Disguise Detection and Face Recognition in Visible and Thermal Spectrums

Tejas I. Dhamecha, Aastha Nigam, Richa Singh, and Mayank Vatsa
IIIT Delhi, India
{tejasd, aastha08001, rsingh, mayank}@iiitd.ac.in

Abstract

Face verification, though for humans seems to be an easy task, is a long-standing research area. With challenging covariates such as disguise or face obfuscation, automatically verifying the identity of a person is assumed to be very hard. This paper explores the feasibility of face verification under disguise variations using multi-spectrum (visible and thermal) face images. We propose a framework, termed as Anāvṛta[1], which classifies the local facial regions of both visible and thermal face images into biometric (regions without disguise) and non-biometric (regions with disguise) classes. The biometric patches are then used for facial feature extraction and matching. The performance of the algorithm is evaluated on the IIITD In and Beyond Visible Spectrum Disguise database that is prepared by the authors and contains images pertaining to 75 subjects with different kinds of disguise variations. The experimental results suggest that the proposed framework improves the performance compared to existing algorithms, however there is a need for more research to address this important covariate.

1. Introduction

Face recognition[2] has been an interesting area of research for more than five decades [29]. The pursuit to find the most accurate face representation and perform recognition has passed through shifts in the *paradigm* [29], as well as shifts in the *challenges* addressed. Earlier research has primarily focused on the covariates of pose, illumination and expression whereas recently, the challenges such as twins [17], face alterations due to plastic surgery [21], sketch-to-photo matching [3, 12], multi-spectrum matching [6, 11, 22], and disguise [18, 24] are also being explored.

Disguise is an interesting and a challenging covariate of face recognition. It involves intentional or unintentional changes on face through which one can either obfuscate his/her identity and/or impersonate someone else's identity.

[1] Anāvṛta is a Sanskrit word which means *uncovered*.

[2] In this paper, recognition and verification are interchangeably used.

Figure 1. Same person can appear differently by the use of disguise accessories. The images belong to a famous Indian actor, Amitabh Bachchan, with various disguise variations. Images are taken from Internet.

In either case, facial disguise falls under the broader category of *biometric obfuscation* [27]. Figure 1 shows an example of *face obfuscation*, where the appearance of a subject can vary by using different disguise accessories. (Note that the images in Figure 1 may be affected by covariates other than disguise, e.g. aging; however, in this work we are concentrating on disguise only). Disguise increases the *within-class* variation (when it is used to hide one's identity) and reduces the *between-class* variation (when it is used to impersonate someone else). Even though the problem of face recognition under disguise is very prevalent in real world applications, it has not been studied extensively. Table 1 presents a brief overview of research papers related to face recognition under disguise variations. Note that most of the research has been performed in visible spectrum using the AR [13] and Yale [2] face databases which contains very limited disguise (sunglasses and scarves only). Researchers have also suggested that the effect of occlusion cannot be circumvented by using only visible spectrum [16, 23, 24]. Though thermal imaging has been used for face recognition [6, 20, 22], to the best of our knowledge, its usefulness has not been explored for addressing disguise variations. Unlike single spectrum face recognition, [6, 22] provide motivation for utilizing visible and thermal spectrum together for improved face recognition.

To make face recognition more usable and secure, it is necessary to address the problem of (at least unintentional) disguise. As discussed earlier, it is difficult to address the problem using visible spectrum images alone. Therefore, this research attempts to combine the information obtained from both visible and thermal spectrum images to

Algorithm	Basic approach	Disguise detection	Disguise / Occlusion detected as	Face recognition	Spectrum	Database
Ramanathan *et al*. [18]	PCA	Yes	Left/right half face	Yes	Visible	National Geographic, AR
Singh *et al*. [24]	2D-log polar Gabor	No	-	Yes	Visible	AR, Private*, Synthetic Disguise^{+}
Martinez [14]	Probabilistic matching	No	-	Yes	Visible	AR
Wright *et al*. [25]	SRC	No	-	Yes	Visible	AR, Yale B [2]
Kim *et al*. [9]	ICA	No	-	Yes	Visible	AR, FERET
Yang and Zhang [26]	Gabor SRC	No	-	Yes	Visible	AR, Yale B
Pavlidis and Symosek [16]	-	Yes	Not explicitly	No	Near-IR	-
Yoon and Kee [28]	PCA + SVM	Yes	Upper/lower half	No	Visible	AR, Private$^{\times}$
Kim *et al*. [10]	PCA + SVM	Yes	Upper/lower half	No	Visible	AR, Private$^{\times}$
Choi and Kim [7]	AdaBoost + MCT-based features	Yes	Left-right eye, mouth	No	Visible	AR
Min *et al*. [15]	Gabor + PCA + SVM, LBP	Yes (Gabor + PCA + SVM)	Upper/lower half	Yes (LBP)	Visible	AR
Proposed	ITE, LBP	Yes (ITE)	Individual patches	Yes (LBP)	Visible and Thermal	I^{2}BVSD

