

CONTRIBUTIONS TO BIOMETRIC RECOGNITION:
MATCHING IDENTICAL TWINS AND LATENT FINGERPRINTS

By

Alessandra Aparecida Paulino

A DISSERTATION

Submitted
to Michigan State University
in partial fulfillment of the requirements
for the degree of

Computer Science– Doctor of Philosophy

2013

ABSTRACT

CONTRIBUTIONS TO BIOMETRIC RECOGNITION: MATCHING IDENTICAL TWINS AND LATENT FINGERPRINTS

By

Alessandra Aparecida Paulino

Automatic recognition of a person by the use of distinctive physical and behavioral characteristics is called biometrics. Two of the important challenges in biometrics are recognition of identical twins and matching of latent fingerprints to their exemplar prints (rolled or slap prints). The contributions of this dissertation are focused on these two topics.

Identical (monozygotic) twins are a result of a single fertilized egg that splits into two cells, each one giving birth to one individual. Identical twins have the same deoxyribonucleic acid (DNA), thus their genotypic features (features influenced by the genetic material) are the same. However, some of their phenotypic features (features influenced by the fetal environment) may be different. Thus, it is essential to determine which biometric traits (either by themselves or in combination with other traits) have the ability to distinguish identical twins and the extent of their ability for this discrimination.

The first contribution of this dissertation is an evaluation of the performance of biometric systems in the presence of identical twins for the three most commonly used biometric modalities, namely fingerprint, face and iris. Identical twins are shown to be a challenge to current face recognition systems. On the other hand, fingerprint and iris matching of identical twins show performance comparable to those with unrelated persons. The fusion of different samples from the same modality of a subject (e.g., left and right iris, fingerprints of multiple fingers) yields the best matching performance for identical twins, similar to what has been shown for unrelated persons. Biometric traits can also be used to determine whether two subjects enrolled in a biometric database are identical twins. By using face and iris modalities together, for example, we can correctly iden-

tify 80% of such identical twin pairs, while only 2% of subject pairs who are not identical twins will be incorrectly considered identical twins.

The second contribution of this work is focused on improving latent fingerprint matching performance. Latent fingerprints are partial fingerprint images that typically contain only a small area of friction ridge pattern and large non-linear distortion, are blurred or smudged, and contain complex background noise. Due to these characteristics, latents are a particularly challenging for matching to their mates (reference prints) in a database. Given a latent print in which minutiae have been marked by a human expert (as is the current practice in forensics), we have proposed two approaches to improve the latent matching performance. The first approach consists of enhancing the latent image and fusing the matching score obtained from the enhanced latents with the score based on manually marked minutiae. This approach outperforms a commercial fingerprint matcher on the public latent database NIST SD 27. The second approach consists of developing a latent fingerprint matcher that utilizes minutiae as well as the orientation field information. The proposed matching algorithm outperforms three fingerprint matching algorithms on two different latent fingerprint databases (NIST SD 27 and WVU latent databases).

The latent fingerprint identification accuracy generally deteriorates as the size of the fingerprint database grows. To alleviate this problem and to reduce the overall search time of latent matching, we propose to combine various level 1 and level 2 features, including minutia cylinder code, minutia triplets, singular points, orientation field and ridge period, to efficiently filter out a large portion of the reference fingerprint database. The proposed approach outperforms state-of-the-art indexing techniques on the public domain latent database NIST SD27 against a large background database of 267K rolled prints. The experimental results also show that the proposed filtering scheme has the desirable property of attaining improved computational efficiency of latent search (20% penetration rate) while maintaining the latent matching accuracy.

To my loved ones.

ACKNOWLEDGEMENTS

I consider myself a very fortunate person for having so many people to be thankful to. Some were part of the beginning of my Ph.D. “journey”, some were there even before, some became part of it recently, and some were there from the beginning to the point we are now. Each and all of them helped me in one way or another and I am grateful to all of them for being there for me. There were a few difficult and frustrating moments for sure, but there was never a moment that I regretted coming here and pursuing my Ph.D.. I am very thankful for all the wonderful people I have met and for everything I have learned.

I should start by thanking Dr. Anil K. Jain, who has been a great advisor and mentor over the five years I have worked under his supervision. I admire him immensely as a researcher, a teacher, an advisor, and as a person as well. We can always go to him with a problem or a question, because he is always willing to help. He says maybe he worries about us more than ourselves; this is just the type of advisor he is, caring and always looking out for us and helping us achieve our best. He always goes the extra mile to help us in any way he can; this is just the type of person he is. I have learned so much from him that no thank you will ever be enough.

I am also grateful to have such a great thesis committee. Thank you Dr. Selin Aviyente, Dr. Xiaoming Liu and Dr. Arun Ross for evaluating my work and providing valuable comments and suggestions. Thank you also to Dr. Pang-Ning Tan and Dr. George Stockman for being part of the committee in the early stages.

An important part of being able to write this dissertation is due to all members of the fingerprint group at PRIP lab. Our weekly meetings, their valuable suggestions and discussions helped me not only with my work directly, but also with becoming a better researcher in general.

I would like to thank Jianjiang Feng, with whom we had a very active work collaboration and from whom I also learned a lot. Thank you Heeseung, Eryun and Kai for revising this dissertation. I am grateful for the administrative support received from the CSE office at MSU, including Linda

Moore, Cathy Davison, Norma Teague and Debbie Kruch, and for the computer support received from Kelly Climer and Adam Pitcher. Thank you so much for your patience and support.

I cannot express how much I have learned from fellow Pripies, both in academic and personal levels. We are friends in the journey and I value all the time we spent together, all the lessons I learned from you, and all the help you provided. Thank you Soweon, Jung-Eun, Brendan, Serhat, Kien, Qijun, Eryun, Kai, Abhishek, Pavan, Radha, Sunpreet, Unsang, Hu, Lacey, Scott, Charles. A special thank you to Soweon for all the discussions and help; she has been someone I could always count on, from helping me find a small bug in my code to discussing life. She has become a great friend and a person who I admire and respect personally and professionally.

I would like to thank my husband for being so understanding and patient over the course of my Ph.D.. His unconditional support has helped me go through the tough times, and his help has kept me balanced. Having him by my side gave me strength and it is wonderful to be able to share my life with him. I can always count on him: from cooking by himself so that I can have more time to work, to listening and talking to me so that I can figure things out and move forward.

I also would like to thank my parents for being so supportive of everything, every decision, every step, every journey. They have always been there for me and I am so grateful. The peace of mind that comes from knowing that I can count on them if I need to has made my life so much easier and smoother. I grew up in a wonderful family, and it was not easy to be far from them, but they were always very supportive and understanding. I am grateful to my sisters who, despite the distance, were always there to share things and laugh together. And I have missed all of them so much, especially my nephew, who was only four years old when I came to the U.S. and now is almost as tall as his mother. I am also grateful to my extended family (in-laws), who has been very supportive throughout this journey, and also very understanding of my husband and I being so far from them.

And last but not least, I would like to thank all my friends for everything. I am very fortunate to have friends to count on, laugh with and learn together. It is always wonderful to be with them.

Thank you to everyone who was a part of this journey. Your support made this possible.

TABLE OF CONTENTS

LIST OF TABLES	x
LIST OF FIGURES	xi
CHAPTER 1 INTRODUCTION	1
1.1 Biometric Systems	5
1.2 Common Biometric Traits	8
1.2.1 Fingerprint Representation and Matching	10
1.2.1.1 Representation	10
1.2.1.2 Matching	16
1.2.1.3 Indexing	18
1.2.1.4 Latent Fingerprints	19
1.2.2 Face Representation and Matching	22
1.2.2.1 Representation	22
1.2.2.2 Matching	22
1.2.3 Iris Representation and Matching	25
1.2.3.1 Representation	26
1.2.3.2 Matching	27
1.3 Applications of Biometric Systems	28
1.4 Contributions	29
1.4.1 Distinguishing Identical Twins	30
1.4.2 Latent Fingerprint Matching	31
1.4.3 Latent Fingerprint Indexing	32
1.5 Dissertation Organization	33
CHAPTER 2 BIOMETRIC TRAITS OF IDENTICAL TWINS	34
2.1 Introduction	34
2.2 Related Work	37
2.3 Multibiometrics	43
2.3.1 Fingerprint Discriminability	44
2.3.2 Face Discriminability	44
2.3.3 Iris Discriminability	46
2.4 Experimental Results and Analysis	47
2.4.1 Databases	47
2.4.2 Distinguishing Identical Twins	49
2.4.2.1 Fingerprint	49
2.4.2.2 Face	51
2.4.2.3 Iris	53
2.4.2.4 Multibiometric Experimental Results	54
2.4.2.5 Discussion	57
2.4.3 Finding Similarities between Identical Twins	60

