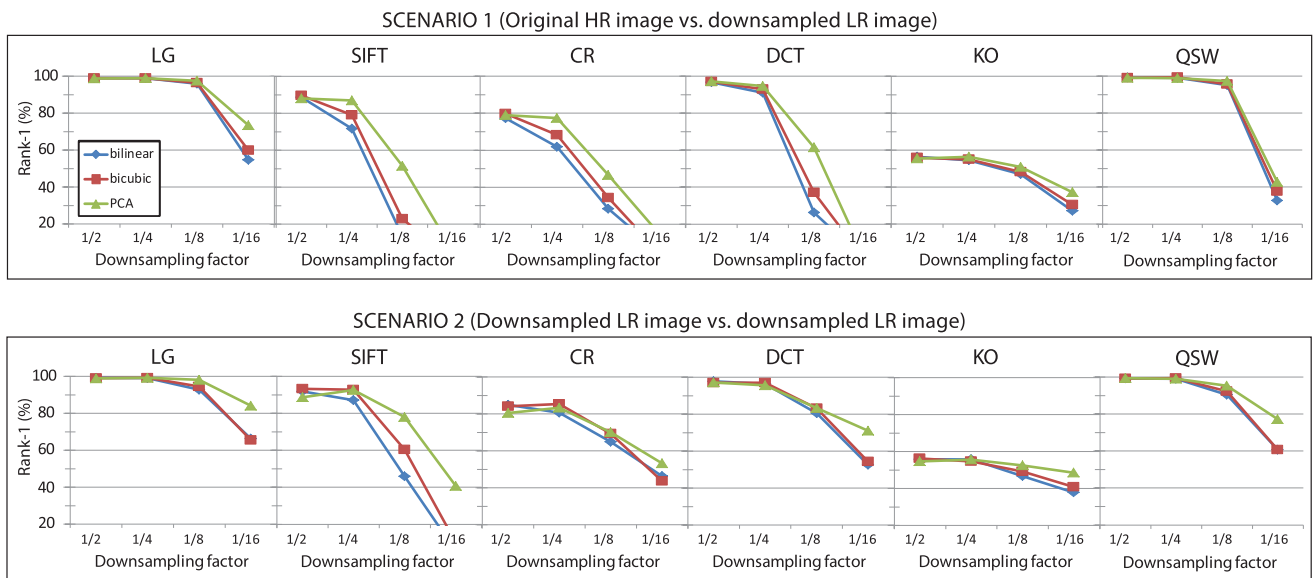
For completeness, we provide in Figure 11 the DET curves of all comparators (only scenario 2 and a downsampling factor of 1/16). Here, it can be observed the gain of the trained PCA enhancement in comparison to bilinear and bicubic interpolations. A consistent improvement across all ranges of the DET curve is obtained with most comparators.

### C. BIOMETRIC IDENTIFICATION

In this section, we report identification experiments using hallucinated iris images. Results are given in

**TABLE 4. Identification results (Top-1 accuracy) of the two scenarios considered for different downsampling factors. The best case for each comparator and for each downsampling factor is marked in bold. The relative Top-1 variation of PCA with respect to bilinear/bicubic (best of the two) is also given.**

Comparator	Method	Scenario 1 (original vs. down-sampled)				Scenario 2 (down-sampled vs. down-sampled)			
		Downsampling				Downsampling			
		115×115 (1/2)	57×57 (1/4)	29×29 (1/8)	15×15 (1/16)	115×115 (1/2)	57×57 (1/4)	29×29 (1/8)	15×15 (1/16)
LG	Bilinear	<b>98.82%</b>	<b>98.95%</b>	95.94%	54.58%	<b>99.21%</b>	99.08%	92.8%	66.23%
	Bicubic	<b>98.82%</b>	<b>98.95%</b>	96.34%	59.82%	99.08%	<b>99.21%</b>	94.5%	65.45%
	Our	98.69%	<b>98.95%</b>	<b>97.51%</b>	<b>73.43%</b>	98.82%	<b>99.21%</b>	<b>98.04%</b>	<b>84.16%</b>
	(var)	(-0.1%)	(0%)	(+1.2%)	(+22.8%)	(-0.4%)	(0%)	(+3.7%)	(+27.1%)
SIFT	Bilinear	88.74%	71.73%	13.48%	0.92%	91.88%	87.17%	45.94%	5.5%
	Bicubic	<b>89.53%</b>	79.06%	22.91%	0.65%	<b>93.19%</b>	<b>92.8%</b>	60.34%	9.95%
	Our	88.09%	<b>86.91%</b>	<b>51.44%</b>	<b>5.37%</b>	88.74%	92.54%	<b>78.01%</b>	<b>40.71%</b>
	(var)	(-1.6%)	(+9.9%)	(+124.5%)	(+483.7%)	(-4.8%)	(-0.3%)	(+29.3%)	(+309.1%)
CR	Bilinear	77.23%	61.78%	28.27%	4.58%	<b>84.82%</b>	80.63%	64.92%	46.34%
	Bicubic	<b>79.71%</b>	68.19%	34.29%	4.32%	84.16%	<b>85.34%</b>	69.24%	43.59%
	Our	78.93%	<b>77.36%</b>	<b>46.6%</b>	<b>14.53%</b>	80.5%	83.38%	<b>70.03%</b>	<b>53.27%</b>
	(var)	(-1%)	(+13.4%)	(+35.9%)	(+217.2%)	(-5.1%)	(-2.3%)	(+1.1%)	(+15%)
DCT	Bilinear	96.73%	90.97%	26.18%	0.39%	<b>97.77%</b>	96.34%	80.5%	52.62%
	Bicubic	96.99%	92.93%	37.04%	0.39%	96.99%	<b>96.99%</b>	83.12%	54.32%
	Our	<b>97.12%</b>	<b>94.63%</b>	<b>61.52%</b>	<b>0.92%</b>	97.12%	95.55%	<b>83.38%</b>	<b>71.07%</b>
	(var)	(+0.1%)	(+1.8%)	(+66.1%)	(+135.9%)	(-0.7%)	(-1.5%)	(+3.3%)	(+30.8%)
KO	Bilinear	<b>56.41%</b>	54.45%	46.99%	27.09%	55.76%	<b>55.76%</b>	46.6%	37.83%
	Bicubic	55.89%	54.97%	48.17%	30.37%	<b>56.15%</b>	54.58%	49.21%	40.84%
	Our	55.5%	<b>56.41%</b>	<b>50.92%</b>	<b>37.17%</b>	54.71%	55.63%	<b>52.49%</b>	<b>48.56%</b>
	(var)	(-1.6%)	(+2.6%)	(+5.7%)	(+22.4%)	(-2.6%)	(-0.2%)	(+6.7%)	(+18.9%)
QSW	Bilinear	<b>99.21%</b>	99.08%	95.16%	32.59%	99.21%	99.08%	90.45%	60.73%
	Bicubic	99.08%	<b>99.35%</b>	95.55%	37.7%	99.21%	<b>99.35%</b>	92.28%	60.6%
	Our	99.08%	98.95%	<b>97.38%</b>	<b>42.8%</b>	<b>99.35%</b>	99.08%	<b>95.29%</b>	<b>77.23%</b>
	(var)	(-0.1%)	(-0.4%)	(+1.9%)	(+13.5%)	(+0.1%)	(-0.3%)	(+3.3%)	(+27.2%)