Table 1. Existing algorithms for addressing disguise variations. AR database [13] contains 3200+ images pertaining to 126 subjects with two kinds of disguises (sunglasses and scarves). The National Geographic (NG) dataset contains 46 images of 1 individual, with various accessories such as hat, glasses, sunglasses, and facial hair. *Private dataset of 150 images pertaining to 15 individuals which contains similar real and synthetic disguise variations as in NG dataset. $^{+}$Synthetic disguise dataset of 4000 images pertaining to 100 individuals. $^{\times}$Private datasets are collected by researches in real world scenarios from ATM (automatic teller machine) kiosk. PCA, SRC, ICA, and SVM are abbreviations of principal component analysis, sparse representation classifier, independent component analysis, and support vector machines, respectively.

improve the performance of face recognition algorithms for disguise. The main contributions of the paper are: (1) a multi-spectrum framework to address the problem of disguise, (2) a descriptor for better encoding of disguise variations, (3) multi-spectrum face database with disguise variations, and (4) performance evaluation of the proposed technique and comparison with existing algorithms and a commercial off-the-shelf (COTS) face recognition system.

2. Anāvṛta: Framework for Recognizing Disguised Faces

Pavlidis and Symosek [16] have suggested that detecting disguise is necessary to efficiently recognize disguised faces. Inspired from [16], our hypothesis is:
"The facial part or patches which are under the effect of disguise (or occluded in most of the cases), are not only un-useful for face recognition, but may also provide misleading information. It is this misleading information that a person uses to hide his/her own identity and/or to impersonate someone else."
Building upon this intuition, we propose a framework, termed as Anāvṛta, for recognizing faces with variations in disguise. As illustrated in Figure 2, there are two cascaded stages in the proposed framework:

1. **Patch Classification**: It comprises dividing face image into patches and classifying them into *biometric* or *non-biometric* classes.

2. **Patch based Face Recognition**: Biometric patches are matched using local binary pattern (LBP) based face recognition algorithm.

2.1. Patch Classification

Several researchers have proposed patch or part-based face recognition in literature [1, 5] and evaluated the performance of individual parts for face recognition. To the

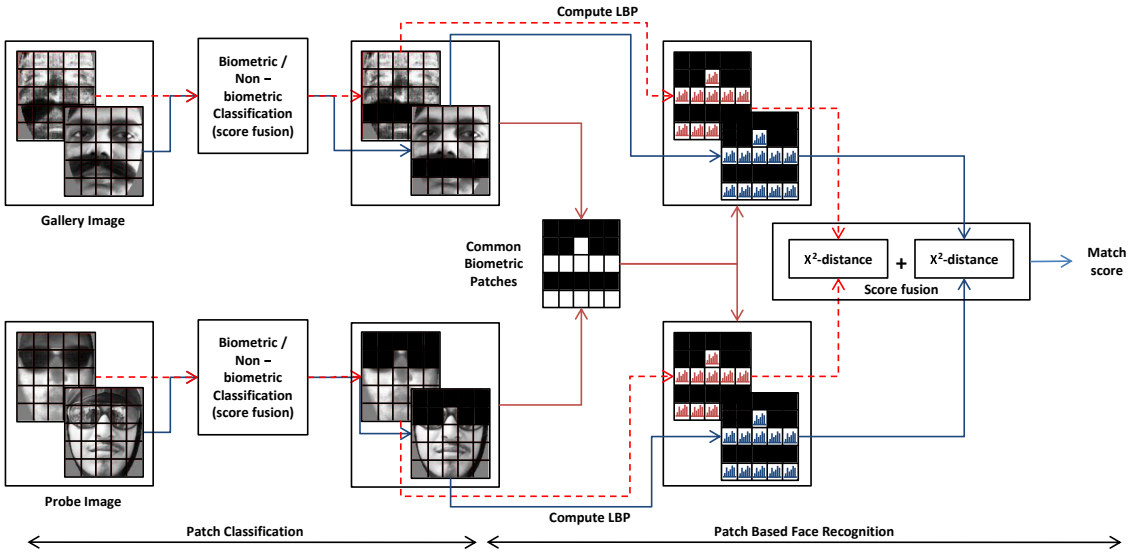

Figure 2. Illustrating the steps involved in the proposed multi-spectrum face recognition framework. The framework provides the flexibility to perform face recognition using either multi-spectrum or visible spectrum only.

best of our knowledge, [15] is the only work in literature which uses occlusion detection, to enhance the recognition performance. In applications such as law-enforcement, analyzing the patches to determine whether they are genuine facial regions or accessories is very important. The proposed patch classification algorithm therefore aims to classify the patches into biometric and non-biometric classes.

