A Coding Scheme for Indexing Multimodal Biometric Databases

Aglika Gyaourova
West Virginia University
Morgantown WV 26506, USA
agyaouro@csee.wvu.edu

Arun Ross
West Virginia University
Morgantown WV 26506, USA
http://www.csee.wvu.edu/~ross

Abstract

In biometric identification systems, the identity associated with the input data is determined by comparing it against every entry in the database. This exhaustive matching process increases the response time of the system and, potentially, the rate of erroneous identification. A method that narrows the list of potential identities will allow the input data to be matched against a smaller number of identities. We describe a method for indexing large-scale multimodal biometric databases based on the generation of an index code for each enrolled identity. In the proposed method, the input biometric data is first matched against a small set of reference images. The set of ensuing match scores is used as an index code. The index codes of multiple modalities are then integrated using three different fusion techniques in order to further improve the indexing performance. Experiments on a chimeric face and fingerprint bimodal database indicate a 76% reduction in the search space at 100% hit rate. These results suggest that indexing has the potential to substantially improve the response time of multimodal biometric systems without compromising the accuracy of identification.

1. Introduction

Multibiometric systems use multiple sources of biometric information to establish the identity of an individual. By adopting multibiometric systems the recognition performance can be improved. The use of multiple sources of information also ensures that the system is harder to forge and the failure to enroll is less probable [18].

The potential of deploying biometric systems has led to the creation of large multibiometric databases. Existing programs that use large multibiometric databases include the US-VISIT program for border security based on the face and fingerprint modalities [3], and the FBI's fingerprint database that currently stores the ten-print information of more than 80 million distinct individuals [4]. The UK National Identity Scheme [1], when enacted, will create a database storing the face, iris, and fingerprint data of every ID card holder in the UK.

At least two problems arise in large-scale biometric databases: (1) searching the database to retrieve an identity can be slow because the input data has to be compared (matched) against the biometric data of every identity in the database, and (2) in most biometric systems, the false accept error grows with the size of the database [13]. (So far, only iris recognition systems have been demonstrated to have no false match errors under certain conditions [6].) Therefore, *filtering* the database in order to reduce the number of potential candidates (i.e., identities) for matching is a desirable component of any large-scale biometric system.

Filtering a biometric database can be accomplished using two distinct schemes: classification and indexing. A classification scheme partitions the database into several classes. The class of the input data is first estimated and, subsequently, the input is compared only against those identities in the database belonging to that class. The main limitation of classifying biometric data is the unbalanced distribution of the identities across the various classes. This problem exists in the traditional Henry fingerprint classification system [12] as well as techniques for face [19], palmprint [20], and iris classification [21].

Indexing schemes, on the other hand, assign an index value to every biometric template. However, the indices of two biometric images pertaining to the same identity are unlikely to be the same because the process of data acquisition and processing is subject to noise. Therefore, indexing systems retrieve those identities whose indices are similar to the index value of the input data. The input image is matched only against the retrieved identities thereby reducing the identification time and, potentially, the identification error rate. Most of the methods for indexing fingerprint images are based on image features such as minutiae triplets [8] and ridge curvature [7, 17]. Indexing of iris images is often based on iris codes [10, 16]. Many fast face retrieval techniques operate in the image domain or use the projection coefficients from a subspace analysis method [14].
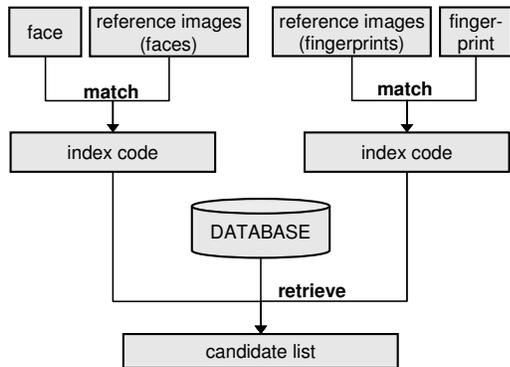
Figure 1. Indexing using two modalities. Two index codes are generated separately using a single modality. The retrieval process uses the two index codes to find candidate identities.

