Mitigating Effects of Plastic Surgery: Fusing Face and Ocular Biometrics

Raghavender Jillela and Arun Ross West Virginia University Morgantown, WV, USA

{Raghavender.Jillela, Arun.Ross}@mail.wvu.edu

Abstract

The task of successfully matching face images obtained before and after plastic surgery is a challenging problem. The degree to which a face is altered depends on the type and number of plastic surgeries performed, and it is difficult to model such variations. Existing approaches use learning based methods that are either computationally expensive or rely on a set of training images. In this work, a fusion approach is proposed that combines information from the face and ocular regions to enhance recognition performance in the identification mode. The proposed approach provides the highest reported recognition performance on a publicly accessible plastic surgery database, with a rank-one accuracy of 87.4%. Compared to existing approaches, the proposed approach is not learning based and reduces computational requirements. Furthermore, a systematic study of the matching accuracies corresponding to various types of surgeries is presented.

1. Introduction

Plastic surgery generally refers to a medical procedure that involves modifying the appearance of external anatomical features using surgical methods¹. Based on their purpose, plastic surgeries can be broadly classified into two categories:

- Reconstructive: These surgeries are performed mainly to reconstruct the generic appearance of a facial feature, so that its functionality is restored or improved. For example, surgical treatment of ptosis (drooping of the upper eyelid due to weak muscles, that can cause vision interference).
- Aesthetic improvement: These surgeries are performed to alter the appearance of a fully functional feature, solely with the purpose of aesthetic improvement.

For example, restoring damaged skin due to burn injuries or accidents.

Facial plastic surgeries have become increasingly popular in the recent past, especially for aesthetic improvement purposes. A report from the American Society of Plastic Surgery states that a total of 13.8 million cosmetic and reconstructive plastic surgeries were performed just in the year 2011². Three of the top five surgeries in this set relate to the modification of facial features³. Some of the major facial plastic surgeries include: rhinoplasty (nose surgery), blepharoplasty (eyelid surgery), brow lift (eyebrow surgery), otoplasty (ear surgery), and rhytidectomy (face lift surgery) (see Figure 1). A detailed, but nonexhaustive list of facial plastic surgeries is provided in [13].



Figure 1. Some of the major facial plastic surgeries. Image taken from the FRGC database.

The degree to which the appearance of a human face can be modified by plastic surgery, depends on the number and the types of surgeries performed. Figure 2 shows two image pairs⁴ containing modifications based on the number of

¹American Society of Plastic Surgeons, *The History of Plastic Surgery*, http://www.plasticsurgery.org, 2012

²http://www.plasticsurgery.org/Documents/news-resources/statistics/2011statistics/2011-cosmetic-procedures-trends-statistics.pdf

³http://www.plasticsurgery.org/Documents/news-resources/statistics/2011statistics/2011-top-5-cosmetic-procedures-statistics.pdf

⁴Top row images: *Facial Plastic Surgery Database*, http://research.iiitd.edu.in/groups/iab/resources.html. Bottom row images: *10 worst celebrity plastic surgery mishaps*, http://www.womansday.com. **Eye regions** in the Facial Plastic Surgery Database images have been blurred in this paper to preserve the privacy of individuals.

surgeries. Humans can recognize such variations in facial appearance with very low, or moderate level of difficulty. However, plastic surgeries can negatively impact the performance of automatic face recognition systems [5] because of the following reasons:

- Most face recognition algorithms take the holistic appearance of the face into account for feature extraction. A wide number of plastic surgeries can alter the overall appearance of the face, thereby reducing the similarity between genuine image pairs.
- Depending on the type and number of surgeries performed, a multitude of variations are possible in the appearance of the face. Such variations are difficult to be modeled by existing face recognition algorithms.

In some cases, facial plastic surgery can unintentionally serve as a method to circumvent automatic face recognition systems.



Figure 2. Images showing the degree to which the appearance of a human face can be modified by plastic surgeries. Top row: (a) before and (b) after a minor plastic surgery (blepharoplasty). Bottom row: (c) before, and (d) after multiple plastic surgeries.