2.5	Summary and Conclusions	66
CHAPTER 3 LATENT FINGERPRINT MATCHING		68
3.1	Introduction	68
3.2	Related Work	74
3.2.1	Rolled Fingerprint Matching	74
3.2.2	Latent Fingerprint Matching	75
3.2.3	NIST Evaluation of Latent Fingerprint Technologies	76
3.2.4	Evaluation of Latent Examiners	77
3.3	Latent Fingerprint Features	78
3.3.1	Orientation Field Estimation	78
3.3.1.1	Minutiae based Orientation Field Reconstruction	80
3.3.2	Local Minutia Descriptor	84
3.3.2.1	Minutia Cylinder Code (MCC)	84
3.4	Latent Matching Approaches	86
3.4.1	Latent Matching with Enhanced Image	86
3.4.1.1	Rank-level Fusion	88
3.4.1.2	Score-level Fusion	89
3.4.2	Latent Matching with Descriptor-Based Hough Transform	90
3.4.2.1	Alignment	91
3.4.2.2	Similarity Measure	95
3.5	Databases	101
3.5.1	NIST Special Database 27 (NIST SD27)	101
3.5.2	West Virginia University Latent Database (WVU LFD)	102
3.6	Experimental Results	104
3.6.1	Latent Matching with Enhanced Images	104
3.6.1.1	Fusion methods	105
3.6.2	Latent Matching with Descriptor-Based Hough Transform	110
3.6.2.1	Commercial Matchers	110
3.6.2.2	Alignment Performance	110
3.6.2.3	Matching Performance	114
3.6.2.4	Fusion of Matchers	119
3.6.2.5	Effect of Fingerprint Quality	123
3.6.2.6	Computational Cost	125
3.6.2.7	Comparison of the two proposed methods	125
3.7	Summary and Conclusions	131
CHAPTER 4 LATENT FINGERPRINT INDEXING		132
4.1	Introduction	132
4.2	Related Work	135
4.3	Feature Extraction	140
4.3.1	Reference Prints	140
4.3.2	Latents	142
4.4	Indexing Approach	143
4.4.1	Triplets	146

4.4.2	Minutia Cylinder Code (MCC) Indexing	148
4.4.3	Orientation Field Descriptor Indexing	149
4.4.4	Singular Points	150
4.4.5	Ridge Period	151
4.4.6	Fusion of Indexing Scores	152
4.5	Experimental Results	153
4.5.1	Databases	153
4.5.2	Indexing Results	154
4.5.3	Implementation Details	162
4.6	Summary	163
CHAPTER 5 SUMMARY AND FUTURE WORK		165
5.1	Summary	165
5.2	Contributions	167
5.3	Future Work	168
BIBLIOGRAPHY		171

LIST OF TABLES

Table 1.1	Comparison of the characteristics of the three most commonly used biometric traits from [1]. (H = high, M = medium, L = low)	8
Table 1.2	False reject and false acceptance rates associated with fingerprint, face and iris verification systems.	9
Table 2.1	Summary of studies on the biometrics of identical twins. Sets can include identical twin pairs as well as non-identical twin pairs.	38
Table 2.2	Equal Error Rate (%) for distinguishing (i) genuine vs. impostor identical twins and (ii) genuine vs. general impostors based on different biometric features.	57
Table 3.1	Comparison of rank-1 accuracies reported in the literature for the NIST SD27 database.	119
Table 3.2	Rank-1 accuracies for various subjective quality levels of latents in NIST SD27.	123
Table 3.3	Rank-1 accuracies for various objective quality values of latents in NIST SD27 (large, medium and small refer to the number of minutiae in the latent).	123
Table 3.4	Rank-1 accuracies for various objective quality values of latents in WVU LFD (large, medium and small refer to the number of minutiae in the latent).	124
Table 3.5	Rank-1 accuracies for latents grouped according to NFIQ quality values of corresponding rolled prints in NIST SD27.	124
Table 3.6	Rank-1 accuracies for latents grouped according to NFIQ quality values of corresponding rolled prints in WVU LFD.	124
Table 4.1	Summary of studies on fingerprint indexing for rolled, plain and latent prints.	137
Table 4.2	Computation time for individual indexing modules.	163

LIST OF FIGURES

Figure 1.1	Illustration of Bertillon system for person identification based on numerous body measurements, which was used in the United States and in Europe from the end of the 19th century to the beginning of the 20th century. From left to right and then top to bottom the figures show measurement of height, reach, trunk, length of head, width of head, right ear, left foot, left middle finger, and left forearm.	2
Figure 1.2	Examples of intra-class variability and inter-class similarity. (a) Two different impressions of the same finger of the same subject and (b) face images of a pair of identical twins. For interpretation of the references to color in this and all other figures, the reader is referred to the electronic version of this dissertation.	5
Figure 1.3	An overview of a biometric system.	6
Figure 1.4	Three types of fingerprint images.	11
Figure 1.5	Examples of fingerprint features belonging to the three levels: (a) orientation field and singular points and (b) minutiae shown in fingerprints from NIST Special Database 4 [2], and (c) sweat pores shown on part of a rolled print image from the WVU Latent Database [3].	11
Figure 1.6	Examples of Level 1 features in fingerprints. (a) Orientation field and (b) singular points (core in red and delta in green) shown in rolled fingerprint images from NIST Special Database 27 [4].	12
Figure 1.7	Example fingerprints belonging to each of the five major fingerprint classes (fingerprints from NIST Special Database 4 [2].)	14
Figure 1.8	An illustration of how minutia direction is defined for (a) minutia ending and (b) minutia bifurcation [5].	15
Figure 1.9	Example of matched minutiae from a fingerprint pair (latent and rolled print).	17
Figure 1.10	Examples of latent fingerprints in NIST SD 27 database [4].	20
Figure 1.11	Examples of facial features at the three levels: (a) facial geometry (Level 1), (b) facial landmarks, Gabor filter, local binary pattern (LBP), and shape (Level 2) and (c) moles (Level 3) [5].	23
Figure 1.12	An illustration of a graph model fitted to a face [5].	24

Figure 1.13	Schematic diagram of (a) SIFT [6] and (b) LBP [5] descriptor construction.	25
Figure 1.14	Examples of two irises of two different subjects from the CASIA twins database. The image on the right (b) is not of very high quality since the iris is partially occluded by hair.	26
Figure 1.15	An overview of an iris recognition system [5].	27
Figure 1.16	A pair of identical twins from CASIA Twins database Version 2.	31
Figure 2.1	Face images of (a) a pair of identical (monozygotic) twins and (b) a pair of non-identical (dizygotic) twins from the University of Notre Dame ND-Twins database [7].	35
Figure 2.2	A diagram of the human eye ¹	46
Figure 2.3	The kiosk for biometric acquisition (a) and the face acquisition device (b).	48
Figure 2.4	Fingerprint images of fingers 1, 2, 3, and 4 of the first twin (a), and the four images of the corresponding fingers of the second twin in an identical twin pair (b); similarly, (c) and (d) show fingerprint images of a non-identical twin pair. Note the similarity in ridge flow pattern between identical twins. All four corresponding fingers of identical twins in (a) and (b) have the same pattern type. But for non-identical twins in (c) and (d), only two corresponding fingers (no. 1 and 3) have the same pattern type.	50
Figure 2.5	Face images of the first and second twin in (a) an identical twin pair, and (b) a non-identical twin pair.	52
Figure 2.6	The left and right iris images of identical ((a) and (b)) and non-identical twin pairs ((c) and (d)).	54
Figure 2.7	Identical twin, general impostor distributions and genuine distributions for fingers 3 and 4, face, and right iris.	55
Figure 2.8	Identical twin, general impostor, and genuine distributions of the 4-finger fusion and 2-iris fusion, respectively.	56
Figure 2.9	ROC curves for fingerprint, iris and face, and a fusion example (face + finger 4). Due to the small number of identical twin impostor (102), FAR less than 1/102 cannot be estimated.	58
Figure 2.10	Examples of face images in CASIA Twins V2 database.	60
Figure 2.11	ROC curves for face using FaceVacs 8.6.	61
Figure 2.12	ROC curves for twin detection using iris and face.	63