**FIGURE 12. Identification results (Top-1 accuracy) with different downsampling factors.**

Table 4 and Figure 12. Here, we adopt as metric the classification accuracy for a hit list size of  $k = 1$  candidate (also called Top-1 or Rank-1). In Table 4, we also mark in bold the best accuracy for each comparator and for each downsampling factor. In addition, the relative Top-1 accuracy variation of PCA with respect to bilinear/bicubic (best of the two) is given in brackets. Finally, we provide in Figure 13 the CMC curves of scenario 2 for a downsampling factor of 1/16.

1) INTERPOLATION METHOD ANALYSIS

Similarly to the verification experiments, we also observe here a superior performance of the trained PCA enhancement when the resolution decreases. At an image size of  $29 \times 29$  pixels (downsampling of 1/8), the PCA method produces better accuracy than bilinear or bicubic interpolations for any comparator. The Top-1 gain for this image size in scenario 1 reaches a value of 66.1% (DCT comparator) and 124.5% (SIFT), whereas in scenario 2, the gain reaches



**FIGURE 13.** Identification results (CMC curves) of the iris comparators. Results are given for scenario 2 and a downsampling factor of 1/16 (image size 15×15 pixels).

29.3% (SIFT). In scenario 1, some comparators show a remarkable improvement even at a resolution of 57×57 pixels. Here, CR has an improvement of 13.4%, whereas SIFT has an improvement of 9.9%. But the advantage of using the PCA trained method becomes more evident again at very low resolution (15×15 pixels), as it can be appreciated in the bigger separation between curves in Figure 12. For example, in scenario 2, a Top-1 improvement of 309.1% is observed with SIFT, followed by 30.8% with DCT, and ~27% with LG and QSW. In scenario 1, improvements of 483.7% and 217.2% can be observed with SIFT and CR, respectively, although the Top-1 accuracy of these comparators is far from being usable in practice (below 15%).

2) SCENARIO ANALYSIS

Regarding the two scenarios of operation employed, similar observations than in the verification experiments can be made here. The performance of a number of comparators (LG, KO, QSW, DCT) only differs between the two scenarios when the downsampling factor is of 1/8 or higher, while the others (SIFT, CR) show differences already at 1/4. Also, when the two scenarios differ, the best performance is always obtained in scenario 2. In scenario 1, only query samples are from low-resolution images, which are matched against high-resolution enrolment samples. As we observed earlier, the similarity between a high-resolution image and its reconstructed version decreases with the resolution, which explains the worse performance obtained in scenario 1.

It is also specially relevant the poor performance of some comparators in scenario 1 at a resolution of 1/16. For example, SIFT and DCT obtain a Top-1 accuracy of ~5% or less. These same comparators also show a high EER in verification mode (above 36%). This makes these two features unsuitable to compare images that have very different resolution, as it is the case in scenario 1. Indeed, the only comparator capable of handling such differences to some extent is LG. Its Top-1 identification accuracy in scenario 1 is of 73.43%, and it also shows the best verification accuracy among all comparators (EER of 7.29%).

3) COMPARATOR ANALYSIS

Finally, we analyse the differences between the individual comparators when they operate in identification mode.

An effect observed in the verification experiments that we also see here is that the performance of any comparator remains without significant degradation until a downsampling factor of at least 1/8. This corresponds to an image size of 29×29 pixels. Some comparators (like LG or QSW), are not significantly degraded even with this small image size, having a Top-1 accuracy of 95-98% in any of the two scenarios.

In terms of absolute performance, LG and QSW are, again, the best comparators, with a Top-1 accuracy higher than 98% across a great range of resolutions. They only degrade at very low-resolution (15×15 pixels), showing a noteworthy Top-1 accuracy of 84.16% (LG) and 77.23% (QSW) in scenario 2. The Top-1 accuracy of DCT at this very low-resolution is 71.07%, while accuracy of the other comparators fall below 53%, making them unfeasible for practical applications in this extreme case.