The techniques mentioned above either use features implicit in the biometric image or statistics extracted from the biometric feature vector. To avoid the complexity of designing new feature extraction routines, we propose an indexing scheme based on match scores. In a related study by Maeda *et al*. [15], filtering is performed based on match scores. The input image is sequentially matched against images in the database. The sequence of matching is dictated by the similarity between the match scores obtained for the input and the corresponding scores for the images in the database. For this purpose, a matrix that contains the pairwise match scores of *all* images in the database has to be permanently stored and updated for each new enrolled identity. The technique of Maeda *et al*. is novel in employing match scores, therefore eliminating the need to perform additional image processing for filtering purposes. However, storing the matrix of match scores for a database containing millions of images can be prohibitive.

Multiple biometric traits can be used to speed-up the filtering process and also to enhance its robustness. Robustness is achieved when the modest performance of one modality is compensated by the good performance of another modality. Thus, the probability of discarding the correct identity during filtering is reduced. In a previous study by Hong *et al*. [11], the identification results obtained by using different modalities were combined hierarchically to achieve a filtering effect. The number of candidate identities is first reduced by using modalities that allow for faster matching. The final identification is then performed in the reduced search space by using a different modality that has a better matching accuracy.

In this paper, we propose a method for indexing multimodal biometrics databases (face and fingerprint) based on index codes generated by a biometric matcher. The proposed indexing technique relies on the use of a *small* set of reference images for each modality. A modality-specific in-

dex code is generated by matching an input image against these reference images, resulting in a set of match scores. During identification, the index code of the input image is compared against the index codes of the enrolled identities in order to find a set of potential matches. The output of multiple modalities are fused to further narrow down the list of candidates and increase the hit rate, thereby resulting in a more robust and efficient system. Our indexing approach relies on a matcher, which is an inherent part of every automated biometric identification system. The generated index codes are compact and so the proposed method has modest storage requirements.

2. Indexing method

In this section we use the face modality as an example. However, the inferred properties are applicable to the fingerprint modality (as observed in our experiments) and perhaps for other biometric modalities as well.

2.1. Creating Index Codes for a Single Modality

Let $\mathcal{R} = \{r_1, r_2, ..., r_n\}$ be an ordered set of face images, which we call *reference images*, and $\mathcal{S}_x(\mathcal{R}) = \{s(x, r_1), s(x, r_2), ..., s(x, r_n)\}$ be the set of match scores obtained when face image x is compared to each reference image in $\mathcal{R}$. ($s(x, y)$ is the match score obtained from comparing image x against image y.) We refer to $\mathcal{S}_x(\mathcal{R})$ as the *index code* of image x. If two images x and y belong to the same identity then their index codes $\mathcal{S}_x(\mathcal{R})$ and $\mathcal{S}_y(\mathcal{R})$ are likely to be strongly correlated (i.e., as measured by Pearson's correlation coefficient). In contrast, if the images x and y belong to different individuals, then the correlation between their index codes is likely to be weak or negative.

The correlation coefficient between index codes can be used to identify a potential list of candidates from the database. During identification, the indexing system first computes the index code, $\mathcal{S}_x$, of the input x. Then it outputs every enrolled identity, y, for which the correlation coefficient between $\mathcal{S}_y$ and $\mathcal{S}_x$ exceeds a specific threshold value.

Algorithm for retrieving candidate identities.

Let $\mathcal{S}_x$ be the index code of the input image, $\mathcal{S}_y$ be the index code of an image y from the database, and T be a predefined threshold.

1. Compute Pearson's correlation coefficient using $\rho(x, y) = \frac{Cov(\mathcal{S}_x, \mathcal{S}_y)}{[Var(\mathcal{S}_x)Var(\mathcal{S}_y)]^{1/2}}$ for all y.
2. Output those y for which $\rho(x, y) > T$.