Only recently, have researchers from the biometric community begun to investigate the effect of plastic surgery on face recognition algorithms [13, 3, 1]. Prior to that, research on this topic was stymied by the lack of databases containing pre- and post-surgery face images. Singh et al. [13] assembled the first database that contains face images related to various types of plastic surgeries. The low recognition accuracies that have been reported on this database seem to suggest that the task of face recognition on plastic surgery images is a challenging problem.

2. Existing Approaches

Singh et al. [13] reported recognition accuracies on the plastic surgery database using six different face recognition algorithms: Principal Component Analysis (PCA), Fisher Discriminant Analysis (FDA), Local Feature Analysis (LFA), Circular Local Binary Patterns (CLBP), Speeded Up Robust Features (SURF), and Neural network Architecture based 2-D Log Polar Gabor Transform (GNN). These algorithms were selected because they provide a combination of appearance-based, feature-based, descriptorbased, and texture-based feature extraction and matching approaches. Despite combining local and global recognition approaches, the matching performance obtained was rather low (see Table 1). Marsico et al. [4] used correlationbased face recognition on pose and illumination normalized images. Bhatt et al. [3] used an evolutionary granular approach with CLBP and SURF features to process tessellated face images. Aggarwal et al. [1] used a combination of face recognition by parts and sparse representation approach. The matching schemes used in the literature, along with their rank-one recognition accuracies are listed in Table 1.

Table 1. List of algorithms used for performing face recognition on plastic surgery images and the corresponding rank-one accuracies.

Authors	Algorithm used	Rank-one
		Accuracy
Singh et al.	PCA	29.1%
	FDA	32.5%
	LFA	38.6%
	CLBP	47.8%
	SURF	50.9%
	GNN	54.2%
Marsico et al.	Correlation based approach	70.6%
Bhatt et al.	Evolutionary granular approach	78.6%
Aggarwal et al.	Combination of recognition-by-	77.9%
	parts & sparse representation	
	approaches	

3. Motivation

A careful study of the existing research in this area reveals the following interesting observations:

- 1. A majority of the algorithms that have been used are learning based which require a carefully selected set of training images. Despite this, it can be observed that the rank-one identification accuracy did not exceed 79%.
- 2. No commercial face recognition systems have been used for evaluating recognition performance.
- 3. No biometric fusion schemes have been explored in an attempt to improve recognition accuracy.

Considering the rapid advancements in the area of face recognition, there is a need to improve recognition accuracy on facial images exhibiting plastic surgeries. To this end, the present work provides the following contributions:

- The recognition performance of two commercial face recognition systems on plastic surgery images is evaluated. It is demonstrated that these systems can provide performance on par with the learning based methods.
- An information fusion approach that combines independently processed ocular information with the face biometric is presented. The proposed approach is observed to provide the current highest reported recognition performance on plastic surgery images.

The usage of ocular information has the following benefits:

An empirical analysis suggests that the number of plastic surgeries that affect the appearance of the ocular region, compared to those that alter the holistic appearance of the face, is very small. Table 2 shows a list of surgeries categorized based on the primary facial region impacted by the surgery. It is apparent, from this table that only a few of the surgeries directly impact the ocular region. Thus, in post-surgery images, the ocular region is likely to be more stable than the global facial appearance. Sample images demonstrating this observation are provided in Figure 3.

Table 2. List of major facial plastic surgeries separated by the corresponding regions whose appearance can be potentially affected.

Primary re-	Type of surgery		
gion of impact			
Entire face (10)	Rhinoplasty, Genioplasty, Cheek implant,		
	Otoplasty, Liposhaving, Skin resurfacing,		
	Rhytidectomy, Lip augmentation, Craniofa-		
	cial surgery, Dermabrasion		
Only the ocular	Blepharoplasty, Brow lift, Non-surgical local		
region (3)	procedures (e.g., BOTOX)		

- 2. Since the ocular region can be directly obtained from the face image, no additional sensors are necessary thereby making it a good choice for fusion.
- 3. Existing research suggests that the fusion of ocular information with the face biometric can lead to improved recognition performance [10].