- **Biometric patches** are those facial parts that are not disguised; and hence they are useful for recognition.
- **Non-biometric patches/artifacts** are facial parts that are disguised. These patches may reduce the performance and should be avoided as far as possible.

Figure 3. Example of biometric and non-biometric patches. The first row contains biometric patches and the second row contains non-biometric patches. These patches are for illustration only, and are not part of our database.

Figure 3 shows some examples of biometric and non-biometric patches. The proposed patch classification algorithm has two steps: feature extraction and classification.

2.1.1 ITE Feature Extraction

It is our assertion that some of the non-biometric patches or occlusions, such as hair and artificial nose, can be distinguished using texture information, while some others, such as scarves and sunglasses, can be distinguished using their intensity values. Therefore, the proposed algorithm uses a concatenation of texture and intensity descriptors as input feature. As shown in Figure 2, the algorithm starts with tessellating the face image. Input face image I is first divided into non-overlapping rectangular patches I_{ij}, $1 \leq i \leq m, 1 \leq j \leq n$, where m and n are the number of horizontal and vertical patches respectively. The intensity and texture descriptors are computed for all the patches using the intensity histogram and Local Binary Patterns (LBP) algorithm [1] respectively. The proposed descriptor is termed as the *Intensity and Texture Encoder* (ITE). For a patch ij of an image I, ITE is defined as

$$\mathbf{E}(I_{ij}) = [intensityHist(I_{ij}); lbpHist(I_{ij})] \quad (1)$$

where $intensityHist(\cdot)$ represents the histogram of an intensity image and $lbpHist(\cdot)$ represents the LBP histogram. We use basic LBP operator with 8 sampling points, that produces 256 dimensional feature vector for each patch. Intensity histogram consists of 256 bins, resulting in a feature vector of the same dimension.

2.1.2 ITE Feature Classification

The ITE features can, potentially, be classified using any of the generative or discriminative classification techniques. Our observation of biometric and non-biometric patches shows that the set of biometric patches is well defined and can be modeled efficiently. However, due to the variety of accessories that can be used for disguise, non-biometric patches have an exhaustive population set which is difficult

to model using a limited training database. Therefore, in this research, we have proposed the use of Support Vector Machine (SVM) [8], a discriminative classifier, for classifying biometric and non-biometric patches.

An SVM model is learned from the ITE descriptors of all the patches from training images (which are annotated manually). This model is used to classify the patches from the testing data. For every patch, a decision score s is computed using SVM. A patch is classified as biometric, if the score is less than the threshold T, i.e. $s < T$; and if score is equal to or greater than threshold, i.e. $s \geq T$, the patch is classified as non-biometric. Accordingly, a flag variable F_{ij} is generated, which represents whether the patch is classified as biometric or non-biometric. The flag value of every patch is then combined to generate the flag matrix, $\mathbf{F}_{m\times n} = [F_{ij}]_{1\leq i\leq m, 1\leq j\leq n}$, using Eq. 2.

$$F(I)_{ij} = \begin{cases} 1 & \text{if } I_{ij} \text{ is classified as biometric} \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

This classification algorithm can be utilized for both single and multiple spectrum images. In single spectrum mode, ITE features of images patches are classified using trained SVM. We propose to employ patch classification in multi-spectrum mode, where patch classification scores of both the spectrums can be combined using sum rule fusion [19]. The fused score, s_f, is computed as the weighted average of the patch classification scores of individual spectrums, i.e. $s_f = (w_v s_v + w_t s_t)/(w_v + w_t)$, where w_i represent the weights and s_i represent the classification score in each spectrum. In this work we have assumed $w_v = w_t$.

2.2. Patch based Face Recognition

Let I^p be the probe image which is to be matched with the gallery image I^g. The corresponding flag matrices $\mathbf{F}(I^p)$ and $\mathbf{F}(I^g)$ are generated using Eq. 2. Here, it is possible that for some gallery patch, I^g_{xy}, which is classified as biometric, the corresponding probe patch, I^p_{xy}, is classified as non-biometric. In other words, $F(I^g)_{xy} = 1$ and $F(I^p)_{xy} = 0$, or $F(I^g)_{xy} = 0$ and $F(I^p)_{xy} = 1$. This renders the particular patch of gallery image not useful for recognition because the corresponding patch from the probe image is under disguise effect and matching a biometric patch with a non-biometric patch may lead to incorrect information.