Regarding the CMC curves at very low-resolution (Figure 13), if we allow a hit list size of  $k = 5$  candidates, both LG and QSW show a classification accuracy higher than 90%, which goes up to 95% for a list of  $k = 10$  candidates. This shows that the proposed PCA approach can be effectively used to improve iris identification under severe downsampling. With bilinear or bicubic interpolations, the accuracy of these two comparators is below 78% (if  $k = 5$ ) and 82% (if  $k = 10$ ). Also, the DCT and CR comparators can reach 90% accuracy with PCA enhancement, but we need to increase the size of the hit list to  $k = 20$  or  $k = 25$  candidates, respectively. KO and SIFT, on the other hand, need a list size of  $k = 50$  candidates to reach 90% accuracy.

VI. CONCLUSIONS

Iris is regarded as one of the most accurate biometric modalities [76]. It provides very high accuracy in controlled environments, but deployment in non-controlled environments such as at-a-distance or on-the-move is not yet mature [1]. The use of more relaxed acquisition environments is pushing image-based biometrics towards the use of low-resolution imagery. This can pose significant problems in terms of reduced performance if not tackled properly. In this context, super-resolution techniques can be used to enhance the quality of low-resolution iris images and therefore, to improve the recognition performance of existing systems [51].

Super-resolution is a core topic in computer vision, with many techniques proposed to restore low-resolution images [6], [77]. However, compared with the existing literature in generic super-resolution, super-resolution in biometrics is a relatively recent topic [4]. This is because most approaches are general-scene, designed to produce an overall visual enhancement. They try to improve the quality of the image by minimizing an objective fidelity measure, such as the Peak Signal-to-Noise Ratio (PSNR), which does not necessarily correlate with better recognition performance [47]. Images from a specific biometric modality have particular local and global structures that can be exploited to achieve a more efficient upsampling [10]. For example, recovering local texture details is essential for iris images due to the prevalence of texture-based recognition in this modality [78].

In this paper, we present an extensive up-to-date survey of super-resolution applied to iris biometrics (Section II). We provide a comprehensive coverage of the existing literature, including a taxonomy of existing iris super-resolution approaches (Figure 3 and Table 1). They can be broadly classified into reconstruction-based and learning-based methods [5]. Reconstruction based methods register and fuse a sequence of low-resolution images by pixel-wise combination of intensity values, in order to estimate a high-resolution image. On the other hand, learning based methods use coupled dictionaries to learn the mapping relations between low- and high-resolution image pairs. The research community has lately focused on the latter category, since they can provide higher quality images and larger magnification factors [8]. They also have the advantage of needing only one image as input. In our survey, we also cover other aspects, from the database employed in the evaluation of each method, to the domain of operation (pixel or feature domain), the data used as input (iris images, polar images, or feature representations), the use of local patches, the smallest size of the input low-resolution image, or the biometric authentication experiments reported in each study. The majority of works have employed near-infrared (NIR) data. Also, the majority of works reconstruct the polar representation of the iris image [53]. It is also very common, specially among learning-based methods, to restore images at the patch level. Each patch is allowed to have its optimal reconstruction coefficients, helping to better recover local texture details, which are essential in iris recognition [8]. Regarding the smallest size of the input low-resolution image, it is common to employ databases where subjects have been naturally captured at a certain distance, such as the MBGC portal, CASIA Mobile, or Q-FIRE databases. Images in these databases have an average iris diameter in the range from 90 to 130 pixels. This is in the limit of the minimum diameter recommended as sufficient for iris recognition, which has been found by experimental studies to be  $\sim 120$  pixels [79]. Due to the lack of databases with smaller resolutions, some works carry out an artificial downsampling to achieve a smaller image size, which is a common approach in the super-resolution

literature [14]. In these studies, the average iris diameter employed is in the range from 11 to 53 pixels.

In the present paper, we also investigate the use of a trained super-resolution enhancement technique based on dictionary learning to improve the resolution of near infrared iris images. We study in depth a technique based on PCA Eigen-transformation of local image patches (eigen-patches), inspired by the system of [13] for face images. Iris images are first resized and aligned such that the eye center and the sclera are aligned. Then, a dictionary is built for each patch position, which is done by applying Principal Component Analysis to a set of collocated patches of low-resolution iris images from a training set. During testing, given a low-resolution patch, it is first projected onto the low-dimensional eigen-space. Then, the reconstruction weights are used to restore the high-resolution patch using collocated patches contained in the coupled high-resolution training set. In the employed method, the structure of the iris image is exploited in two ways: *i*) by building a patch position-dependent dictionary, which caters for a specific region of the iris; and *ii*) by allowing that each patch has its own optimum reconstruction weights, so the solution is locally optimized. In addition, unlike other methods that reconstruct images in the polar domain, we reconstruct the iris image directly, although our method is general enough to be applicable to polar images too. This makes our approach agnostic of the feature extraction method employed, given that there are comparators which do not transform the iris image to the polar domain, see [50], [51].

To evaluate the proposed method, we conducted extensive experiments with a database of 1,872 iris images. Low-resolution images were simulated by down-sampling high-resolution irises. High-resolution images, of size  $231 \times 231$ , were sub-sampled by  $1/n$ , with  $n \in \{2, 4, 6, 8, 10, 12, 14, 16\}$ . This resulted in low-resolution images of  $115 \times 115$ ,  $57 \times 57$ ,  $39 \times 39$ ,  $29 \times 29$ ,  $23 \times 23$ ,  $19 \times 19$ ,  $17 \times 17$ , and  $15 \times 15$  pixels respectively. The latter corresponds to an iris diameter of 13 pixels. Such a small resolution has not been previously employed in any iris super-resolution study, apart from ours. This paper expands our previous studies [11], [12] with a more comprehensive experimental framework. For the authentication experiments, we used publicly available feature extraction methods from popular and state-of-the-art schemes [52]. In particular, we included in this article four new iris recognition algorithms [15]–[18], which have been added to the two employed in our previous work [19], [20]. We also report identification experiments, which are lacking in the majority of iris super-resolution studies.