4. Ocular Recognition

The ocular region refers to a small region around the eye, containing the eye, the eyebrows, and the surrounding skin. Recent research has shown that the ocular information can be used as a soft biometric [10, 8, 7]. It has been experimentally demonstrated that the ocular information can be used in lieu, or to improve the matching accuracy, of the iris [12] and face [10] under non-ideal conditions. While there are no specific guidelines for the dimensions of the periocular region, Park et al. [10] suggest that including the eyebrows



Figure 3. Facial images of a subject (a) before, and (b) after undergoing rhytidectomy. (c) and (d): Corresponding ocular images of the same subject. Note that the variation in the appearance of the face, from a visual perspective, is much larger than that of the ocular region.

can result in higher matching accuracy. Most existing approaches use monocular information from either the left or right side of an individual's face. In this study, information corresponding to both the eyes (bi-ocular [11]) is considered. The reasons for using bi-ocular information are:

- 1. Park et al. [10] showed that the fusion of the left and right periocular region improves matching accuracy.
- 2. The spatial resolution of the face images used in this work is very low (explained in Section 6). Thus, utilizing the bi-ocular region ensures an effective use of information.

Some examples of the bi-ocular images used in this work are shown in Figure 4.



Figure 4. Sample bi-ocular images used in this work. Note that the images have been resized for the purpose of clarity.

5. Proposed Approach

Based on the initial hypothesis, the proposed approach combines the information from the face and ocular regions at score level to improve the recognition performance. Two commercial face recognition software, Verilook 3.2^5 and PittPatt⁶, were used in this work. The use of these software

 ⁵ Verilook 3.2, Neurotechnology, http://www.neurotechnology.com
⁶ PittPatt, Pittsburgh Pattern Recognition, now acquired by Google

helps in establishing baseline performances due to commercial face recognition systems on plastic surgery images. This also helps in avoiding computationally expensive training based methods.

To perform automatic cropping of ocular regions from face images, a face detector based on the Viola-Jones Adaboost algorithm [14] was used. This step also serves as a basic quality check, where challenging images that could cause Failure To Enroll (FTE) error are discarded (e.g., images containing very small inter-ocular distances, partial faces, etc.). Ocular regions extracted from low-resolution face images could be very noisy and impact the recognition performance. To perform feature extraction from ocular regions, two techniques, viz., Scale Invariant Feature Transform (SIFT) [6] and Local Binary Patterns (LBP) [9] were used. The combination of SIFT and LBP techniques allows for image feature extraction at both local and global levels, respectively. Furthermore, SIFT and LBP have been the most significantly used techniques' in the ocular recognition literature [10, 12]. The use of these techniques helps in maintaining uniformity for performance comparisons.

SIFT The Scale Invariant Feature Transform (SIFT) technique works by detecting and encoding information around local keypoints that are invariant to scale and orientation changes of an image. Given an image I(x, y), the corresponding scale space image $L(x, y, \sigma)$, at a scale σ , is obtained as $L(x, y, \sigma) = G(x, y, \sigma) * I(x, y)$, where $G(x, y, \sigma)$ is a Gaussian filter and the symbol * represents a convolution operation. A set of Difference of Gaussian (DoG) images, between scales separated by a multiplicative factor k, are obtained by the equation DoG = $(G(x, y, k\sigma) - G(x, y, \sigma)) * I(x, y)$. From this set of images, extrema points are detected by choosing the local maxima or minima among eight neighbors of a pixel in the current image, and nine neighbors each in the scales above and below the current DoG image. These extrema points correspond to image discontinuities and are further processed to exclude unstable extrema points. A 36 bin orientation histogram covering the [0, 360] interval around each keypoint is then generated using the gradient magnitude m(x, y) and orientation $\theta(x, y)$ information, where m(x, y) = [((L(x + y)))] $1,y)-L(x-1,y))^2+(L(x,y+1)-L(x,y-1))^2)]^{\frac{1}{2}},$ and $\theta(x,y) = tan^{-1} \left(\frac{(L(x,y+1)-L(x,y-1))}{(L(x+1,y)-L(x-1,y))} \right).$ The orientation of the keypoint is computed as the highest peak in the orientation histogram associated with it. The feature vector is obtained by sampling the gradient magnitude and orientations within a descriptor window of size 16×16 around a keypoint. The final keypoint descriptor of dimension $4 \times 4 \times 8$

is generated by computing an 8 bin orientation histogram over 4×4 sample regions within the descriptor window. In this work, a publicly available MATLAB implementation⁸ of SIFT was used.