$$\mathbf{F}^u(I^p, I^g) = \mathbf{F}(I^p) \wedge \mathbf{F}(I^g) \quad (3)$$

The patch classification algorithm explained in Section 2.1 classifies the patches into biometric and non-biometric, and Eq. 3 provides information that *for a given gallery-probe pair, which patches should be used for face recognition*. Note that, in order to take advantage of patch classification, the face recognition approach has to be patch-based. Therefore, we propose to use LBP [1] which is one of the widely used patch-based descriptors for face recognition. If $desc^I_{ij}$ represents the LBP descriptor of ij patch of image I, and the χ^2-distance between two LBP descriptors is represented as $dist(\cdot,\cdot)$, then the distance between two images, I^p and I^g, is calculated as:

$$Dist(I^p, I^g) = \frac{1}{\eta}\sum_i \sum_j dist(desc^{I^p}_{ij}, desc^{I^g}_{ij})\mathbf{F}^u(I^p, I^g)_{ij}$$
$$\text{where } \eta = \sum_i \sum_j \mathbf{F}^u(I^p, I^g)_{ij} \quad (4)$$

and $\mathbf{F}^u(I^p, I^g)_{ij}$ is calculated using Eq. 3. As we propose to use multiple spectrums, the distances are computed for each spectrum individually and then fused using sum ruleross2006multimodal. So, $Dist(I^p, I^g)$ in multiple spectrum mode is defined as,

$$Dist(I^p, I^g) = Fusion(Dist(VI^p, VI^g), Dist(TI^p, TI^g))$$

subject to Eq. 4, where VI and TI represent the visible and thermal spectrum images respectively. Since both the spectrums encode different kinds of information, it is our assertion that fusion of visible and thermal spectrums should result in improved recognition performance.

3. Database

To the best of our knowledge, there is no publicly available face database that contains multi-spectrum (visible and thermal) images with disguise variations. Further, the visible spectrum databases generally used for disguise related research (AR [13] and Yale [2] face databases) contain very limited disguise variations, such as scarves and/or sunglasses. Therefore, to evaluate the effectiveness of the proposed algorithms, we have prepared the IIITD In and Beyond Visible Spectrum Disguise (I^2BVSD) face database[3] of disguised/obfuscated face images. The database contains visible and thermal spectrum images of 75 participants with disguise variations. The number of images per person varies from 6 to 10. For every subject, there is at least one frontal neutral[4] face image and at least five frontal disguised face images. For each spectrum, there are 681 images. The visible spectrum images are captured using Nikon D-90 and thermal images are captured using a thermal camera having micro-bolometer sensor operating at $8 - 14\mu m$. The size of visible images is 4288×2848, while the thermal images are of size 720×576 . All the face images are captured under (almost) constant illumination with neutral expressions and frontal pose. The disguise variations included in the database are categorized into the following categories. (1) **Without disguise**: neutral image, (2) **Variations in hair styles**: different styles and colors of wigs, (3) **Variations**

[3]The database will be publicly available to the research community.
[4]Face images without any disguise is referred as *neutral* face image.

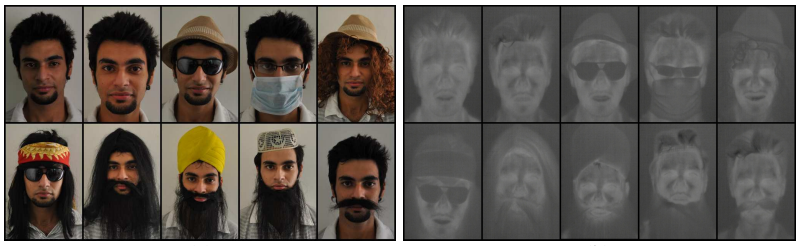

(a) (b)

Figure 4. Sample images from the I^{2}BVSD database (a) images captured in visible spectrum, and (b) corresponding images acquired in thermal spectrum.

due to beard and mustache: different styles of beard and mustaches, (4) **Variations due to glasses**: sunglasses and spectacles, (5) **Variations due to cap and hat**: different kinds of caps, turbans, veil (also known as hijab which covers hair), and bandanas, (6) **Variation due to mask**: disposable doctor's mask, and (7) **Multiple variations**: A combination of disguise accessories.

Figure 4 shows a sample of both visible and thermal images from the database. The disguises are chosen in such a way that they result in more realistic appearances and (almost) every part of the face is hidden at least once. The subjects are asked to disguise themselves using the given accessories. This allows different subjects to have different types of disguises thus providing more variations across individuals in the database.

4. Experiments and Results

This section demonstrates the results of the proposed face recognition framework which includes the patch classification algorithm and LBP based face recognition. The experiments are performed on the I^{2}BVSD face database. Eye coordinates are manually annotated to avoid any errors in face detection. Face images are first preprocessed using the CSU Face Identification Evaluation System [4] to obtain normalized images, which includes: (a) locating left and right eye at fixed positions, (b) resizing target face image to 130×150 pixels, (c) applying an elliptical face mask to remove the background information, (d) converting the image to gray scale, (e) histogram equalization, and (f) pixel normalization so that, the average of all pixel intensities is 0 and standard deviation is 1. The result section is divided into three subsections: patch classification using ITE, evaluation of the proposed algorithm, and comparison with sparse representation classifier (SRC) and COTS.