Our experimental section starts by reporting the Peak Signal-to-Noise Ratio (PSNR) and the Structural Similarity index (SSIM) between

the hallucinated and the corresponding high-resolution reference images. The proposed method was compared with bilinear and bicubic interpolation. Our experiments show the superiority of the presented PCA approach over these

two interpolation methods. As resolution decreases, the gain of the PCA method over bilinear or bicubic becomes more prominent, at least as measured by the PSNR. At a resolution of  $115 \times 115$  pixels, the PSNR of PCA is, on average, 0.61dB higher than the PSNR of bicubic over the whole image. At  $15 \times 15$  pixels, the difference increases to 1.46dB. This is consistent with a subjective assessment of the reconstructed images (Figure 6). When the resolution decreases, we observe that images reconstructed with bilinear or bicubic methods have much more blur than those reconstructed with PCA. On the other hand, this subjective difference is not captured by the SSIM. As we observe in our experiments, the SSIM gain of PCA over bilinear or bicubic remains around 0.01-0.02 across the whole range of resolutions evaluated. One of the drawbacks of these image fidelity metrics is precisely that they are not expected to have the same sensitivity to image degradations [21], [22]. In our case, the SSIM is not a good predictor of the subjective differences observed between images reconstructed with the PCA algorithm and with the bilinear/bicubic methods. Another drawback acknowledged in the literature is that while they may reflect the *goodness* of the enhancement (in the visual sense), they are not necessarily good predictors of the recognition accuracy [47]. In other words, an overall visual enhancement of the image does not necessarily correlate with a better performance when such image is used for recognition purposes [9].

Consequently, to evaluate the usefulness of the proposed PCA reconstruction method to improve authentication accuracy, we conducted a series of verification and identification experiments with six different iris recognition algorithms based on different features. We considered two operational scenarios: one where original high-resolution images are matched against hallucinated images ('controlled' enrolment scenario, or scenario 1), and another scenario where only hallucinated images are used ('uncontrolled' scenario, or scenario 2). The benefit of our trained approach becomes evident in comparison to bilinear or bicubic interpolations when the resolution becomes very low. For example, with an image size of  $29 \times 29$  pixels or smaller, the PCA method gives better accuracy than bilinear or bicubic interpolations for any given comparator, highlighting the superiority of our trained PCA approach to enhance very low-resolution iris images. With some comparators, the superiority of PCA is also appreciable even at a resolution of  $57 \times 57$  pixels. It is also worth noting the resilience of some comparators to severe downsampling (image size of  $15 \times 15$  pixels) when using PCA reconstruction. Two particular comparators based on Gabor wavelets showed impressive EER values of  $\sim 5\%$  and a Top-1 accuracy of 77-84% for this extreme case.

Another observation is the different behavior exhibited by the comparators employed when resolution is reduced. The two mentioned Gabor-based comparators do not show a significant degradation in performance until an image size of  $15 \times 15$  pixels is used. The majority of comparators, on the other hand, show appreciable degradation at a resolution of  $29 \times 29$  pixels. However, one particular comparator shows

a high resilience to changes in resolution, with an EER in the range of 12-13% and a Top-1 accuracy of 48-55% for any given resolution. This reinforces our previous assumption that enhancing an image in the visual sense, or measured by a metric of *fidelity*, does not necessarily contribute to a better recognition performance. Since recognition algorithms are based on different features, and their performance is not affected by the same factors, a metric of image fidelity may be useful for a particular algorithm only [9]. For this reason, focusing on the recognition performance of the algorithms employed becomes necessary.

Another interesting phenomenon is that since the performance of most comparators is not significantly affected until an image size of  $29 \times 29$  pixels, it would be feasible to use both query and test images of reduced size. They could be then up-sampled simply with bicubic or bilinear interpolation prior to recognition. This is of importance for example in scenarios where storage or data transmission capacity is limited. We also observe that the performance in scenario 2 is better than in scenario 1 for any comparator, specially at low resolutions. To further study this phenomenon, we looked into the score distributions of each comparator, observing that genuine scores shift towards the impostor ones in scenario 1 as resolution decreases, an effect not observed in scenario 2. In scenario 1, enrolment samples were taken from high-resolution images, whereas query samples were taken from reconstructed images. As resolution decreases, it is understandable that the similarity between the two types of images decreases as well, which explains the shift of the distributions of genuine scores. In scenario 2, both enrolment and query images are reconstructed images, therefore it is expected that their relative similarity does not change to the same extent than in scenario 1.

The proposed enhancement algorithm assumes that hallucination weights are the same in the low- and high-resolution manifolds. While this simplifies the problem, the low-resolution manifold is usually distorted by the one-to-many relationship between low- and high-resolution patches, so this assumption may not hold always true [14]. Another simplification is the assumption of linearity in the combination of patches from the training dictionary. In future work, we will explore methods to remove these two simplifications in order to increase performance of the hallucination algorithm. In our most recent work, for example, we are adapting methods which simultaneously consider the low- and high-resolution manifolds during the hallucination process [41]. We will also look into strategies to iteratively update the low-resolution dictionary to reduce the modality gap between low- and high-resolution patches contained within the dictionary [80]. Augmentation of the training set is another strategy that we are considering by adding spatial offset during the extraction of collocated patches, allowing to better cope with eye alignment inaccuracies and local image distortions. Lastly, we are also studying the feasibility of deep learning methods to provide an end-to-end mapping between low- and high-resolution iris images [44], [45].

## VII. ACKNOWLEDGMENT

This work was initiated while F. A.-F. was a visiting researcher at the University of Malta, with a Grant from the EU COST Action IC1106.