LBP Given an image *I*, sample points are first determined by uniformly sampling the image at a fixed frequency. A block of size 8×8 pixels around every sampling point is considered as a region of interest (ROI). For each pixel *p* within the ROI, a neighborhood of size 3×3 pixels is considered for LBP value generation, as shown in Figure 5.

p_1	p_2	p_3
p_0	p	p_4
p_7	p_6	p_5

Figure 5. Neighborhood for computing the LBP of pixel p.

The mathematical equation for computing the LBP value at a pixel p is given by:

$$LBP(p) = \sum_{k=0}^{k=7} 2^k f(I(p) - I(p_k)), \qquad (1)$$

where $I(p_k)$ represents the intensity value of pixel p_k , and

$$f(x) = \begin{cases} 1 & \text{if } x \ge 0, \\ 0 & \text{if } x < 0. \end{cases}$$
(2)

The LBP values of all the pixels within a given ROI are then quantized into an 8 bin histogram. Histograms corresponding to all sampling points are then concatenated to form a final feature vector. Euclidean distance was used to measure the similarity between two feature vectors. In this work, to perform LBP feature extraction and matching, every RGB ocular image was first decomposed into its individual R, G, and B channels. Each channel was sampled at a frequency of 16 pixels, yielding a total of 465 sampling points. The final LBP feature vectors for each channel were of size 1×3720 (concatenating 8 bin histograms for 465 sampling points).

Score-level fusion For a given image, let S_{VL} and S_{PP} denote the face match scores obtained using Verilook and PittPatt, respectively. S_{SIFT} , represents the SIFT ocular score and S_{LBP-R} , S_{LBP-G} , and S_{LBP-B} represent the LBP ocular scores for each of the R, G, and B channels of an ocular image, respectively. A final LBP ocular score, S_{LBP} , was computed by considering the average of S_{LBP-R} , S_{LBP-G} , and S_{LBP-B} . The averaging operation was chosen because it provided relatively better performance, when compared to the other operators (e.g., min,

⁷Gradient Orientation Histogram (GO), another global level feature extraction technique, has also been widely used in ocular recognition literature. However, it was excluded in this study because LBP outperformed GO.

⁸http://www.vlfeat.org/overview/sift.html

max, etc.). Score-level fusion was then performed to combine the face and ocular information. A schematic representation of the proposed score-level fusion approach is shown in Figure 6.



Figure 6. A schematic representation of the proposed approach.

6. Database

Images from the plastic surgery database described in [13] are used in this work. Currently, this is the only publicly available database that contains images of subjects captured before and after various types of plastic surgeries. Biometric databases are typically assembled through a concerted data collection process by acquiring the required data from the subjects directly. On the contrary, this database was generated by downloading facial images from two different plastic surgery information websites⁹. This introduces significant challenges in working with this database, such as: (a) low resolution, (b) variations in scale and expression, and (c) duplicate entries. Figure 7 shows sample images illustrating these challenges.

Three different datasets are considered in this work. The details of each dataset are listed as follows:

Face dataset A All the images contained in the plastic surgery database were used in this dataset. This dataset contains frontal face images of 900 subjects. For each subject, there is 1 pre-surgery facial image and 1 post-surgery facial image. The resolution of the images range from 163×131 to 288×496 pixels, and the inter-ocular distance varies from 20 to 100 pixels. These images are divided into a gallery (containing 900 pre-surgery images), and a probe set (containing the corresponding 900 post-surgery images). This dataset helps in performing a direct comparison of recognition performances obtained by commercial recognition systems, with those reported in the existing literature.