When using multiple modalities, the index codes are generated separately for each modality and combined during the retrieval process (Figure 1). When k modalities are

available, the architecture of the proposed indexing scheme is defined by the k ordered sets of reference images (one set for each modality) and the k thresholds that specify the minimum value of the correlation coefficient needed to include an enrolled identity in the candidate list. The number of reference images for each modality may be different. The k index codes for each enrolled identity are stored in the database and used during the retrieval process. Thus, our indexing approach operates using only the information available in the match scores. The proposed method has modest storage requirements because the generated index codes are compact.

2.2. Computation Time

The retrieval process performs an exhaustive search across the index codes of the enrolled identities, which may raise questions about the contribution of the indexing to the overall response time of the biometric system. In fact, an increase in the speed of identification can be achieved only if the search space is substantially reduced and if the correlation coefficient between two index codes can be computed in a fraction of the time needed to match two biometric templates. In the case of face and fingerprints, which are the modalities commonly present in large biometric databases, matching two templates is slower than computing the correlation coefficient between two vectors (codes).

Let P be the fractional reduction in the number of candidate identities achieved by the indexing scheme when applied on a database of size M. Let n denote the dimensionality of the index code. Then, the overall computation time of the identification system can be approximated by the sum of the n matching operations between the input image and the reference images, the M operations for computing the correlation coefficients, and the $P * M$ matching operations required for identification. Similarly, the time needed for identification without indexing consists of only M matching operations. If t_m is the time needed to perform a single matching operation and t_ρ is the time needed to compute a single correlation coefficient, we are interested in determining the values of n, M, t_m, and t_ρ that will reduce the overall response time, *i.e.*, we determine the conditions under which the following inequality holds:

$$M\, t_m > n\, t_m + P\, M\, t_m + M\, t_\rho \,. \tag{1}$$

Let $\alpha = t_m/t_\rho$. Since the number of operations required for matching two biometric templates (at least for face and fingerprints) can be a magnitude larger than the computation of the correlation coefficient, we can assume for now that $\alpha = 10$. If we also assume that $n = 256$ (as used in our experiments) and $P = 1/2$, then

$$M\, t_m \;>\; 256\, t_m + \frac{M}{2}\, t_m + M\, \frac{t_m}{\alpha} \tag{2}$$

$$\Rightarrow M \;>\; 640\,. \tag{3}$$

Therefore, for databases having over 640 identities, the use of the proposed indexing method will reduce the identification time. In addition, when the candidate list represents a smaller fraction of the database (i.e., $P < 1/2$), indexing will be beneficial for even smaller databases.

2.3. Selecting reference images

The number of reference images and the degree of diversity among them are important considerations for good indexing performance. A larger number of reference images should ensure better indexing performance but will also increase the computation time. Therefore, this number can be chosen empirically according to the size of the database, and the desired accuracy and speed-up. Another important factor is the choice of reference images. A diversity of reference images is needed to ensure a wide range of index codes. We do not rely on implicit image features to ensure diversity. Instead, we select the images whose impostor match scores exhibit large variances, (i.e., the *max-variation* rule):

Algorithm for selecting n reference images: *max-variation* rule.

Let $\boldsymbol{F} = \{f_1, f_2, ..., f_q\}$ be the candidate pool of reference images, and $s(x, y)$ be the match score between images x and y.

1. For each f_i, compute $v_i = Var\{s(f_i, f_j)\}_{j=1, j\neq i}^{L}$.
2. Let V be the list of sorted v_i values in descending order.
3. Use the images corresponding to the top n values in V as reference images.

Another diversity measure that we tested was based on the mean value of the impostor scores (i.e., the *max-mean* rule). The *max-mean* rule is analogous to the *max-variation* rule, the only difference being that the Var operator is replaced by the mean operator.