## REFERENCES

- [1] I. Nigam, M. Vatsa, and R. Singh, "Ocular biometrics: A survey of modalities and fusion approaches," *Inf. Fusion*, vol. 26, pp. 1–35, Nov. 2015.
- [2] A. K. Jain and A. Kumar, "Biometrics of next generation: An overview," in *Second Generation Biometrics*. Dordrecht, The Netherlands: Springer, 2011.
- [3] G. W. Quinn, P. J. Grother, and M. Ngan, "IREX IV: Part 2 compression profiles for iris image compression," NISTIR, Gaithersburg, MD, USA, Tech. Rep. 7978, 2014. [Online]. Available: <http://www.nist.gov/itl/iad/ig/irexiv.cfm>
- [4] K. Nguyen, C. Fookes, S. Sridharan, M. Tistarelli, and M. Nixon, "Super-resolution for biometrics: A comprehensive survey," *Pattern Recognit.*, vol. 78, pp. 23–42, Jun. 2018.
- [5] S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: A technical overview," *IEEE Signal Process. Mag.*, vol. 20, no. 3, pp. 21–36, May 2003.
- [6] K. Nasrollahi and T. B. Moeslund, "Super-resolution: A comprehensive survey," *Mach. Vis. Appl.*, vol. 25, no. 6, pp. 1423–1468, 2014.
- [7] L. Yue, H. Shen, J. Li, Q. Yuanc, H. Zhang, and L. Zhang, "Image super-resolution: The techniques, applications, and future," *Signal Process.*, vol. 128, pp. 389–408, Nov. 2016.
- [8] R. A. Farrugia and C. Guillemot, "Face hallucination using linear models of coupled sparse support," *IEEE Trans. Image Process.*, vol. 26, no. 9, pp. 4562–4577, Sep. 2017.
- [9] F. Alonso-Fernandez, J. Fierrez, and J. Ortega-Garcia, "Quality measures in biometric systems," *IEEE Security Privacy*, vol. 10, no. 6, pp. 52–62, Nov./Dec. 2012.
- [10] S. Baker and T. Kanade, "Limits on super-resolution and how to break them," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 9, pp. 1167–1183, Sep. 2002.
- [11] F. Alonso-Fernandez, R. A. Farrugia, and J. Bigun, "Eigen-patch iris super-resolution for iris recognition improvement," in *Proc. Eur. Signal Process. Conf. (EUSIPCO)*, Aug./Sep. 2015, pp. 76–80.
- [12] F. Alonso-Fernandez, R. A. Farrugia, and J. Bigun, "Very low-resolution iris recognition via Eigen-patch super-resolution and matcher fusion," in *Proc. IEEE Conf. Biometrics, Theory, Appl. Syst. (BTAS)*, Sep. 2016, pp. 1–8.
- [13] H.-Y. Chen and S.-Y. Chien, "Eigen-patch: Position-patch based face hallucination using eigen transformation," in *Proc. IEEE Int. Conf. Multimedia Expo*, Jul. 2014, pp. 1–6.
- [14] N. Wang, D. Tao, X. Gao, X. Li, and J. Li, "A comprehensive survey to face hallucination," *Int. J. Comput. Vis.*, vol. 106, no. 1, pp. 9–30, 2014.
- [15] C. Rathgeb and A. Uhl, "Secure iris recognition based on local intensity variations," in *Proc. ICIAR*, vol. 6112, 2010, pp. 266–275. [Online]. Available: [https://link.springer.com/chapter/10.1007/978-3-642-13775-4\\_27#citeas](https://link.springer.com/chapter/10.1007/978-3-642-13775-4_27#citeas)
- [16] D. M. Monro, S. Rakshit, and D. Zhang, "DCT-based iris recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 4, pp. 586–595, Apr. 2007.
- [17] J.-G. Ko, Y.-H. Gil, J.-H. Yoo, and K.-I. Chung, "A novel and efficient feature extraction method for iris recognition," *ETRI J.*, vol. 29, no. 3, pp. 399–401, 2007.
- [18] L. Ma, T. Tan, Y. Wang, and D. Zhang, "Personal identification based on iris texture analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 12, pp. 1519–1533, Dec. 2003.
- [19] L. Masek, "Recognition of human iris patterns for biometric identification," M.S. thesis, School Comput. Sci. Softw. Eng., Univ. Western Australia, Crawley, WA, Australia, 2003.
- [20] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, 2004.
- [21] Y. A. Y. Al-Najjar and D. C. Soong, "Comparison of image quality assessment: PSNR, HVS, SSIM, UIQI," *Int. J. Sci. Eng. Res.*, vol. 3, no. 8, pp. 1–5, 2012.
- [22] A. Hore and D. Ziou, "Image quality metrics: PSNR vs. SSIM," in *Proc. Int. Conf. Pattern Recognit.*, Aug. 2010, pp. 2366–2369.
- [23] P. Grother and E. Tabassi, "Performance of biometric quality measures," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 4, pp. 531–543, Apr. 2007.
- [24] R. Barnard et al., "High-resolution iris image reconstruction from low-resolution imagery," *Proc. SPIE*, p. 63130D, Aug. 2006. [Online]. Available: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/6313/1/High-resolution-iris-image-reconstruction-from-low-resolution-imagery/10.1117/12.681930.short>