4.1. Patch Classification using ITE

All the images in the database are divided into 5×5 non-overlapping rectangular patches of size 26×30 pixels. Every patch is manually annotated as biometric or non-biometric to create the ground truth for training as well as evaluation. If more than half of the patch is covered with accessories, it is annotated as a non-biometric patch. Images of randomly chosen 35 subjects form the training set and the images from the remaining 40 subjects are used for testing. The training set thus contains 8050 patches (322 images$\times$25 patches) in each band, out of which 6324 patches are biometric and 1726 patches are non-biometric. Similarly, the testing set comprises of 8975 patches (359 images$\times$25 patches) amongst which 6944 are biometric and 2031 are non-biometric. Figure 5 shows the distribution of (annotated) biometric patches in the training and testing splits. Notice that there are face images with different amounts (in terms of number of biometric patches) of disguises.

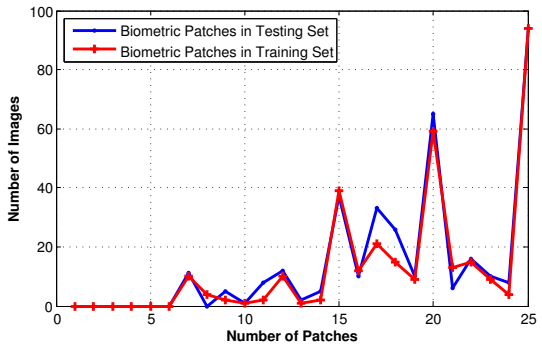

Figure 5. The distribution of biometric patches in the training and test sets.

As explained earlier, for each patch, the ITE features are computed using Eq. 1; and min-max normalization is performed to map the values in the interval $[-1, 1]$. The normalized descriptor is provided as input to SVM with Radial Basis Function kernel for patch classification. The kernel parameter and misclassification cost are estimated using grid search along with 5-fold cross validation. In grid search, parameters that provide the maximum training accuracy are considered as optimum.

Since ITE is a concatenation of LBP and intensity values, the efficacy of ITE is compared with LBP and pixel intensity values. LBP histograms, intensity histograms, and ITE histograms are computed and provided as input to SVM separately for classification. The performance of all three

histograms is evaluated for (a) visible spectrum images, (b) thermal images, and (c) feature fusion (concatenation) of descriptors extracted from visible and thermal images. Receiver Operating Characteristics (ROC) curves for patch classification representing the results of these experiments are shown in Figure 6.

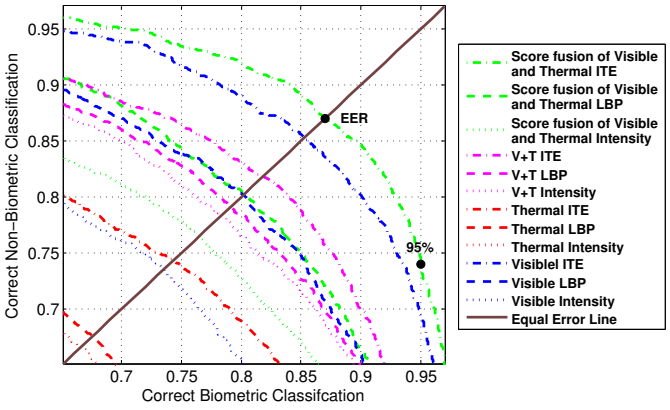

Figure 6. ROC curves for patch classification. Here, (V+T) represents feature level fusion (concatenation) of descriptors.

In case of both visible and thermal spectrums, ITE provides better results compared to texture or intensity information for patch classification. This supports our hypothesis that *concatenation of texture and intensity should yield better patch classification results*. The proposed score level fusion of both the bands outperforms all other approaches. It is also observed that visible spectrum performs better than thermal, this may be due to the fact that thermal images contain less texture information. Feature level fusion performed by concatenation of visible and thermal ITE, does not increase the accuracy.

4.2. Performance Evaluation of Proposed Framework

The output of patch classification yields biometric patches which are used for feature extraction and matching. For evaluating the proposed face matching approach, the testing set is divided into two parts: gallery and probe. For each subject, one neutral face image, and four other randomly selected images are taken as gallery and the remaining images constitute the probe/query set. Hence, there are total 200 gallery images and 159 probe images. We have performed experiments with five random cross validation trials. The experiments are performed in verification mode and the results are reported in terms of ROC curve and verification accuracy at 0.1%, 1.0% and 10% False Accept Rate (FAR). The performance of the proposed approach (multi spectrum score fusion) is compared with the performance of individual spectrums. Note that such an evaluation is important to understand the importance of fusion over individual spectrums. For each of these three scenarios (visible only, thermal only, and fusion), we performed following three experiments.