In the current implementation, the reference images are selected from the database itself (although they could also be some type of digital artifacts, such as synthetic fingerprints, faces, etc.). While the entire database can be viewed as a candidate pool for selecting reference images, practical considerations dictate the use of a small random subset of images for this purpose.

2.4. Fusing Index Codes of Multiple Modalities

Different combination (or fusion) techniques are appropriate for achieving different goals. Database retrieval techniques have a tradeoff between a low miss rate and a large reduction in search space. Fusion schemes are often helpful in achieving one of those goals.

2.4.1 Concatenation of index codes

Let $S_x(\mathcal{R}^i) = \{s(x, r_1^i), s(x, r_2^i), ..., s(x, r_n^i)\}$ be the index code of image x from modality i (r_j^i is the j^{th} reference image of the i^{th} modality). The fused index code, $\mathcal{F}_x$, is obtained by concatenating the index codes from different modalities: $\mathcal{F}_x = \{s(x, r_1^1), ..., s(x, r_n^1), s(x, r_1^2), ..., s(x, r_n^2)\}$. Using the fused index code, the indexing proceeds as in the case of a single modality.

This fusion scheme results in longer index codes and a larger diversity among the reference images. When the number of reference images, n, is the same in each modality, the length of the fused index codes will be $n*k$ but only the top n reference images from each modality will be used to generate index codes. When the number of reference images increases, the added diversity may decrease (since the reference images are sampled from a finite set of images). The weakness of this fusion scheme is that poor indexing performance in one of the modalities can negatively affect the overall indexing performance since the two index codes are forcibly concatenated.

2.4.2 Union of candidate lists

Let C^i be the set of retrieved identities according to modality i. The final set of identities, $\mathcal{C}$, retrieved from the indexing will be $\mathcal{C} = \bigcup_{i=1}^{k} C^i$ (i.e., the identities retrieved by each modality are combined to obtain the final set of candidates). This fusion scheme has the potential to eliminate errors in the candidate list introduced by individual modalities. Thus, the poor indexing performance of one modality will not affect the overall indexing performance. This fusion system fails only when all the modalities perform poorly.

2.4.3 Intersection of candidate lists

Similar to the union fusion scheme, here the final fused output is the intersection of the candidate lists of the individual modalities, or $\mathcal{C} = \bigcap_{i=1}^{k} C^i$. This type of fusion can further reduce the size of the search space. However, to achieve such a reduction, the indexing performance of multiple modalities have to be comparable (i.e., performance should vary little across modalities).

3. Databases

We performed experiments using 2024 frontal images (2 images per subject) from the FERET database [2] and 1740 fingerprint images (2 images per subject) taken from the West Virginia University multimodal database [5]. For every identity in these databases, one of the two images was used to compose the database while the other was used as
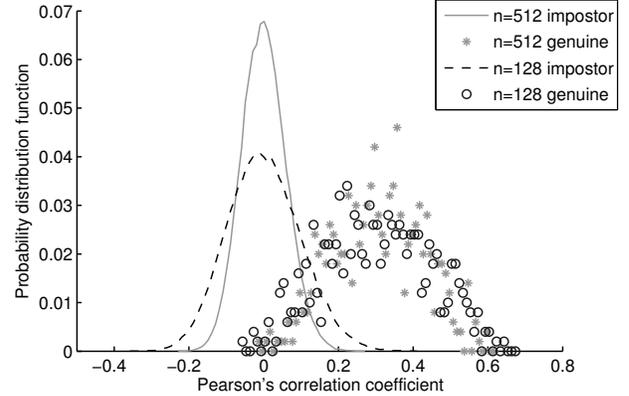

Figure 2. Distribution of the correlation coefficients between index codes belonging to the same identity (genuine class) and different identities (impostor class). Increasing the dimensionality of the index codes (n) results in a better separation between the genuine and the impostor classes.

the input image for testing. Match scores were produced by the commercial face and fingerprint matchers developed by Neurotechnology.