Figure 2.13	ROC curves for twin detection using fingerprint and face.	64
Figure 2.14	Face images of pairs of (a) identical and (b) non-identical twins with extremely high face match scores considering the range of scores is $[0, 1]$ using FaceVacs 8.6.)	65
Figure 3.1	Three types of fingerprint impressions. Rolled and plain fingerprints are also called reference fingerprints.	69
Figure 3.2	An example of a tenprint card containing two types of impressions (rolled and plain) of the ten fingers. The first portion of the card contains personal information related to the person being fingerprinted, such as name, date of birth, place of birth, gender, race, height, weight, etc. The second portion of the card contains impressions of the ten fingers. The first and second rows of impressions consist of rolled impressions of the fingers from the left and right hands, respectively (from left to right: thumb, index, middle, ring, little). The third row contains plain impressions of the left four fingers taken simultaneously, left thumb, right thumb, and right four fingers taken simultaneously.	70
Figure 3.3	Examples of estimated orientation field ((b) and (d)) by gradient-based approach [8] for a (a) latent and (c) its mated rolled print in NIST SD27.	73
Figure 3.4	Latent fingerprints of three different quality levels in NIST SD27.	76
Figure 3.5	Illustration of the orientation field computed at a small block.	79
Figure 3.6	Local ridge orientation estimation based on the nearest minutiae in each sector.	80
Figure 3.7	Comparison of orientation field estimation methods: (a) Original latent fingerprint image, (b) orientation field estimated directly from the given image, (c) orientation field estimated from minutiae and singular points, and (d) manually marked orientation field.	82
Figure 3.8	Comparison of orientation field estimation methods: (a) Original latent fingerprint image, (b) orientation field estimated directly from the given image, (c) orientation field estimated from minutiae and singular points, and (d) manually marked orientation field.	83
Figure 3.9	Sections of two cylinders associated with the two corresponding minutiae, one in latent and other in rolled print.	85
Figure 3.10	Overview of the latent matching with enhanced image approach.	87

Figure 3.11	Examples of enhanced latent images. (a) and (e) are original latent images, (b) and (f) are the skeletons of the original latent images extracted by VeriFinger 4.2 SDK, (c) and (g) are enhanced images using reconstructed orientation field, and (d) and (h) are the skeletons of the enhanced images extracted by VeriFinger.	88
Figure 3.12	Overview of the proposed latent matching approach.	91
Figure 3.13	Illustration of the MCC descriptor-based Hough Transform alignment. ²	93
Figure 3.14	Final score computation scheme combining minutiae and orientation field information.	96
Figure 3.15	Latent print for which the matching score of the genuine pair is slightly reduced when a smaller threshold value is used compared to a larger threshold value, while impostor score is greatly reduced. (a)-(c) show the latent, the true mate, and the rank-1 non-mate according to large threshold, respectively. (d)-(g) show latent minutiae that were matched to rolled print minutiae in the following cases: (d) true mate using small threshold, (e) true mate using large threshold, (f) non-mate using small threshold, and (g) non-mate using large threshold. In (d)-(g), the scores corresponding to each case are included. Filled-in minutiae indicate the matching minutiae.	98
Figure 3.16	Latent print identified at a higher rank after fusing minutiae matching scores with orientation field matching scores. The rank of the true mate was improved from 2 to 1 after the fusion, and the rank of the highest ranked non-mate was 3 after the fusion. (a)-(c) show minutiae and the image of (a) a latent, (b) its true mate, and (c) the highest ranked non-mate according to minutiae matching. (d) and (f) show latent minutiae and orientation field (in blue) aligned with minutiae and orientation field of the true mate. (e) and (g) show latent minutiae and orientation field (in blue) aligned with minutiae and orientation field of the rank-1 non-mate.	100
Figure 3.17	A latent and its corresponding rolled print in the WVU latent database. The NFIQ quality [9] of the rolled print is 4, which is one above the worst NFIQ quality.	102
Figure 3.18	Histograms of NFIQ values of rolled prints in NIST SD27 and WVU databases.	103
Figure 3.19	CMC curves for manually marked minutiae, enhanced image using reconstructed orientation field, automatically extracted minutiae by Verifinger from latent image and enhanced image using manually marked orientation field.	105
Figure 3.20	CMC curves for manually marked minutiae, enhanced image using reconstructed orientation field, highest score rank-level fusion and Borda count fusion.	106

Figure 3.21	CMC curves for score level fusion for latents of all qualities.	107
Figure 3.22	Matched minutiae shown for manually marked, rolled and enhanced latent images. (a) Boosted max retrieved the true mate for this latent at rank one even though one of the retrieved ranks (enhanced image) was 2,403 and (b) the retrieved rank for the true mate for this latent is rank one after boosted max was applied even though neither the manually marked minutiae nor enhanced image retrieved the true mate at rank one.	109
Figure 3.23	Alignment Accuracy: percentage of correctly aligned latents vs. misalignment threshold.	111
Figure 3.24	Example in which DBHT (Descriptor-Based Hough Transform) alignment is better than GHT (Generalized Hough Transform) alignment. (a) Latent with manually marked minutiae, (b) corresponding rolled print with automatically extracted minutiae, (c) rolled print with latent minutiae aligned by GHT, and (d) aligned by DBHT.	112
Figure 3.25	Example in which DBHT (Descriptor-Based Hough Transform) alignment is better than GHT (Generalized Hough Transform) alignment. (a) Latent with manually marked minutiae, (b) corresponding rolled print with automatically extracted minutiae, (c) rolled print with latent minutiae aligned by GHT, and (d) aligned by DBHT.	113
Figure 3.26	Example of alignment error due to the small number of true mated minutia pairs in the overlapping area between a latent and its mated rolled print. Note that there is only one aligned minutiae pair here.	114
Figure 3.27	Performance of Morpho, MCC SDK, and Proposed Matcher when the union of manually marked minutiae (MMM) extracted from latents and automatically extracted minutiae by Morpho from rolled prints is input to the matchers.	115
Figure 3.28	Performance comparison using manually marked minutiae (MMM) and automatically extracted minutiae from latents.	118
Figure 3.29	Two latent prints from the WVU database correctly identified at rank-1 by the proposed matcher but ranked below 20 by Morpho.	120
Figure 3.30	Latent prints ((a) from NIST SD27 and (b) from WVULFD) whose mates were not retrieved in the top 20 candidates by the proposed matcher but correctly identified at rank-1 by Morpho matcher.	121
Figure 3.31	Score-level fusion of the proposed matcher and Morpho for NIST SD27 and WVU databases.	122

Figure 3.32	Latent print mate from NIST SD 27 identified at rank 1 after score-level fusion of Morpho and proposed matcher. The first column shows (a) a latent, (c) its true mate, (e) rank-1 non-mate by the proposed matcher, and (g) rank-1 non-mate by Morpho matcher. The second column shows (b) latent minutiae, (d), (f), and (h) latent minutiae (in blue) aligned by the proposed matcher to the rolled print minutiae shown in (c),(e) and (g). In (c), (e), and (g), the numbers in parentheses indicate the ranks at which true mated rolled print was retrieved by the proposed matcher and Morpho matcher, respectively. . . .	126
Figure 3.33	Latent print mate from WVU LFD identified at rank 1 after score-level fusion of Morpho and proposed matcher. The first column shows (a) a latent, (c) its true mate, (e) rank-1 non-mate by the proposed matcher, and (g) rank-1 non-mate by Morpho matcher. The second column shows (b) latent minutiae, (d), (f), and (h) latent minutiae (in blue) aligned by the proposed matcher to the rolled print minutiae shown in (c),(e) and (g). In (c), (e), and (g), the numbers in parentheses indicate the ranks at which true mated rolled print was retrieved by the proposed matcher and Morpho matcher, respectively. . . .	128
Figure 3.34	Score-level fusion of the two proposed approaches: by enhancement and using descriptor-based Hough Transform (DBHT) matching approach. By combining the two approaches, the rank-100 accuracy is 81%.	130
Figure 4.1	Examples of fingerprints. (a) Rolled print (Michigan State Police database), (b) plain or slap print (FVC 2002) and (c) latent print (NIST SD27).	133
Figure 4.2	(a) Reference print from Michigan State Police database, (b) its skeleton extracted by a commercial matcher, (c) orientation field extracted from skeleton image, and (d) minutiae extracted by a commercial matcher.	141
Figure 4.3	(a) Skeleton of a reference print from Michigan State Police database, (b) the grayscale image obtained after applying image filtering and (c) the extracted orientation field.	142
Figure 4.4	(a) Latent print from NIST SD27 latent database, and some manually marked features: (b) orientation field, (c) singular points and (d) minutiae.	143
Figure 4.5	Overview of the proposed indexing scheme.	145
Figure 4.6	Illustration of triplets features.	147
Figure 4.7	Illustration of the key generation in the orientation field descriptor indexing scheme.	150
Figure 4.8	Rolled prints in (a) NIST SD27, (b) NIST SD14, and (c) Michigan State Police databases.	153