- [25] G. Fahmy, "Super-resolution construction of iris images from a visual low resolution face video," in *Proc. 9th Int. Symp. Signal Process. Appl.*, Feb. 2007, pp. 1–4.
- [26] K. Nguyen, C. Fookes, S. Sridharan, and S. Denman, "Focus-score weighted super-resolution for uncooperative iris recognition at a distance and on the move," in *Proc. 25th Int. Conf. Image Vis. Comput. New Zealand (IVCNZ)*, Nov. 2010, pp. 1–8.
- [27] K. Nguyen, C. Fookes, S. Sridharan, and S. Denman, "Quality-driven super-resolution for less constrained iris recognition at a distance and on the move," *IEEE Trans. Inf. Forensics Security*, vol. 6, no. 4, pp. 1248–1258, Dec. 2011.
- [28] K. Nguyen, C. Fookes, and S. Sridharan, "Robust mean super-resolution for less cooperative NIR iris recognition at a distance and on the move," in *Proc. Symp. Inf. Commun. Technol. (SoICT)*, 2010, pp. 122–127.
- [29] K. Hollingsworth, T. Peters, K. W. Bowyer, and P. Flynn, "Iris recognition using signal-level fusion of frames from video," *IEEE Trans. Inf. Forensics Security*, vol. 4, no. 4, pp. 837–848, Dec. 2009.
- [30] R. Jillela, A. Ross, and P. J. Flynn, "Information fusion in low-resolution iris videos using principal components transform," in *Proc. IEEE Workshop Appl. Comput. Vis. (WACV)*, Jan. 2011, pp. 262–269.
- [31] H. Ren, Y. He, J. Pan, and L. Li, "Super resolution reconstruction and recognition for iris image sequence," in *Proc. 7th Chin. Conf. Biometric Recognit. (CCBR)*, 2012, pp. 193–201.
- [32] N. Othman, N. Houmani, and B. Dorizzi, "Improving video-based iris recognition via local quality weighted super resolution," in *Proc. 2nd Int. Conf. Pattern Recognit. Appl. Methods (ICPRAM)*, Feb. 2013, pp. 623–629.
- [33] N. Othman and B. Dorizzi, "Impact of quality-based fusion techniques for video-based iris recognition at a distance," *IEEE Trans. Inf. Forensics Security*, vol. 10, no. 8, pp. 1590–1602, Aug. 2015.
- [34] S.-H. Hsieh, Y.-H. Li, C.-H. Tien, and C.-C. Chang, "Extending the capture volume of an iris recognition system using wavefront coding and super-resolution," *IEEE Trans. Cybern.*, vol. 46, no. 12, pp. 3342–3350, Dec. 2016.
- [35] A. Deshpande and P. P. Patavardhan, "Super resolution and recognition of long range captured multi-frame iris images," *IET Biometrics*, vol. 6, no. 5, pp. 360–368, Sep. 2017.
- [36] A. Deshpande and P. Patavardhan, "Multi-frame super-resolution for long range captured iris polar image," *IET Biometrics*, vol. 6, no. 2, pp. 108–116, Mar. 2017.
- [37] J. Huang, L. Ma, T. Tan, and Y. Wang, "Learning based resolution enhancement of iris images," in *Proc. BMVC*, 2003, pp. 1–10.
- [38] K. Y. Shin, K. R. Park, B. J. Kang, and S. J. Park, "Super-resolution method based on multiple multi-layer perceptrons for iris recognition," in *Proc. 4th Int. Conf. Ubiquitous Inf. Technol. Appl.*, Dec. 2009, pp. 1–5.
- [39] J. Liu, Z. Sun, and T. Tan, "Distance metric learning for recognizing low-resolution iris images," *Neurocomputing*, vol. 144, pp. 484–492, Nov. 2014.
- [40] R. Aljadaany, K. Luu, S. Venugopalan, and M. Savvides, "IRIS super-resolution via nonparametric over-complete dictionary learning," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2015, pp. 3856–3860.
- [41] F. Alonso-Fernandez, R. A. Farrugia, and J. Bigun, "Iris super-resolution using iterative neighbor embedding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jul. 2017, pp. 655–663.
- [42] F. Alonso-Fernandez, R. A. Farrugia, and J. Bigun, "Learning-based local-patch resolution reconstruction of iris smart-phone images," in *Proc. Int. Joint Conf. Biometrics (IJCB)*, Oct. 2017, pp. 787–793.
- [43] Q. Zhang, H. Li, Z. He, and Z. Sun, "Image super-resolution for mobile iris recognition," in *Proc. 11th Chin. Conf. Biometric Recognit. (CCBR)*, 2016, pp. 399–406.
- [44] E. Ribeiro, A. Uhl, F. Alonso-Fernandez, and R. A. Farrugia, "Exploring deep learning image super-resolution for iris recognition," in *Proc. 25th Eur. Signal Process. Conf. (EUSIPCO)*, Aug./Sep. 2017, pp. 2176–2180.
- [45] E. Ribeiro, A. Uhl, and F. Alonso-Fernandez, "Iris super-resolution using CNNs: Is photo-realism important to iris recognition?" *IET Biometrics*, vol. 8, no. 1, pp. 69–78, Jan. 2019.