Face dataset B This dataset was obtained by discarding images from *face dataset A* corresponding to: (a) failures in face detection using the Adaboost algorithm, and (b) very low image resolution that can yield noisy ocular regions (as



Figure 7. Images exhibiting some of the challenges in the facial plastic surgery database. (a) and (d): images with varying resolution, scale and inter-ocular distances corresponding to the same subject. (b) and (e): variations in expressions of a subject. (c) and (f): duplicate entries. The image in (c) is listed as ID #26300 and its duplicate image in (f) is re-listed as ID #28519. Note the difference in identification labels, although they belong to the same subject who has undergone multiple surgeries. This incorrect labeling can negatively impact the perceived matching accuracy.

described in Section 5). As a result, a total of 478 images corresponding to 239 subjects were selectively discarded from *face dataset A*. The remaining 1322 images are divided into a gallery (containing 661 pre-surgery images), and a probe set (containing the corresponding 661 post-surgery images). A set of 568 face images corresponding to 568 unique subjects from the FRGC database ¹⁰ were added to the gallery. These images have a resolution of 1704×2272 pixels, with an average inter-ocular distance of 260 pixels. These additional images help in (a) compensating for the effect of discarded images, (b) observing the robustness of the proposed feature extraction and matching techniques by increasing the number of impostor scores, and (c) providing a heterogenous combination of surgically modified and unmodified face images.

Ocular dataset This dataset was generated by automatically cropping the bi-ocular regions from images in *face image dataset B*. The average resolutions of the cropped bi-ocular regions range from 115×54 to 842×392 pixels. All the ocular images in both the gallery and probe sets were resized to a fixed resolution of 500×250 pixels. This helps in ensuring a fixed-size feature vector when global feature extraction schemes are used.

The total number of images used in the face and ocular datasets, along with their spatial resolutions are summarized

⁹www.locateadoc.com and www.surgery.org

¹⁰NIST, Face Recognition Grand Challenge (FRGC) Database, http://www.nist.gov/itl/iad/ig/frgc.cfm

in Table 3.

		# of images	Resolution
Face dataset A	Gallery	900	$\begin{array}{c} 163 \times 131 \text{ to } 334 \times \\ 466 \end{array}$
	Probe	900	147×226 to 288×496
Face dataset B	Gallery	1229(661 + 568)	288×250 to 1704×2272
	Probe	661	$\begin{array}{c} 288 \times 250 \text{ to } 288 \times \\ 485 \end{array}$
Ocular dataset	Gallery	1229	500×250
Oculai uataset	Probe	661	500×250

Table 3. Number of images used in each dataset, along with their spatial resolutions.

7. Experiments and Results

To determine the face recognition performance, every image in the probe set of the face image dataset was matched against the gallery. The same protocol was used for the ocular image dataset to generate the ocular match scores. When performing score-level fusion, the score matrices corresponding to *face dataset B* and *ocular dataset* were normalized in the [0, 1] range using min-max normalization.

7.1. Face recognition performance

In this work, Cumulative Match Characteristic (CMC) curves were used to summarize the identification performances. Figure 8 shows the CMC curves obtained using the commercial face recognition systems on the considered face datasets. The rank-one recognition accuracies obtained using Verilook and PittPatt on *face dataset A* were observed ¹¹ to be (a) 70.3% (b) 65.8%, respectively. Similarly, the corresponding rank-one recognition accuracies obtained on *face dataset B* were observed to be (a) 73.9% and (b) 81.4%, respectively. From the figure, it can be observed that PittPatt provides better recognition performance than VeriLook when low resolution images are discarded.

7.2. Ocular recognition performance

The rank-one accuracies obtained using LBP and SIFT on the *ocular database* were observed to be 45.6% and 48.1%, respectively. The CMC curves for both the techniques are shown in Figure 9. From the figure, it can be observed that SIFT provides better ocular recognition performance compared to LBP. This is because SIFT depends on local key-point information that is scale and rotation invariant. On the other hand, the LBP match score is dependent on the similarity of global level information that is affected by misalignment of gallery and probe images.



Figure 8. CMC curves showing the recognition performances of VeriLook and PittPatt on *face dataset A* and *face dataset B*.



Figure 9. CMC curves showing the recognition performances of LBP and SIFT on *ocular dataset*.