Figure 4.9	Comparison of the indexing performance of basic triplets, constrained triplets and MCC indexing.	154
Figure 4.10	Comparison of indexing performance based on the proposed approach against Yuan <i>et al.</i> [10] and Feng and Jain [11] on the NIST SD27 latent database. Note that Feng and Jain [11] used only a small background database of 10,258 versus the proposed approach and [10] that use 267,258 reference prints in the background database. Further, Feng and Jain [11] did not provide the complete hit rate vs. Penetration rate curve, but only a single point at a penetration rate of 39% (●).	155
Figure 4.11	Increase in the cumulative hit rate at a fixed penetration rate (20%) as we incrementally add features for indexing (from left to right), where CTrip denotes constrained triplets, OFDesc denotes orientation field descriptor, SP denotes singular points, RP denotes ridge period. The plot shows that we can achieve about 91% hit rate at 20% penetration rate when we use all five features (CTrip, OFDesc, MCC, SP and RP) for indexing.	156
Figure 4.12	Latent matching performance of two different matchers (DBHT in-house latent matcher and Morpho matcher) before and after applying the proposed indexing to filter out 80% of the background database. Note that after the filtering and the fusion of indexing and match scores, the overall matching performance is improved while the matching time is reduced by a factor of five (20% penetration rate).	158
Figure 4.13	Examples of latents for which the true mate was correctly retrieved when 20% of the background database was considered are shown. Note that by applying the indexing, the retrieval rank went from (a) 3 to 1 and (b) 142 to 49 due to the exclusion of non-mated reference prints with high match scores to the latent. In (b), the true mate only appeared in the top-100 candidate list after the filtering.	159
Figure 4.14	Latent matching performance of a commercial latent fingerprint matcher before and after applying the proposed indexing and retaining one third of the background database. The filtering improves the overall matching performance of the commercial latent fingerprint matching while reducing the matching time.	160
Figure 4.15	Examples of latents for which the true mate was correctly retrieved when one third of the background database was considered are shown. Note that by applying the indexing, the retrieval rank went from last to 1 due to the fusion of indexing and match scores.	161

Figure 4.16 Examples of latents for which the true mate was not retrieved when one third of the background database was retained. The true mates retrieval ranks before the filtering were 1, so the indexing degrades the matching performance for these two latents. 162

CHAPTER 1

INTRODUCTION

The first systematic capture of biometric data for identification purpose was done by Sir William Herschel in 1858, when he captured a handprint of each worker so that, later at payday, he could verify each worker's identity to avoid a worker receiving someone else's pay [12]. In 1870, Alphonse Bertillon [13] developed a method of person identification based on body measurements known as "Bertillonage" (see Fig. 1.1). This method was developed to identify criminals who have been previously arrested, repeat offenders, and who were using a different name to avoid harsher penalties. The Bertillon system recorded the precise measurements of various attributes of human body such as height, length of the arm and geometry of the head, as well as a listing of marks present on the body surface such as scars, moles and tattoos. This information was then stored and separated into categories for fast matching and retrieval. This system was believed to be very reliable and it was indeed used in law enforcement from the late 19th century to the early 20th century. According to [14], the French police used the Bertillon system to capture 241 repeat offenders in the year of 1884; after that, this system was adopted by law enforcement agencies both in Europe and in the Americas. However, it was later discovered that those measurements were not unique to each person; furthermore, two measurements of the same subject taken by different officers were not always consistent. Thus, Bertillon system was considered unreliable and abandoned by the police authorities to be replaced by fingerprint identification.

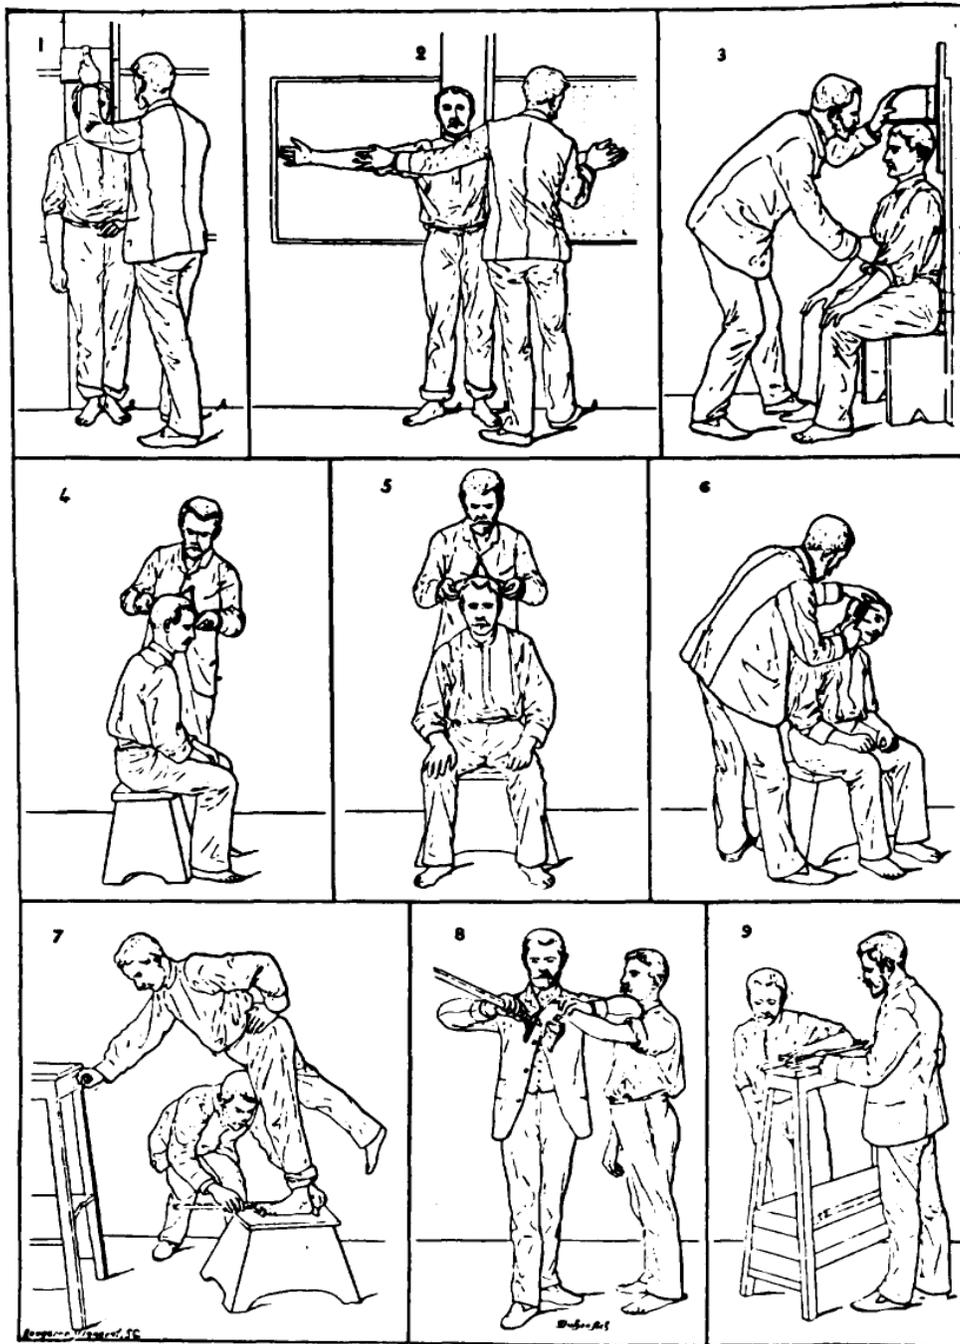


Figure 1.1: Illustration of Bertillon system for person identification based on numerous body measurements, which was used in the United States and in Europe from the end of the 19th century to the beginning of the 20th century. From left to right and then top to bottom the figures show measurement of height, reach, trunk, length of head, width of head, right ear, left foot, left middle finger, and left forearm.

A few years after Bertillon system was developed, several publications on the use of fingerprints as a mean of person identification started to appear [15, 16, 17]. According to Cummins and Midlo [18], the first scientific publication related to fingerprint identification appeared in *Nature*, in 1880 [15]. In [15], Faulds suggested that fingerprints left at crime scenes could be used to identify criminals or to exclude suspects. Soon after this article was published, a letter written by Herschel was published in *Nature* [16] stating that he had been using fingerprint as a method of identification in India for about 20 years, with different applications such as to avoid personification. In 1888, Sir Francis Galton introduced the idea of using minutiae features to compare two fingerprints [17]. Sir Francis Galton also introduced a categorization for fingerprints into three major classes (arch, loop and whorl), and the classes were further divided into subclasses.

After reading Sir Francis Galton's work on using fingerprints as a mean of identification, Juan Vucetich started collecting fingerprints from arrested men, along with measurements from the Bertillon system. Juan Vucetich was a police official in Argentina [19], and, in 1892, he used fingerprint evidence to identify a suspect of first-degree murder, who was then convicted. This was the first use of fingerprint as forensic evidence.