- [46] K. Nguyen, C. Fookes, S. Sridharan, and S. Denman, "Feature-domain super-resolution for iris recognition," in *Proc. 18th IEEE Int. Conf. Image Process.*, Sep. 2011, pp. 3197–3200.
- [47] K. Nguyen, S. Sridharan, S. Denman, and C. Fookes, "Feature-domain super-resolution framework for Gabor-based face and iris recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2012, pp. 2642–2649.
- [48] K. Nguyen, C. Fookes, S. Sridharan, and S. Denman, "Feature-domain super-resolution for iris recognition," *Comput. Vis. Image Understand.*, vol. 117, no. 10, pp. 1526–1535, 2013.
- [49] J. Liu, Z. Sun, and T. Tan, "Code-level information fusion of low-resolution iris image sequences for personal identification at a distance," in *Proc. 6th Int. Conf. Biometrics, Theory, Appl. Syst. (BTAS)*, Sep. 2013, pp. 1–6.
- [50] F. Alonso-Fernandez, P. Tome-Gonzalez, V. Ruiz-Albacete, and J. Ortega-Garcia, "Iris recognition based on SIFT features," in *Proc. IEEE Int. Conf. Biometrics, Identity Secur. (BIDS)*, Sep. 2009, pp. 1–8.
- [51] K. Nguyen, C. Fookes, A. Ross, and S. Sridharan, "Iris recognition with off-the-shelf CNN features: A deep learning perspective," *IEEE Access*, vol. 6, pp. 18848–18855, 2017.
- [52] C. Rathgeb, A. Uhl, and P. Wild, *Iris Biometrics: From Segmentation to Template Security* (Advances in Information Security), vol. 59. New York, NY, USA: Springer, 2013.
- [53] J. Daugman, "How iris recognition works," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 14, no. 1, pp. 21–30, Jan. 2004.
- [54] *CASIA Databases*. Accessed: Aug. 2018. [Online]. Available: <http://biometrics.idealltest.org/>
- [55] P. J. Phillips et al., "Overview of the multiple biometrics grand challenge," in *Proc. Int. Conf. Biometrics (ICB)*, vol. 5558, 2009, pp. 705–714.
- [56] P. A. Johnson, P. Lopez-Meyer, N. Sazonova, F. Hua, and S. Schuckers, "Quality in face and iris research ensemble (Q-FIRE)," in *Proc. IEEE Int. Conf. Biometrics, Theory, Appl. Syst. (BTAS)*, Sep. 2010, pp. 1–6.
- [57] K. B. Raja, R. Raghavendra, V. K. Vemuri, and C. Busch, "Smart-phone based visible iris recognition using deep sparse filtering," *Pattern Recognit. Lett.*, vol. 57, pp. 33–42, May 2015.
- [58] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 2, pp. 295–307, Feb. 2016.
- [59] B. J. Kang and K. R. Park, "Real-time image restoration for iris recognition systems," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 37, no. 6, pp. 1555–1566, Dec. 2007.
- [60] N. D. Kalka, J. Zuo, N. A. Schmid, and B. Cukic, "Estimating and fusing quality factors for iris biometric images," *IEEE Trans. Syst., Man, Cybern. A, Syst. Humans*, vol. 40, no. 3, pp. 509–524, May 2010.
- [61] E. R. Dowski and W. T. Cathey, "Extended depth of field through wave-front coding," *Appl. Opt.*, vol. 34, no. 11, pp. 1859–1866, Apr. 1995.
- [62] M. K. Ng, H. Shen, E. Y. Lam, and L. Zhang, "A total variation regularization based super-resolution reconstruction algorithm for digital video," *EURASIP J. Adv. Signal Process.*, vol. 2007, no. 1, p. 074585, Jun. 2007.
- [63] W. T. Freeman, T. R. Jones, and E. C. Pasztor, "Example-based super-resolution," *IEEE Comput. Graph. Appl.*, vol. 22, no. 2, pp. 56–65, Mar. 2002.
- [64] B. Li, H. Chang, S. Shan, and X. Chen, "Locality preserving constraints for super-resolution with neighbor embedding," in *Proc. 16th Int. Conf. Image Process. (ICIP)*, Nov. 2009, pp. 1189–1192.
- [65] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in *Computer Vision—ECCV*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. Cham, Switzerland: Springer, 2014, pp. 184–199.
- [66] S. Schuler, C. Leistner, and H. Bischof, "Fast and accurate image upscaling with super-resolution forests," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 3791–3799.
- [67] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, Aug. 2013.
- [68] J. Kim, J. K. Lee, and K. M. Lee, "Accurate image super-resolution using very deep convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Las Vegas, NV, USA, 2016, pp. 1646–1654, doi: [10.1109/CVPR.2016.182](https://doi.org/10.1109/CVPR.2016.182).
- [69] C. Ledig et al., "Photo-realistic single image super-resolution using a generative adversarial network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Honolulu, HI, USA, 2017, pp. 105–114, doi: [10.1109/CVPR.2017.19](https://doi.org/10.1109/CVPR.2017.19).
- [70] M. Turk and A. Pentland, "Eigenfaces for recognition," *J. Cognit. Neurosci.*, vol. 3, no. 1, pp. 71–86, 1991.
- [71] A. Serrano, I. M. de Diego, C. Conde, and E. Cabello, "Recent advances in face biometrics with Gabor wavelets: A review," *Pattern Recognit. Lett.*, vol. 31, no. 5, pp. 372–381, 2010.
- [72] J. Daugman, "New methods in iris recognition," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 37, no. 5, pp. 1167–1175, Oct. 2007.
- [73] F. Alonso-Fernandez and J. Bigun, "Near-infrared and visible-light periorcular recognition with gabor features using frequency-adaptive automatic eye detection," *IET Biometrics*, vol. 4, no. 2, pp. 74–89, 2015.
- [74] H. Hofbauer, F. Alonso-Fernandez, P. Wild, J. Bigun, and A. Uhl, "A ground truth for iris segmentation," in *Proc. 22nd Int. Conf. Pattern Recognit. (ICPR)*, Aug. 2014, pp. 527–532.
- [75] C. Rathgeb, A. Uhl, P. Wild, and H. Hofbauer, "Design decisions for an iris recognition SDK," in *Handbook of Iris Recognition* (Advances in Computer Vision and Pattern Recognition), 2nd ed, K. W. Bowyer and M. J. Burge, Eds. London, U.K.: Springer, 2016.
- [76] A. K. Jain, K. Nandakumar, and A. Ross, "50 years of biometric research: Accomplishments, challenges, and opportunities," *Pattern Recognit. Lett.*, vol. 79, pp. 80–105, Aug. 2016.
- [77] D. Thapa, K. Raahemifar, W. R. Bobier, and V. Lakshminarayanan, "A performance comparison among different super-resolution techniques," *Comput. Elect. Eng.*, vol. 54, pp. 313–329, Aug. 2016.
- [78] K. W. Bowyer, K. P. Hollingsworth, and P. J. Flynn, "A survey of iris biometrics research: 2008–2010," in *Handbook of Iris Recognition* (Advances in Computer Vision and Pattern Recognition), M. Burge and K. Bowyer, Eds. London, U.K.: Springer, 2013, pp. 15–54.
- [79] E. Tabassi, P. Grother, and W. Salamon, "IREX II-IQCE iris quality calibration and evaluation: Performance of iris image quality assessment algorithms," NISTIR, Gaithersburg, MD, USA, Tech. Rep. 7296, 2011. [Online]. Available: <http://iris.nist.gov/irex/>
- [80] J. Jiang, R. Hu, Z. Wang, and Z. Han, "Face super-resolution via multilayer locality-constrained iterative neighbor embedding and intermediate dictionary learning," *IEEE Trans. Image Process.*, vol. 23, no. 10, pp. 4220–4231, Oct. 2014.