7.3. Score-level fusion performance

Weighted score-level fusion is used to combine the normalized scores from the following scenarios: (a) face (VeriLook and PittPatt scores obtained using face dataset B), (b) ocular (LBP and SIFT scores obtained using ocular dataset), and (c) face and ocular (VeriLook and PittPatt scores obtained using *face dataset B*, LBP and SIFT scores obtained using *ocular dataset*). These normalized scores were combined using the simple sum rule with different weights, with an objective of maximizing the rank-one accuracy. The rank-one recognition accuracies obtained for the above mentioned scenarios are: (a) 85.3%, (b) 63.9%, and (c) 87.4%. Figure 10 shows the corresponding CMC curves, along with the weights used for fusion in each case. From the results obtained, it can be observed that score-level fusion clearly improves the recognition performances when combining both inter-modality scores and intra-modality scores. The rank-one recognition performance obtained by the proposed approach (87.4%) reflects the highest recognition accuracy observed in the literature

¹¹Only these recognition performances should be considered when making a direct comparison with results from existing literature.

for this database. The rank-two recognition accuracy for the fusion scheme is observed to be 94.4%. This significant increase ($\sim 7\%$) in performance was due to the presence of duplicate entries, as described in Section 6. In such cases, a probe image would first match with the duplicate sample of the same subject (with different identification tag), and then with the the corresponding sample with the same identification tag. Such an effect causes a reduction in performance at rank-one. Some of the duplicate images that match at ranktwo but not at rank-one are shown in Table 4. If such duplicate images are accounted for (either removed, or given the same identification tags), a higher rank-one recognition performance can be expected. The benefit of the proposed technique can be observed in Table 5, showing example face and ocular images that were not correctly matched at rank-1 by the face recognition systems, but were correctly matched at rank-1 after performing fusion.



Figure 10. CMC curve showing the recognition accuracies obtained using score-level fusion of face scores, ocular scores, and a combination of the two.

7.4. Effect of individual surgeries

The effect of individual surgeries on the recognition performances was studied. Depending on the type of surgery performed, the images were categorized into two main groups: global and local [13]. Images corresponding to global surgeries show variations in the overall appearance of the face (e.g., rhytidectomy). Local surgeries, however, typically modify the appearance of a single facial feature, and may minimally impact the overall appearance of the face (e.g., otoplasty, rhinoplasty, etc.).

For this experiment, images corresponding to only major surgeries are considered. Images related to surgeries that do not provide clear information about which facial region they affect were excluded. For example, botox injections can be Table 4. **Duplicate image pairs that reduce the recognition per**formance at rank-one. Notice the difference in the identification tags, that causes the genuine pairs to be reckoned as impostors.

Input probe image	Corresponding gallery image that the probe has to match with	Instead matches with
ID # 03918 (after)	ID # 03918 (before)	ID # 13176 (before)
ID # 22517 (after)	ID # 22517 (before)	ID # 10228 (before)

Table 5. Example face and ocular image pairs (pre- and postsurgery) that were not correctly matched at rank-1 by the face recognition systems, but were correctly matched at rank-1 after performing fusion.



used to modify both local (say, around the lips), as well as the global appearance. Since the database does not provide meta-data that clearly explains these details, such images were excluded from this experiment. The rank-one recognition accuracies corresponding to individual surgeries obtained using face, ocular, and fusion schemes are provided in Table 6.

Table 6. Rank-one recognition accuracies corresponding to individual surgeries obtained using the face, ocular, and fusion schemes on images from face dataset B and ocular dataset.

Type of surgery	Face		Ocular		Proposed
	(VL)	(PP)	(SIFT)	(LBP)	
Browlift	88.2%	100%	64.7%	58.8%	97.0%
Otoplasty	85.4%	90.9%	69.0%	65.4%	94.5%
Blepharoplasty	74.2%	92.8%	64.2%	45.7%	94.2%
Rhinoplasty	79.1%	85.9%	54.3%	54.3%	85.9%
Rhytidectomy	78.8%	90.0%	48.4%	46.7%	92.2%

From the table, it can be observed that PittPatt and SIFT provide comparatively better face and ocular recognition performances, respectively. Once again, the proposed approach improves the recognition performance compared to individual techniques. Singh et al. [13] performed a similar study and concluded that face recognition algorithms cannot handle global facial plastic surgeries. Similar observation can be made from the results in this work. The recognition performance is more negatively impacted by global surgeries (rhinoplasty and rhytidectomy) than local surgeries (browlift, otoplasty, and blepharoplasty).