In 1893, the acceptance of the hypothesis by the Home Ministry Office, UK, that any two individuals have different fingerprints made many law enforcement agencies aware of the potential of using fingerprints as a mean of identification [1]. Some of the law enforcement agencies started collecting fingerprints from offenders so that they could identify them later in case they used an alias (changed their names) to evade harsher penalties. The fingerprints collected from crime scenes were compared to fingerprints collected from previous offenders in order to identify repeat offenders.

In 1900, Edward Henry has refined the fingerprint classification scheme introduced by Galton by adding more categories [1]. This resulting classification scheme (referred to as Galton-Henry classification) was adopted by several law enforcement agencies in various countries and it has been widely used.

The initial interest in recognition technologies were mainly by law enforcement agencies; more

recently, the concerns about security and identity fraud have increased, therefore creating a need for biometric recognition technologies in non-forensic applications [1] (e.g. border control, national ID, etc.). Over the years, several physical and behavioral characteristics have been explored, leading to the development of new recognition technologies based on face, fingerprint, hand geometry and voice that have been successfully deployed.

The use of physical or behavioral characteristics (e.g., fingerprint, iris, signature, etc.) to automatically recognize a person is referred to as biometrics or biometric system. A biometric characteristic (or trait) is a measurable physical or behavioral characteristic of an individual that is distinguishable and repeatable. A biometric trait should have four main characteristics: (i) be universal, meaning most people should have the trait; (ii) be distinctive, meaning it should be different from person to person; (iii) be permanent, meaning it should be invariant (relative to matching criteria) over time; and (iv) be collectable (measurable), meaning it can be measured quantitatively [1]. These characteristics make biometrics a reliable solution to person recognition. Additionally, biometric traits cannot be transferred, forgotten, guessed, shared or lost like other means of person recognition such as passwords, ID-cards, etc.

Biometric recognition can be very challenging because of intra-class variability and inter-class similarity of biometric traits [5]. Intra-class variability refers to how a biometric trait can appear very different in multiple acquisitions of the same individual, and inter-class similarity refers to how a biometric trait can be very similar for different persons. Intra-class variability is exemplified in Fig. 1.2a, in which two different impressions of the same finger are shown, and inter-class similarity is exemplified in Fig. 1.2b, in which two face images of a pair of identical twins are shown. Although the two images shown in Fig. 1.2a are impressions of the same finger of the same person (same class), the second impression in the figure contains a large amount of distortion compared to the first one. In Fig. 1.2b, the two face images are from different persons (different classes) but since they are identical twins, their facial appearance is extremely similar.



(a)



(b)

Figure 1.2: Examples of intra-class variability and inter-class similarity. (a) Two different impressions of the same finger of the same subject and (b) face images of a pair of identical twins. For interpretation of the references to color in this and all other figures, the reader is referred to the electronic version of this dissertation.

1.1 Biometric Systems

A biometric recognition system involves two basic phases: enrollment and recognition [5]. In the enrollment phase, a sensor collects the biometric data from which a set of features are extracted (template) and stored in a database along with the individual's identity (name, ID number, etc.). In the recognition phase, the identity of an individual is either confirmed (verification) or determined (identification) by collecting the biometric data again, extracting the same features and comparing

them to the features stored in the database. From this comparison, a match (similarity) score is generated and used to make a decision to whether the two sets of features came from the same subject or not. An overview of a biometric system is shown in Fig. 1.3.

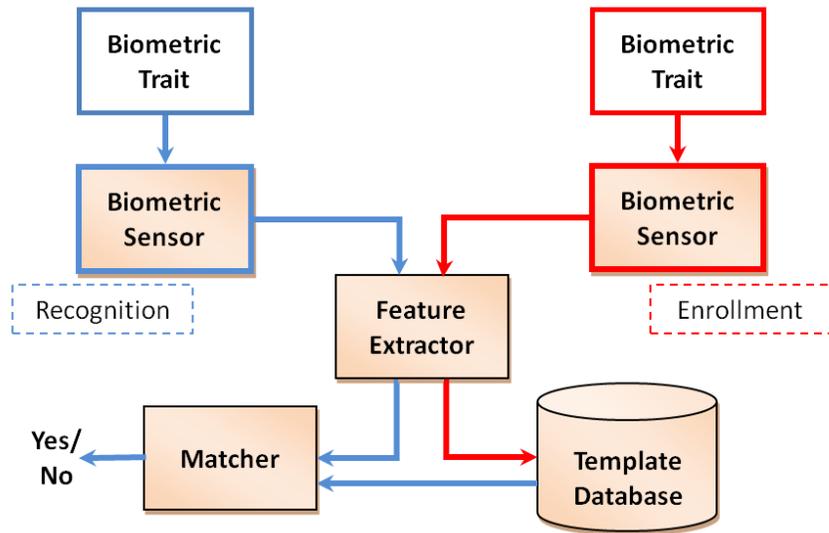


Figure 1.3: An overview of a biometric system.

A biometric system can be designed to work in two different modes: (i) verification mode or (ii) identification mode. In the verification mode, the user will claim his/her identity by using a personal identification number, a user name, a token, etc; then, the system will verify whether the subject is the person he/she is claiming to be (genuine) or not (impostor) by using biometric traits. In this case, only the stored data related to the claimed identity is compared to the data of the person making the claim to an identity. A threshold on the match score is usually applied to decide whether the two biometric samples are true mates. If the match score is higher than the system threshold, then the samples are considered to be from the same source. Otherwise, they are considered to be an impostor pair. An example of the verification scenario occurs when you try to use the ATM at a bank and you have to provide biometrics data along with your ATM card to verify your identity. In this case, the owner of the ATM card is known, so the biometric system needs to ensure that the true owner is the one who is using the card to perform the transaction.

In the identification mode, the user does not claim any identity. The user only provides his/her

biometric data, and the data is compared to the stored template of every subject in the system database. In the ATM example cited above, now the ATM account holder will not insert his ATM card in the machine, but simply provide his biometric data for identification. If the system can find a subject in the database for whom the query biometric data is close enough, then the query is considered to be from the same subject. Otherwise, the system will output that the user is not enrolled in the database. For example, when a fingerprint impression is found at a crime scene, forensic examiner usually does not know to whom the fingerprint belongs. So, the crime scene fingerprint is compared to the reference (rolled) prints stored in the database. If a match is found, the identity of the suspect is determined.

A match score is called a genuine score if it is computed between mated samples and an impostor score if it is computed between non-mated samples. In order to evaluate and compare the performance of biometric systems, the most commonly used quantitative measures are Receiver Operating Characteristic (ROC) curve in the verification mode and Cumulative Match Characteristic (CMC) curve in the identification mode [5].

The ROC is a plot of the Genuine Accept Rate (GAR) versus the False Acceptance Rate (FAR) as the match score threshold is varied. The GAR refers to the portion of true mated subjects who are correctly accepted by the system, while FAR refers to the portion of impostor mates who are falsely accepted by the system at a specified threshold. While the desired values of GAR and FAR are application dependent, a biometric system, operating in the verification mode with high GAR at low FAR is, in general, preferred.

In the identification mode, given a query, it is desirable that the match score between this query and its true mate be the highest score compared to the match scores between the query and any impostor. For each query, the database can be ordered based on the match score. The identity in the database corresponding to the highest score is considered the rank-1 match, while the identity corresponding to the second highest score is considered the rank-2 match and so on. Ideally, we want the true mate to be in the rank-1 position. The CMC curve is a plot of the identification rate versus the rank. Given a number of queries and a rank t , we can measure the percentage of the true

mates that are retrieved at rank- t or higher; this percentage is the identification rate at rank- t . The system with higher identification rates for given queries is the most desirable system.

1.2 Common Biometric Traits

Several biometric traits have been used for recognition, such as face, fingerprint, iris, palmprint, hand geometry, voice, ear shape, signature, key stroke and gait [5]. We focus on fingerprint, face and iris, which are the most prominent traits and are used in this dissertation.

Each biometric trait has its strengths and weaknesses, and the choice of a biometric trait is usually dictated by the application requirements. Table 1.1 summarizes advantages and disadvantages of the three biometric traits (fingerprint, face and iris) [1]. For each trait, the authors in [1] provided a perceived measure of the four characteristics: universality, distinctiveness, permanence and collectability. For example, the high collectability of face means it is very easy to collect this biometric trait.

It is generally known that fingerprints are very distinctive, and the overall ridge structure does not change significantly over time (from infants to old age); even after superficial cuts on the finger, the fingerprint pattern reappears after the healing process. This is supported by the high matching performance of state of the art systems in rolled/slap fingerprint matching [20]. Fingerprints are relatively easy to acquire, but require some degree of cooperation from the subjects (collectability is “medium” in Table 1.1). Also, some people might not have strong fingerprint pattern suitable for automatic recognition due to genetic factors, occupation, aging, etc [5].