**FERNANDO ALONSO-FERNANDEZ** received the M.S. and Ph.D. degrees in telecommunications engineering from the Universidad Politécnica de Madrid, Spain, in 2003 and 2008, respectively. Since 2010, he has been with the Centre for Applied Intelligent Systems Research, Halmstad University, Sweden. Since 2017, he has been an Associate Professor with Halmstad University. He has been actively involved in multiple EU (such as FP6 Biosecure NoE and COST IC1106) and national projects focused on biometrics and human-machine interaction. He has over 70 international contributions at refereed conferences and journals and has authored several book chapters. His research interests include signal and image processing, pattern recognition, and biometrics, with an emphasis on facial cues and body biosignals. He was a recipient of a Marie Curie IEF and a Postdoctoral Fellowship from the Swedish Research Council and a Project Research Grant for Junior Researchers of the Swedish Research Council. He has co-chaired the 9th IAPR International Conference on Biometrics, in 2016. Since 2018, he has been serving on the Editorial Board of the *IET Biometrics* and an Associate Editor for the IEEE Biometrics Council Newsletter.



**REUBEN A. FARRUGIA** (S'04–M'09–SM'17) received the degree in electrical engineering and the Ph.D. degree from the University of Malta, Msida, Malta, in 2004 and 2009, respectively. In 2008, he was appointed as an Assistant Lecturer with the University of Malta, where he is currently a Senior Lecturer. He has been serving in technical and organizational committees for several national and international conferences. He served as the General Chair for the IEEE International Workshop on Biometrics and Forensics and the Technical Programme Co-Chair for the IEEE Visual Communications and Image Processing, in 2014. He has been contributing as a reviewer for several journals and conferences, including the IEEE TRANSACTIONS ON IMAGE PROCESSING, the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO AND TECHNOLOGY, and the IEEE TRANSACTIONS ON MULTIMEDIA. In 2013, he was appointed as the National Contact Point of the European Association of Biometrics.



**JOSEF BIGUN** (M'88–SM'98–F'03) received the M.S. and Ph.D. degrees from Linköping University, Sweden, in 1983 and 1988, respectively. From 1988 to 1998, he was with the École Polytechnique Fédérale de Lausanne, Switzerland. He was a Professor with the Signal Analysis Chair, Halmstad University, and the Chalmers University of Technology, in 1998. His scientific interests include a broad field in computer vision, texture and motion analysis, biometrics, and the understanding of biological recognition mechanisms. He has co-chaired several international conferences. He has been contributing as a Referee or as an Editorial Board Member of journals, including *Pattern Recognition Letters* and the IEEE TRANSACTIONS ON IMAGE PROCESSING. He served on the executive committees for several associations, including IAPR. He has been an elected Fellow of IAPR.



**ESTER GONZALEZ-SOSA** received the B.S. degree in computer science and the M.Sc. degree in electrical engineering from the Universidad de Las Palmas de Gran Canaria, in 2012 and 2014, respectively, and the Ph.D. degree from the Biometric Recognition Group, Universidad Autonoma de Madrid, in 2017. She has carried out several research internships in worldwide leading groups in biometric recognition, such as TNO, EURECOM, or Rutgers University. Her research interests include pattern recognition, signal processing, and biometrics, with an emphasis on face, body, soft biometrics, and millimeter imaging. He was a recipient of the UNITECO AWARD from the Spanish Association of Electrical Engineers, the Competitive Obra Social La CAIXA Scholarship, and the European Research Award 2018 of the European Association for Biometrics.

ests include pattern recognition, signal processing, and biometrics, with an emphasis on face, body, soft biometrics, and millimeter imaging. He was a recipient of the UNITECO AWARD from the Spanish Association of Electrical Engineers, the Competitive Obra Social La CAIXA Scholarship, and the European Research Award 2018 of the European Association for Biometrics.

...



**JULIAN FIERREZ** received the M.Sc. and Ph.D. degrees in electrical engineering from the Universidad Politecnica de Madrid, Spain, in 2001 and 2006, respectively. Since 2002, he has been with the Biometrics and Data Pattern Analytics Lab-ATVS, Universidad Autonoma de Madrid, where he has been an Associate Professor since 2010. From 2007 to 2009, he held a Postdoctoral position at Michigan State University under a Marie Curie Fellowship. He has been actively involved in the last 15 years in large EU projects focused on biometrics (such as BIOSECURE, TABULA RASA, and BEAT). He was a recipient of a number of distinctions, including the EAB European Biometric Industry Award 2006, the EURASIP Best Ph.D. Award 2012, and the IAPR Young Biometrics Investigator Award 2017. Since 2016, he has been an Associate Editor of the IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY and the IEEE Biometrics Council Newsletter.

last 15 years in large EU projects focused on biometrics (such as BIOSECURE, TABULA RASA, and BEAT). He was a recipient of a number of distinctions, including the EAB European Biometric Industry Award 2006, the EURASIP Best Ph.D. Award 2012, and the IAPR Young Biometrics Investigator Award 2017. Since 2016, he has been an Associate Editor of the IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY and the IEEE Biometrics Council Newsletter.