We selected 256 reference images from the database using the *max-variation* rule. After removing the reference images from the databases and reducing the size of the face database to that of the fingerprint database, a chimeric multimodal database consisting of 614 identities was created in order to evaluate the performance of the proposed indexing scheme.

4. Experiments

We evaluated the indexing performance by using the conventional hit rate and penetration rate measures. The hit rate represents the percentage of test inputs for which the corresponding identity is correctly retrieved from the database:

$$Hit\ rate = \frac{N_h}{N}, \qquad (4)$$

where N_h is the number of tests in which the correct identity is present in the candidate list and N is the total number of tests. The penetration rate is the average reduction in the search space achieved by the indexing scheme:

$$Penetration\ rate = \frac{1}{N} \sum_{i=1}^{N} \frac{L_i}{M}, \qquad (5)$$

where L_i is the number of identities in the candidate list of the i^{th} input image and M is the number of identities in the database. In our experiments $N = M$.

The distribution of correlation coefficients between the index codes of face images from different people (i.e., the impostor class) is centered around zero (Figure 2). On the
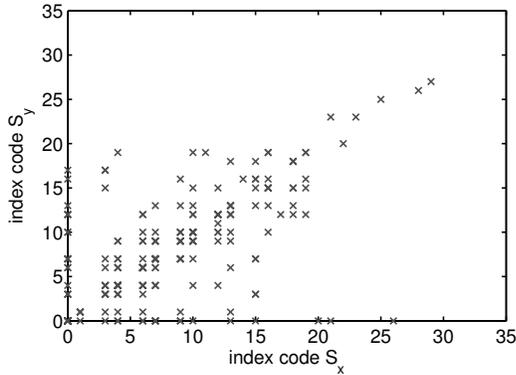
Figure 3. Scatter plot of the corresponding components of two index codes belonging to the same identity.
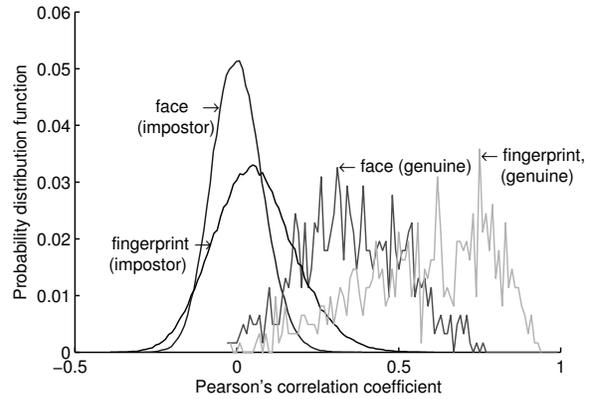

Figure 5. Distributions of correlation coefficients for the face and the fingerprint modality (n=256).
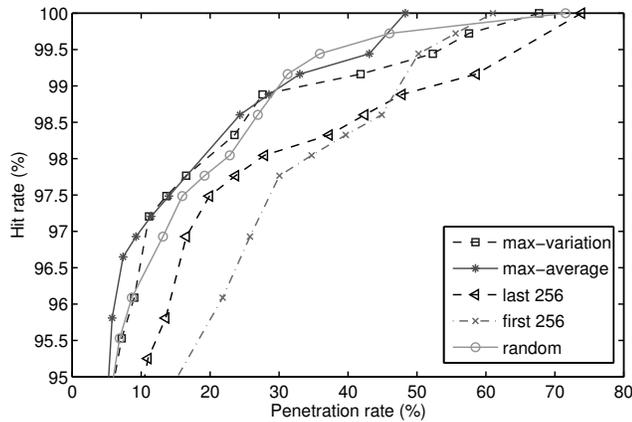