Face is one of the most convenient biometrics in terms of acquisition and it does not even re-

Table 1.1: Comparison of the characteristics of the three most commonly used biometric traits from [1]. (H = high, M = medium, L = low)

Biometric Trait	Universality	Distinctiveness	Permanence	Collectability
Fingerprint	M	H	H	M
Face	H	L	M	H
Iris	H	H	H	M

Table 1.2: False reject and false acceptance rates associated with fingerprint, face and iris verification systems.

Biometric trait	Test	Test conditions	False reject rate (%)	False accept rate (%)
Fingerprint	FVC 2006	Heterogeneous population including manual workers and elderly people	4.2	0.1
	FpVTE 2003	Operational quality	0.6	0.01
Face	FRVT 2012	Mugshots	4.0	0.1
	FRVT 2012	Visa Photos	1.0	0.1
Iris	ICE 2006	Controlled illumination	1.1-1.4	0.1

quire cooperation from the subject. Face images can be easily acquired, even from a distance and without the subject's knowledge. This advantage is specially useful in surveillance applications. However, many variations in face images such as pose, lighting, expression, and changes in appearance such as make-up and accessories, make face recognition a very challenging problem [5]. This explains the low distinctiveness of state-of-the-art face recognition systems in unconstrained scenarios. Furthermore, face characteristics might not be stable over time (e.g., due to weight gain or aging).

Iris image capture requires a more sophisticated and expensive sensor because the useful texture patterns are better captured in near-infrared images; it also requires the subject's cooperation (subject must stand at a specified distance from the iris camera) and the iris quality can be influenced by a number of factors such as partially closed eyelids, eyelashes, contact lenses, etc. However, studies on large scale databases suggest that iris recognition systems can achieve extremely low error rates [21].

Table 1.2 shows false reject and false acceptance rates associated with fingerprint, face and iris verification systems [5].

In the next section, we review the representation (feature extraction) and matching techniques of fingerprint, face and iris biometric modalities.

1.2.1 Fingerprint Representation and Matching

A fingerprint is the impression of the friction ridge skin on a finger tip. Friction ridge skin presents raised ridges because their function is related to grasping and gripping; this explains their presence in the palms of our hands and sole of our feet. The characteristics of a fingerprint are determined during fetal development and its formation starts at approximately 6 or 7 weeks of gestational age [22]. The fingerprint pattern is mostly determined by the gestational environment, since minor changes in the flow of amniotic fluids are responsible for the differences in the skin structures around palm or finger tips. Thus, the fingerprints of every person are different among themselves, and different from fingerprints of other persons.

There are essentially three types of fingerprint images that are acquired for matching: (i) rolled, which is obtained by rolling the finger from “nail-to-nail” either on a paper (in this case ink is first applied to the finger surface) or the platen of a scanner; (ii) plain, which is obtained by placing the finger flat on a paper or the platen of a scanner without rolling; and (iii) latent, which is lifted from surfaces of objects that are inadvertently touched or handled by a person typically at crime scenes (see Fig. 1.4). Rolled and slap fingerprints are generally captured under controlled conditions and by cooperative subjects. On the other hand, latent prints are left by criminals at crime scenes. For this reason, rolled and plain prints (collectively referred to as reference prints) are of much better quality than latents. The two first types (rolled and plain) are usually collected in the form of a tenprint card, which is a card that contains rolled and plain impressions of the ten fingers of a subject, along with their identity information (see Fig. 3.2).

1.2.1.1 Representation

The structure present in a fingerprint is composed of ridges and valleys. In a rolled fingerprint image obtained by using ink (e.g. Fig. 1.4 (a)), the ridges are the dark areas corresponding to the raised ridges in our fingers and valleys are the bright areas that correspond to the space between the raised ridges.

Fingerprint characteristics or features can be categorized into three different levels (from coarse

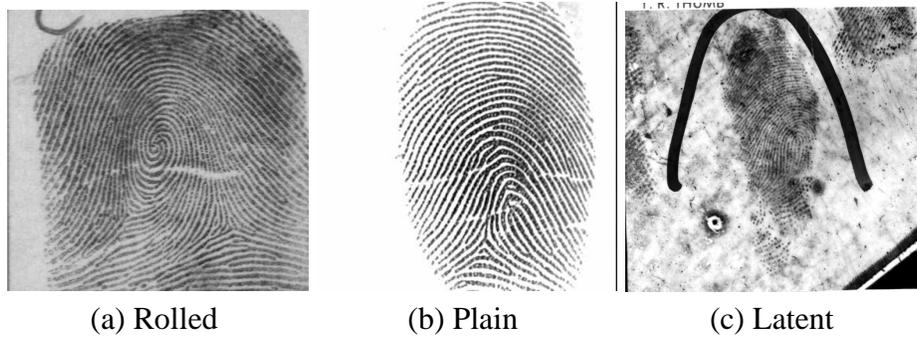


Figure 1.4: Three types of fingerprint images.

to fine): Level 1 (ridge flow), Level 2 (minutiae) and Level 3 (pores, incipient ridges, dots, etc) [1].

Fig. 1.5 shows examples of features in the three levels.

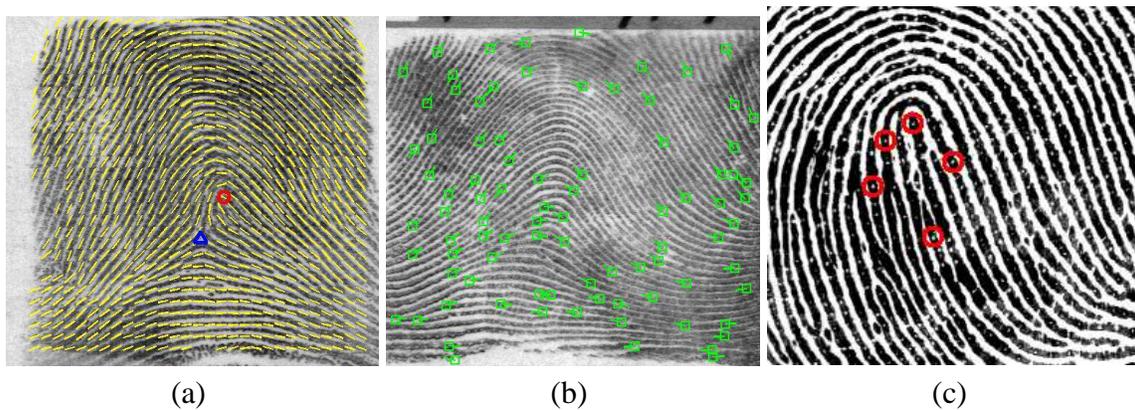


Figure 1.5: Examples of fingerprint features belonging to the three levels: (a) orientation field and singular points and (b) minutiae shown in fingerprints from NIST Special Database 4 [2], and (c) sweat pores shown on part of a rolled print image from the WVU Latent Database [3].

- Level 1 features

This is the coarsest level representation and consists mainly of the ridge orientation map and ridge frequency map of the fingerprint. The ridge orientation (ridge frequency) map consists of the local orientation (frequency) of the ridges in the fingerprint. The local orientation of a ridge at point (x,y) in the fingerprint image is defined as the orientation of the line tangent to the ridge passing through (x,y) , and it is in the range $[0, \pi)$. The local ridge frequency

at point (x,y) is defined as the average number of ridges passing through a line segment centered at (x,y) and of unit length that is normal to the local ridge orientation.

Ridge orientation is mostly smooth in fingerprints. However, there can be a few areas of the fingerprint in which the ridge orientations change abruptly. These locations are called singular points. There are two types of singular points: core, in which a set of ridges enters and exits from the same direction, and delta, in which three sets of ridges appear to meet. The number of cores and deltas in a fingerprint is determined by the fingerprint type. The maximum number of cores and deltas in a fingerprint is 4, and a fingerprint always contains an even number of singular points because cores and deltas appear in pairs. Therefore, a fingerprint might have no singular points (arch or tented arch type fingerprint), one core and one delta (loop type fingerprint), or two cores and two deltas (double loop or whorl type). Examples of ridge orientation map and singular points are shown in Fig. 1.6.