1. Face recognition with biometric patches is classified using ITE and SVM classifier
2. Face recognition with manually annotated biometric patches
3. Face recognition with all the patches (normal LBP approach)

The results of face recognition are shown in Figure 7. In both the spectrums individually (left and middle graphs), for FAR>1%, using only ground truth biometric patches results in better accuracy than using all the patches for face recognition. The performance of the proposed framework depends significantly on the performance of the patch classification algorithm. Intuitively, rejecting a non-biometric patch is less benefitting than the loss incurred by wrongly rejecting a biometric patch. From the ROC curve of patch classification shown in Figure 6, it can be analyzed that at EER, 13% of the biometric patches are being misclassified. Since, the number of biometric patches provided as input to the face recognition algorithm reduces, the ROC curves in Figure 7 show that the performance of face recognition reduces when the threshold of patch classification is chosen at EER. However, for 95% correct biometric patch classification (Figure 6), even though, the number of correctly classified non-biometric patches decreases, the face recognition algorithm is receiving more biometric patches as input and the results show that at the threshold pertaining to 95% biometric patch classification, the proposed face recognition framework yields the best performance for FAR>0.1%. This support our hypothesis that not using non-biometric patches for recognition can result in better accuracy. As shown in Figure 7(c), using the threshold obtained at 95% correct biometric patch classification accuracy, the proposed score fusion of visible and thermal spectrums further improves the verification accuracy.

4.3. Comparison with COTS and Sparse Representation Classification

In this section, we present the comparison with FaceVacs, a commercial-off-the-shelf system (referred as COTS) and sparse representation classification (SRC) [25]. Note that, SRC is one of the most important techniques in literature for addressing occlusion/disguise. In SRC, the residual is considered as the dissimilarity measure of the gallery-probe pair. In case of SRC, for fusion of two spectrums, sum rule fusion of residuals of both individual spectrums is performed. Since, COTS is designed to work with visible spectrum images only, all the results pertaining to COTS are computed for visible spectrum images only.

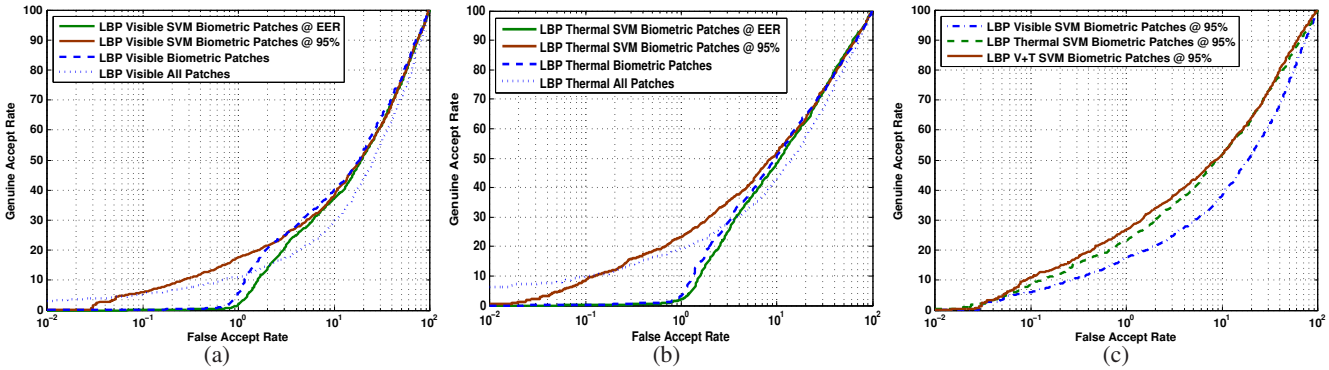

Figure 7. The results of the proposed face recognition framework using LBP descriptor (a) on visible spectrum images, (b) on thermal spectrum images, and (c) sum fusion of both spectrums under three different cases of patch classification scenarios.

For evaluating the performance of the proposed framework, we have utilized all the gallery and probe images irrespective of the information content or image quality. However, the COTS used in this research has inbuilt algorithms for quality assessment and enrollment. The thresholds for enrolling a gallery image are very strict whereas for probe images, it is relaxed. Out of the 200 gallery images, COTS registered approximately 60% of the gallery images and the remaining images were considered as *failure to enroll* whereas all the probe images were processed successfully. It is also observed that if the face image does not contain any non-biometric patch, then the probability of getting enrolled in the COTS is higher. However, for a fair comparison, we have overridden the COTS to include all 200 images in the gallery.