8. Conclusions and Future Work

This work proposes a fusion approach that combines the face and ocular information to improve biometric identification using images corresponding to facial plastic surgeries. The proposed approach yields a rank-one recognition accuracy of 87.4%, which quickly increases to 94.4% at rank-two. The performance obtained using the proposed approach reflects the current best rank-one accuracy reported on the considered plastic surgery database. Compared to existing approaches, the proposed scheme presents a method to improve recognition performance without using trainingbased methods. Based on the results, it is opined that the problem of face recognition using the publicly available plastic surgery database could be further improved if the non-ideal factors (e.g., duplicate entries, low image resolutions, etc.) of the database are accounted for. Future work would include an adaptive fusion scheme (face only, or a combination of face and ocular) for improved identification performance.

9. Acknowledgements

This work was funded by US NSF CAREER Grant No. IIS 0642554. The authors would like to thank Dr. Vinod Kulathumani and Sriram Sankar at West Virginia University for their support in using the PittPatt software. They are grateful to Dr. Mayank Vatsa, Dr. Richa Singh, and Himanshu S. Bhatt for their valuable comments [2].

References

- G. Aggarwal, S. Biswas, P. J. Flynn, and K. W. Bowyer. A Sparse Representation Approach to Face Matching Across Plastic Surgery. *IEEE Workshop on Applications of Computer Vision*, January 2012. 2
- [2] H. S. Bhatt, S. Bharadwaj, R. Singh, and M. Vatsa. Recognizing Surgically Altered Face Images using Multi-objective Evolutionary Algorithm. *IEEE Transactions on Information Forensics and Security*, 2012. 8
- [3] H. S. Bhatt, S. Bharadwaj, R. Singh, M. Vatsa, and A. Noore. Evolutionary Granular Approach for Recognizing Faces Altered due to Plastic Surgery. *IEEE International Conference on Automatic Face Gesture Recognition and Workshops*, pages 720–725, March 2011. 2
- [4] M. De Marsico, M. Nappi, D. Riccio, and H. Wechsler. Robust Face Recognition after Plastic Surgery using Local Region Analysis. *Proceedings of International Journal of Computer Vision*, 6754:191–200, 2011. 2
- [5] W. D. Jones. Plastic Surgery 1, Face Recognition 0. *IEEE Spectrum*, 46(9):17, September 2009. 2
- [6] D. G. Lowe. Distinctive Image Features from Scale Invariant Keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. 4
- [7] J. Merkow, B. Jou, and M. Savvides. An Exploration of Gender Identification using only the Periocular Region. pages 1–5, September 2010. 3
- [8] P. E. Miller, A. W. Rawls, S. J. Pundlik, and D. L. Woodard. Personal Identification using Periocular Skin Texture. *Proceedings of the 2010 ACM Symposium on Applied Computing*, pages 1496–1500, 2010. 3
- [9] T. Ojala, M. Pietikainen, and T. Maenpaa. Multiresolution Gray-scale and Rotation Invariant Texture Classification with Local Binary Patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(7):971–987, July 2002. 4
- [10] U. Park, R. R. Jillela, A. Ross, and A. K. Jain. Periocular Biometrics in the Visible Spectrum. *IEEE Transactions* on *Information Forensics and Security*, 6(1):96–106, March 2011. 3, 4
- [11] V. P. Pauca, M. Forkin, X. Xu, R. Plemmons, and A. Ross. Challenging Ocular Image Recognition. *Proceedings of SPIE Biometric Technology for Human Identification*, 8029, 2011. 3
- [12] A. Ross, R. R. Jillela, J. Smereka, V. N. Boddeti, B. V. K. VijayaKumar, R. Barnard, X. Hu, P. Pauca, and R. Plemmons. Matching Highly Non-ideal Ocular Images: An Information Fusion Approach. *Proceedings of the Fifth IAPR International Conference on Biometrics*, March 2012. 3, 4
- [13] R. Singh, M. Vatsa, H. S. Bhatt, S. Bharadwaj, A. Noore, and S. S. Nooreyezdan. Plastic Surgery: A New Dimension to Face Recognition. *IEEE Transactions on Information Forensics and Security*, 5(3):441–448, 2010. 1, 2, 5, 7, 8
- [14] P. Viola and M. Jones. Robust Real-time Face Detection. International Journal of Computer Vision, 57:137–154, 2004.
 4