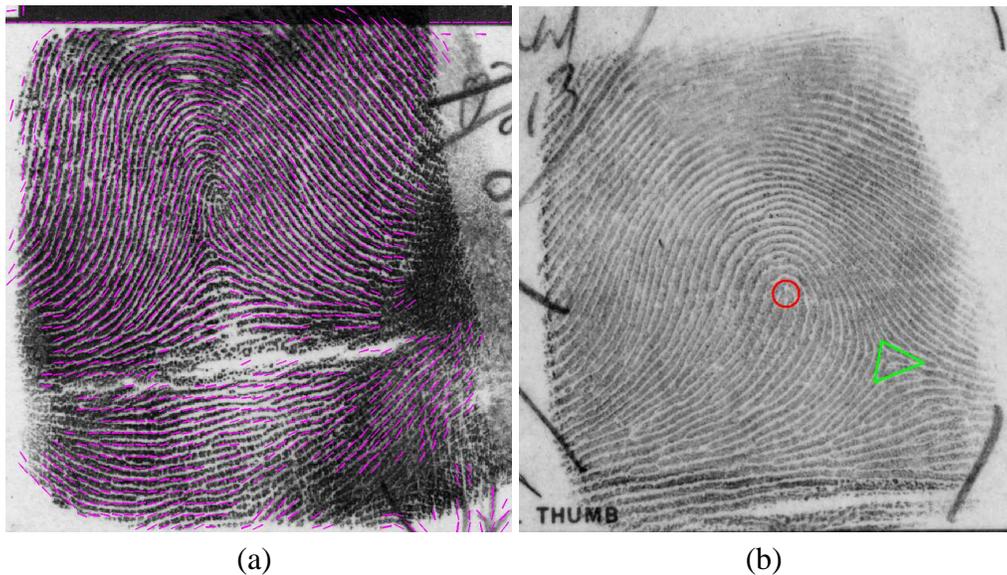


Figure 1.6: Examples of Level 1 features in fingerprints. (a) Orientation field and (b) singular points (core in red and delta in green) shown in rolled fingerprint images from NIST Special Database 27 [4].

Ridge orientation is also referred to as orientation field or ridge flow. We will mainly use the term orientation field throughout this dissertation to refer to this fingerprint feature. The

orientation field is commonly extracted by using a gradient-based method in local neighborhoods [8]. The fingerprint image is usually divided into small non-overlapping blocks (usually 8×8) and one dominant orientation is computed for each block as the orthogonal orientation of the gradient angle, the latter being the direction of the maximum intensity change. Ridges and valleys are more than one pixel wide, so to consider only one angle at a given pixel position makes it sensitive to noise. Therefore, an average of the gradient in a small neighborhood (window) is computed as the gradient angle estimate at a given pixel. This averaging is performed with following restrictions: (i) the angles are doubled so that the circularity of the orientations is maintained — for example the average of the angles 5° and 175° is 0° instead of 90° and (ii) the sin and cos components are used instead of angle values.

Based on the ridge structure, fingerprints can be classified into five classes: left loop, right loop, whorl, arch and tented arch [23] (see Fig. 1.7). Loops (left and right) are characterized by the ridge flow that enters from one side of the fingerprint, curves and returns on the same side from which it came; whorls are characterized by the ridge flow forming a complete circuit; arches are characterized by the ridge flow that enters from one side and exits the opposite side; and tented arches are similar to arches, but at least one ridge presents a high curvature.

- Level 2 features

The ridges in a fingerprint exhibit discontinuities in various ways. The locations where the discontinuities occur are called minutiae (minute or small detail). The most commonly used discontinuities are the locations where a ridge ends and the locations where a ridge bifurcates. A rolled print typically has about 100 minutiae points; the exact number depends upon the finger surface condition and image capture process. A minutia can be represented by its location (x and y coordinates), direction (the orientation of the ridge associated with the minutia), and type (ending or bifurcation). Minutiae characteristics of permanence, unique-

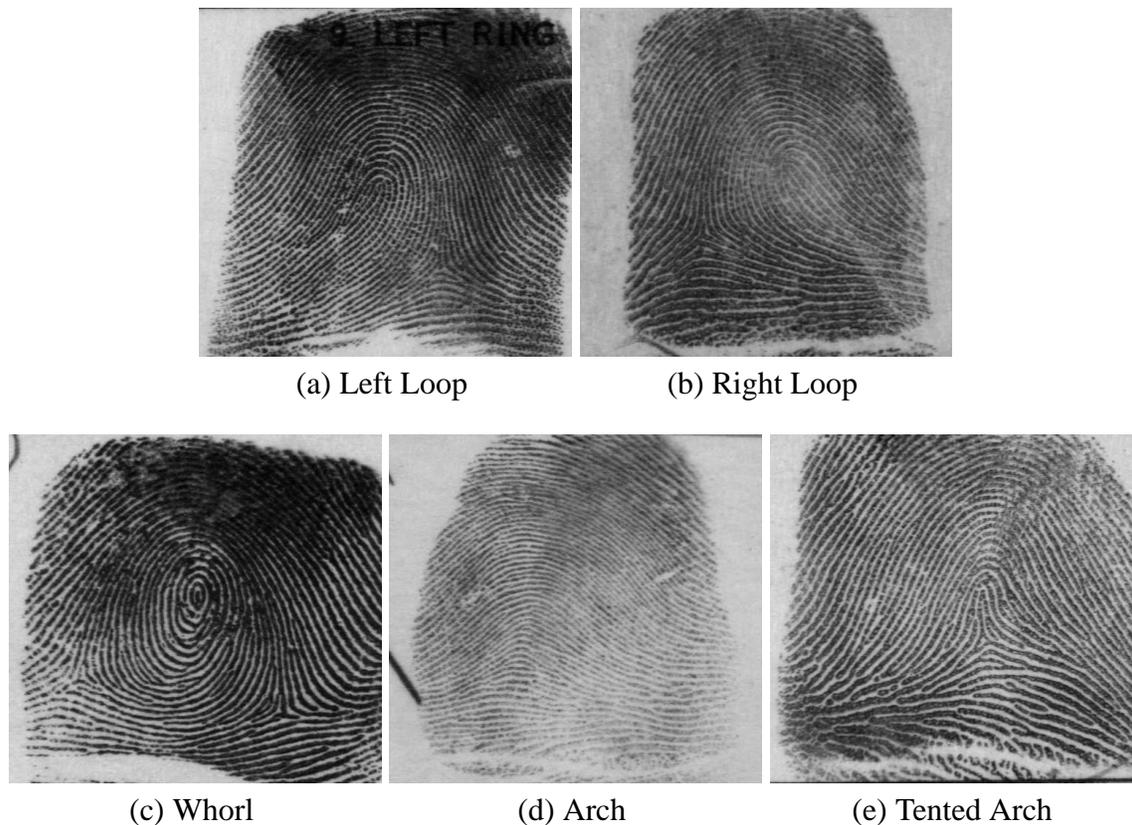


Figure 1.7: Example fingerprints belonging to each of the five major fingerprint classes (fingerprints from NIST Special Database 4 [2].)

ness and ease of representation make minutiae the most commonly used feature in fingerprint matching. Minutia type is often not reliable because depending on the pressure that is applied to make the impression, a bifurcation can look like an ending and vice versa. The minutia direction definition is illustrated in Fig. 1.8; this definition ensures that the error in minutia type does not influence the direction estimation.

The most common method of extracting minutiae is to use binarization and thinning techniques to obtain the fingerprint ridge skeleton (one pixel wide dark ridges on a white background). The binarization consists of thresholding the fingerprint image so that the ridges are separated from the rest of the image (black ridges on a white background). Then, morphological operations using the ridge orientation are applied to reduce the width of the ridges to one pixel. From the skeleton, minutiae can be extracted by finding the endings and bi-

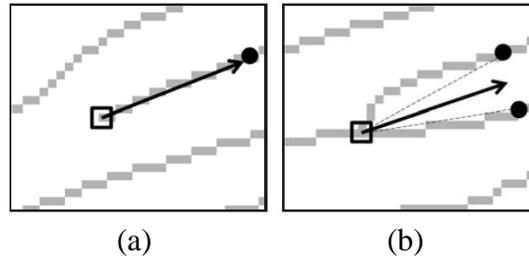


Figure 1.8: An illustration of how minutia direction is defined for (a) minutia ending and (b) minutia bifurcation [5].

furcations of the one pixel wide ridges. Usually some type of image enhancement (such as applying Gabor filters) is necessary before the binarization and thinning procedures. The goal of this enhancement step is to correct degradations occurring in ridge patterns that are very noisy and/or corrupted. As an example, parallel ridges might not be well separated because of the noise.

- Level 3 features

The Level 3 features can be viewed as micro level characteristics of a fingerprint. Level 3 features include pores (sweat pores), incipient ridges, dots, dimensional attributes of ridges such as width and shape, etc. [1, 5]. These features are not usually observed in low resolution images (less than 1000 ppi). Level 3 features are very important for latent fingerprint examiners, especially when the number of minutiae in latent is too small (e.g. less than 5), because they are highly distinctive and easily observed in high resolution images (over 1000 ppi). However, they are not widely used in automatic fingerprint matching systems because, even in high resolution images, their extraction is both computationally demanding and not as reliable as level 1 and level 2 features. In [24], the use of level 3 features is supported for latent print matching because forensic experts often rely on additional information, especially when the number of minutiae in latents is small, to compare latents to rolled prints. Fig. 1.5 (c) shows examples of sweat pores, which are regularly spaced along the ridges.