Figure 8 and Table 2 demonstrate the results of COTS, SRC in individual spectrums, and score fusion of SRC in visible and thermal spectrums. Comparisons with the corresponding experiments of proposed approach are also shown. It shows that COTS is not able to classify the faces under disguises very well. Moreover, it provides evidence of challenging nature of dataset itself. Note that, for lower FAR($<$0.05%), all approaches shown in comparison exhibit very poor performance. From approximately 0.2% till 5% FAR, the verification rate of COTS changes from 16% to 20% GAR. This is attributed to COTS rejecting many samples while recognition. For the same range of FAR, the proposed approach yields up to 45% GAR. In case of SRC, at 1.0% FAR, score fusion improves the performance by 5-7% GAR.

As shown in Table 2, although the performance reported by the proposed approach is not as high as it is usually reported in face recognition literature, it outperforms one of the state-of-art commercial system, and a widely used technique, SRC, by at least 3% at 1.0% FAR.

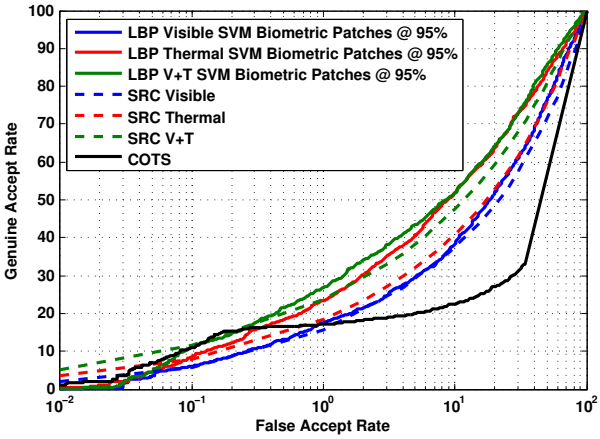

Figure 8. Comparison of COTS and SRC with the proposed approach on visible, thermal spectrum and their score fusion.

Approach	Verification accuracy @ FAR		
	0.1%	1.0%	10%
SRC(V)	5.6 ± 1.3	15.5 ± 1.6	37.7 ± 1.8
SRC(T)	7.8 ± 1.1	18.4 ± 1.3	40.7 ± 1.9
SRC(V+T)	11.7 ± 1.3	23.7 ± 1.7	47.5 ± 1.9
COTS	10.9 ± 2.4	17.1 ± 1.5	22.5 ± 1.2
LBP(V)	6.0 ± 0.5	17.5 ± 0.1	38.3 ± 0.1
LBP(T)	8.3 ± 0.8	23.3 ± 1.0	51.7 ± 2.0
LBP(V+T)	$\mathbf{10.9 \pm 1.0}$	$\mathbf{26.8 \pm 1.2}$	$\mathbf{51.7 \pm 1.6}$

Table 2. Genuine accept rates and their standard deviations at different false accept rates of the proposed approach along with comparison to COTS and SRC.

5. Conclusion and Future Work

This research focuses on addressing the covariate of disguise in face verification using multi-spectrum images. The major contributions of this research are two folds. We designed a framework, Anāvṛta, which is based on the observation that *artifacts are not part of biometric information*

of face and they should not be used for recognition. The framework consists of the ITE based patch classification (in biometric/non-biometric classes) and LBP based face recognition applied on classified biometric patches. The second contribution is creation of a multi-spectrum face disguise database, I^{2}BVSD database, to encourage research in this domain. The proposed approach improves verification accuracy over direct application of LBP, COTS, and SRC. It is also observed that for both, patch classification and face recognition, score fusion of visible and thermal spectrum match scores leads to improvements. It is our assertion that the I^{2}BVSD database will boost the research in face recognition with disguise variations. Though Anāvṛta provides improved performance over existing state-of-the-art algorithms, it is still an open research problem. As a future research direction, we plan to approximate the set of disguise artifacts and use human knowledge for better recognition.

6. Acknowledgement

A. Nigam, R. Singh, and M. Vatsa were partially supported through a research grant from Department of Information Technology, Government of India. T. Dhamecha was partially supported through TCS Research.