Figure 4. Different rules for selecting reference images lead to different indexing performances. The "first 256" and the " last 256" correspond to simply using the first and the last images in the database without applying any selection. See the text for a description of the other selection rules.
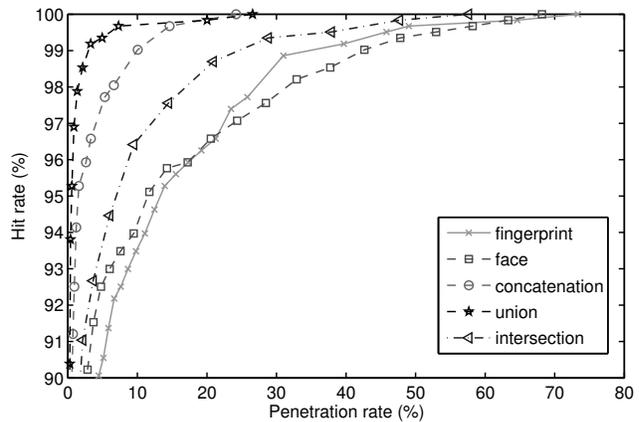

Figure 6. Indexing results using index codes of dimensionality 256. The performance using multiple modalities is substantially better than the performance using any single modality.

other hand, correlations between index codes of face images from the same identities (the genuine class) has a mean greater than zero. Furthermore, the separation between the genuine and the impostor distributions improves as the size of the index code increases (i.e., by increasing the number of reference images). The use of a correlation coefficient to perform indexing is also supported by examining the scatter plot of the corresponding elements of pairs of index codes belonging to the same identity (Figure 3).

The *max-variation* and the *max-mean* rules led to superior indexing performance compared to a random selection of the reference images or when no selection is employed (Figure 4).

Because the genuine and impostor distributions of the correlation coefficients differ for different modalities (Figure 5), care should be taken when selecting the decision

threshold for each modality (T^i). A poorly chosen threshold has a negative impact on the performance of the Union and Intersection fusion rules. The thresholds are expected to be different across biometric modalities and databases. We used standard deviation distances to select the threshold for each modality, where the means and the standard deviations were estimated empirically from the distribution of the impostor class.

Our experiments indicated that the Union and the Concatenation fusion rules resulted in the best performance (Figure 6). The Concatenation fusion rule achieved the best penetration rate (24%) at a hit rate of 100%. Its penetration rate decreases slowly when miss errors are allowed (10% penetration rate at 99% hit rate). In contrast, the penetration rate of the Union fusion rule is 27% at a 100% hit rate but decreases rapidly down to 3% at a 99% hit rate. These two fusion methods achieved substantial reduction of the size of the search space compared to using a single modality. (In-

dexing using either the fingerprint and face modalities results in a 70% penetration rate at a hit rate of 100%.) These results demonstrate that the performance of indexing improves substantially when combining multiple modalities.

5. Discussion and future work

We proposed an indexing technique for multimodal biometric databases and showed its effectiveness in reducing the search space during identification. Thus, the proposed scheme is capable of reducing the response time of biometric identification systems. Our technique only relies on the availability of a matcher and can be incorporated into any biometric system without the need to implement additional image processing algorithms. The proposed indexing scheme is universal and is applicable to any type of multibiometric system, such as those using multiple classification algorithms, multiple biometric traits, or different samples of the same biometric trait (*e.g.* left and right index fingers).

The scatter plot of the index codes of two images that belong to the same identity, Figure 3, indicates that the Pearson's correlation coefficient may not be the best choice since the correlation for small match scores is weak. Correlation through ranking — Spearman's rank correlation coefficient or Kendall tau rank correlation coefficient — may provide better results. Preliminary results using rank correlation have shown better performance than Pearson's correlation coefficient. However, rank correlation is slower and may not be appropriate when rapid response time is more important than a reduced search space.

To speed-up the retrieval process and reduce the storage requirements, a discretization function can be applied to the index codes [9] and an appropriate similarity function used to facilitate the retrieval of identities. In the extreme case, binary index codes and the Hamming distance can be utilized to perform indexing; this can result in rapid identification even for small databases.