References

[1] T. Ahonen, A. Hadid, and M. Pietikäinen. Face recognition with local binary patterns. In *ECCV*, pages 469–481, 2004. 2, 3, 4

[2] P. Belhumeur and D. Kriegman. The Yale face database. *URL: http://cvc. yale. edu/projects/yalefaces/yalefaces. html*, 1997. 1, 2, 4

[3] H. Bhatt, S. Bharadwaj, R. Singh, and M. Vatsa. On matching sketches with digital face images. In *IEEE BTAS*, pages 1–7, 2010. 1

[4] D. Bolme, J. Ross Beveridge, M. Teixeira, and B. Draper. The CSU face identification evaluation system: its purpose, features, and structure. *Computer Vision Systems*, pages 304–313, 2003. 5

[5] J. Chen, S. Shan, C. He, G. Zhao, M. Pietikainen, X. Chen, and W. Gao. Wld: a robust local image descriptor. *IEEE TPAMI*, 32(9):1705–1720, 2010. 2

[6] X. Chen, P. J. Flynn, and K. W. Bowyer. Ir and visible light face recognition. *CVIU*, pages 332–358, 2005. 1

[7] I. Choi and D. Kim. Facial fraud discrimination using detection and classification. *Advances in Visual Computing*, pages 199–208, 2010. 2

[8] C. Cortes and V. Vapnik. Support-vector networks. *ML*, 20(3):273–297, 1995. 4

[9] J. Kim, J. Choi, J. Yi, and M. Turk. Effective representation using ICA for face recognition robust to local distortion and partial occlusion. *IEEE TPAMI*, 27(12):1977–1981, 2005. 2

[10] J. Kim, Y. Sung, S. Yoon, and B. Park. A new video surveillance system employing occluded face detection. In M. Ali and F. Esposito, editors, *Innovations in Applied Artificial Intelligence*, volume 3533 of *LNCS*, pages 65–68. 2005. 2

[11] B. Klare and A. Jain. Heterogeneous face recognition: matching NIR to visible light images. In *IEEE ICPR*, pages 1513–1516, 2010. 1

[12] B. Klare, Z. Li, and A. Jain. Matching forensic sketches to mug shot photos. *IEEE TPAMI*, 33(3):639–646, 2011. 1

[13] A. Martinez. The AR face database. *CVC Technical Report*, 24, 1998. 1, 2, 4

[14] A. Martinez. Recognizing imprecisely localized, partially occluded, and expression variant faces from a single sample per class. *IEEE TPAMI*, 24(6):748–763, 2002. 2

[15] R. Min, A. Hadid, and J. Dugelay. Improving the recognition of faces occluded by facial accessories. In *IEEE F&G Workshop*, pages 442–447, 2011. 2, 3

[16] I. Pavlidis and P. Symosek. The imaging issue in an automatic face/disguise detection system. In *IEEE Workshop CVBVS*, pages 15–24, 2000. 1, 2

[17] P. Phillips, P. Flynn, K. Bowyer, R. Bruegge, P. Grother, G. Quinn, and M. Pruitt. Distinguishing identical twins by face recognition. In *IEEE F&G*, pages 185–192, 2011. 1

[18] N. Ramanathan, R. Chellappa, and A. Roy Chowdhury. Facial similarity across age, disguise, illumination and pose. In *IEEE ICIP*, volume 3, pages 1999–2002, 2004. 1, 2

[19] A. Ross, K. Nandakumar, and A. Jain. *Multimodal Biometrics*. International Series on Biometrics. Springer Science+Business Media, LLC, 2006. 4

[20] A. Selinger and D. A. Socolinsky. Appearance-based facial recognition using visible and thermal imagery: a comparative study. Technical Report 02-01, Equinox Corporation, February, 2002. 1

[21] R. Singh, M. Vatsa, H. Bhatt, S. Bharadwaj, A. Noore, and S. Nooreyezdan. Plastic surgery: a new dimension to face recognition. *IEEE TIFS*, 5(3):441–448, 2010. 1

[22] R. Singh, M. Vatsa, and A. Noore. Hierarchical fusion of multi-spectral face images for improved recognition performance. *Information Fusion*, 9(2):200–210, 2008. 1

[23] R. Singh, M. Vatsa, and A. Noore. Recognizing face images with disguise variations. In M. G. K, Delac and M. S. Bartlett, editors, *Recent Advances in Face Recognition*, pages 149–160. 2008. 1

[24] R. Singh, M. Vatsa, and A. Noore. Face recognition with disguise and single gallery images. *IVC*, 27(3):245–257, 2009. 1, 2

[25] J. Wright, A. Yang, A. Ganesh, S. Sastry, and Y. Ma. Robust face recognition via sparse representation. *IEEE TPAMI*, 31(2):210–227, 2009. 2, 6

[26] M. Yang and L. Zhang. Gabor feature based sparse representation for face recognition with gabor occlusion dictionary. In *ECCV*, pages 448–461, 2010. 2

[27] S. Yoon, J. Feng, and A. Jain. Altered fingerprints: analysis and detection. *IEEE TPAMI*, 34(3):451–464, 2012. 1

[28] S. M. Yoon and S.-C. Kee. Detection of partially occluded face using support vector machines. In *MVA*, pages 546–549, 2002. 2

[29] W. Zhao, R. Chellappa, P. Phillips, and A. Rosenfeld. Face recognition: a literature survey. *ACM Computing Surveys*, 35(4):399–458, 2003. 1