The Union fusion rule achieved a 3% penetration rate when 99% of the input images were indexed correctly. We are currently investigating the cause for the increase in the penetration rate to 27% for the remaining 1% of the images in the database. We are also testing our scheme on databases having a larger number of identities.

Acknowledgments

References

[1] British National Identity Scheme. http://www.ips.gov.uk/identity/.

[2] The FERET database. http://www.frvt.org/FERET/.

[3] US-VISIT program. http://www.dhs.gov/xprevprot/programs/gc_1208531081211.shtm.

[4] FBI press release, October 10, 2008. http://www.fbi.gov/page2/oct08/iafis_101008.html.

[5] S. Crihalmeanu, A. Ross, S. Schuckers, and L. Hornak. A protocol for multibiometric data acquisition, storage and dissemination. Technical report, WVU, LDCSEE, 2007.

[6] J. Daugman. The importance of being random: statistical principles of iris recognition. *Pattern Recognition*, 36(2):279–291, February 2003.

[7] J. Feng and A. Cai. Fingerprint indexing using ridge invariants. *ICPR*, 4:433–436, 2006.

[8] R. S. Germain, A. Califano, and S. Colville. Fingerprint matching using transformation parameter clustering. *IEEE Computational Science & Engineering*, 4(4):42–49, 1997.

[9] A. Gyaourova and A. Ross. A novel coding scheme for indexing fingerprint patterns. In *Proc. Joint IAPR International Workshop on Structural, Syntactic, and Statistical Pattern Recognition*, pages 755–764. Springer-Verlag, 2008.

[10] F. Hao, J. Daugman, and P. Zielinski. A fast search algorithm for a large fuzzy database. *IEEE Trans. Information Forensics and Security*, 3(2):203–212, June 2008.

[11] L. Hong and A. Jain. Integrating faces and fingerprints for personal identification. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 20:1295–1307, 1998.

[12] Henry Classification system. International Biometric Group, 2003. http://www.biometricgroup.com/Henry Fingerprint Classification.pdf.

[13] A. Jain and S. Pankanti. *Automated fingerprint identification and imaging systems*, pages 275–326. Advances in Fingerprint Technology, 2nd Edition, (H. C. Lee and R. E. Gaensslen). CRC Press, 2001.

[14] K.-H. Lin, K.-M. Lam, X. Xie, and W.-C. Siu. An efficient human face indexing scheme using eigenfaces. *Proc. International Conference on Neural Networks and Signal Processing*, 2:920–923, Dec. 2003.

[15] T. Maeda, M. Matsushita, and K. Sasakawa. Identification algorithm using a matching score matrix. *IEICE Trans. Information and Systems, Special Issue on Biometric Person Authentication*, 84(7):819–824, 2001.

[16] R. Mukherjee and A. Ross. Indexing iris images. In *International Conference on Pattern Recognition*, 2008.

[17] A. Ross and R. Mukherjee. Augmenting ridge curves with minutiae triplets for fingerprint indexing. *Biometric Technology for Human Identification IV*, 6539C(1):65390C, 2007.

[18] A. Ross, K. Nandakumar, and A. Jain. *Handbook of Multibiometrics*. International Series on Biometrics, Vol. 6. Springer, 2006.

[19] T. Vikram, K. Chidananda Gowda, D. Guru, and S. Urs. Face indexing and retrieval by spatial similarity. *Congress on Image and Signal Processing*, 1:543–547, May 2008.

[20] X. Wu, D. Zhang, K. Wang, and B. Huang. Palmprint classification using principal lines. *Pattern Recognition*, 37(10):1987–1998, 2004.

[21] L. Yu, D. Zhang, K. Wang, and W. Yang. Coarse iris classification using box-counting to estimate fractal dimensions. *Pattern Recognition*, 38(11):1791–1798, 2005.