1.2.1.2 Matching

Most published fingerprint matching algorithms are based on minutiae. A few techniques are based on image correlation methods (e.g., [25, 26]). In the latter, the features are mostly the pixel intensities. In this case, the correlation between two fingerprint images is computed for alignment, which can be performed globally or locally. The images of different impressions of the same finger might appear very different due to pressure variations (ridge thickness, contrast, global structure, etc), which largely affects the correlation between two images. Also, the methods included in this category might be computationally very expensive [1].

There are other methods that are feature-based, where the features include singular points, level 3 features, texture information, etc. (e.g., [27, 28]). These methods sometimes are used in combination with minutiae to improve the matching performance and they are especially useful in cases where it is very difficult to extract minutiae or the number of minutiae is small, as often is the case with latents. However, these non-minutiae features are not as distinctive as minutiae, so they can only be used in conjunction with minutiae.

As mentioned before, the most common fingerprint matching algorithms are minutiae-based. There are three main steps in fingerprint matching using minutiae: alignment, pairing and score computation. Alignment refers to estimating the parameters between two minutiae sets that can be used to transform one of the sets to the same coordinate system as the second set; pairing refers to finding corresponding minutiae; and score computation refers to assigning a match score to a pair of fingerprints usually based on the number of corresponding minutiae.

Two common fingerprint alignment methods are based on (i) pairs of matched minutiae and (ii) Generalized Hough Transform. In the first method, a pair of matched minutiae is found using, for example, a minutia descriptor. A minutia descriptor is a structure that contains information about the local neighborhood of the minutia. This information can be based on neighboring minutiae, or the texture around minutia, etc. The local minutiae descriptors are matched and a small number of most similar minutiae candidate pairs are used for initial alignment. The rotation in the alignment is derived from the direction difference between the minutiae pair, and translation parameters are

estimated as the distance in x and y coordinates between the two minutiae.

Alignment based on the Generalized Hough Transform [29] consists of finding the peaks in the parameter space associated with the rigid transformation between the two sets. In this case, each pair of minutiae (one in each fingerprint) will vote for a specific set of translation and rotation parameters computed as mentioned above. Then, the peak in the parameter space (Hough Transform) is obtained — usually more than one peak is selected for robustness.

After the fingerprints are aligned, the pairing consists of finding the minutiae correspondences between the two fingerprints. The simplest way of finding the correspondences is to consider a pair of minutiae as matched minutiae if the distance between them and their directional difference are smaller than some pre-specified thresholds (e.g., 15 pixels in translation and 20° in rotation). Fig. 1.9 shows an example of minutiae correspondences in a fingerprint pair (latent and its true mate).

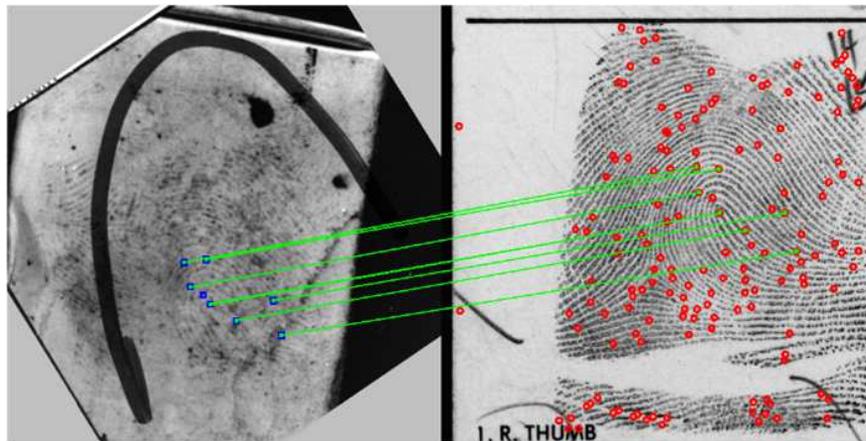


Figure 1.9: Example of matched minutiae from a fingerprint pair (latent and rolled print).

Score computation can be done in a number of ways. However, the most influential feature in the scoring process is the number of matched minutiae. Other major features used in match score computation include ridge flow, ridge frequency, etc. One simple way of generating a score between two fingerprints is to count the number of matched minutiae, and to normalize this number by the average number of minutiae in both fingerprints. In cases where the overlapping area between the two fingerprints is small, the score will be excessively penalized by this normalization scheme; so, the number of minutiae in the overlapping area should be used in the averaging instead

of the total number of minutiae in the fingerprints. Some other features might also be useful in the scoring, such as the quality of minutiae and the similarity of local minutiae descriptors.

1.2.1.3 Indexing

In the identification mode, a search fingerprint is compared to the reference prints in the background database to find the true identity of the subject, if present. When the background fingerprint database is extremely large (say, tens of millions), fingerprint matching becomes a challenge in terms of computational load. A common approach to alleviate this problem is to quickly exclude a large number of reference fingerprints with low similarity to the search fingerprint before performing the more detailed one to one matching. This process of filtering out a large percentage of the background database quickly (e.g., based on level 1 features) before fine level matching is referred to as database filtering, fingerprint indexing or fingerprint retrieval. Indexing usually refers to an approach that provides continuous classification rather than exclusive classes or categories, while the terms filtering and retrieval can be used in both cases.

Approaches that have been proposed for fingerprint indexing can be mainly divided into three categories: minutiae-based, orientation field-based, and based on other features (e.g. SIFT [30]).

In the minutiae-based indexing case, descriptors based on minutiae are utilized (e.g. triplets, minutia cylinder code). The most common approach is based on minutiae triplets, or the triangles formed by sets of three minutiae in the fingerprint. Rotation and translation invariant features such as the side lengths of the triangles, the difference between minutiae direction and one side of the triangle, handedness of the triangle, minutiae type, etc., are quantized into bins to generate a unique key or index for each triangle. Given a search or query fingerprint, triplets in the query are generated and quantized in the same way to obtain the keys. Then, triplets in the background database are retrieved as matched triplets if they have the same key as the search triplets (e.g. [31, 32]). The number of matched triplets between the search and each fingerprint in the database can be used to retrieve the fingerprints that are most similar to the search print.

Orientation field-based indexing usually involves aligning the fingerprints with respect to some

reference point (e.g. singular points), and then clustering the background database so that the orientation field of the search print only needs to be compared to a limited number of representative orientation fields (e.g. [33]). All fingerprints in the closest cluster center are retrieved. Another approach involves representing the orientation field by coefficients of an orientation field model, using the coefficients to search the background database, and retrieving the fingerprints with orientation field coefficients similar to the ones of the search print [34].

Indexing performance is usually measured by computing the hit rate at various penetration rates. A hit rate at a given penetration rate p refers to the portion of the search prints for which the true mates are retrieved within $p\%$ of the background database (penetration rate). The desirable outcome for fingerprint indexing is high hit rates at low penetration rates, which means that only a small portion of the database needs to be searched (or fine matched) in order to find the true mate for a high percentage of the given search prints; high computational efficiency is also desired so that the overall search time is reduced compared to fine matching the search print to every reference print in the background database.

In order to be useful, indexing must be much faster than the more detailed one-to-one matching. However, it still needs to have a good performance so that the matching performance is not degraded. There is a tradeoff between the penetration rate and the hit rate. If a very small penetration rate is chosen, some true mate fingerprints might be excluded from the finer matching, thus the identification performance will drop. On the other hand, if a large penetration rate is chosen, the true mate is more likely to be in the retained background database, but a large portion of the database will need to be searched, thus making the total gain in speed very small.

1.2.1.4 Latent Fingerprints

The sweat and the contact of a finger touching a surface leaves an impression of the ridges on these surfaces. This type of fingerprint impressions are called latents. Although latents can refer to any impressions lifted from surfaces touched by a person, we usually use this term for impressions of fingers lifted at crime scenes. The moisture in the fingers that is transferred to object surface is

not usually visible. Thus, some processing methods, typically chemical [35], must be applied to make them visible so that a copy of the impression can be lifted from the objects and scanned. Alternatively, a digital photograph can be taken directly to generate a latent fingerprint image.

Latent prints can be lifted from a variety of objects and materials such as metal, paper, plastic, human skin, glasses, etc. Some of the characteristics of latent fingerprints that makes latent matching a very difficult problem compared to reference to reference print matching include: (i) the uneven pressure of the finger and the non-flatness of the surface can generate a distorted impression of the finger; (ii) the area of the fingerprint ridge structure contained in a latent impression is usually very small; (iii) the way the object was touched by the finger might generate a blurred or smudged impression; and, (iv) in the case of a digital photograph of the latent, the image might contain a very noisy background. All these factors heavily influence the quality of the latent, and pose a major challenge to latent matching. Some examples of latent fingerprints are shown in Fig. 1.10.

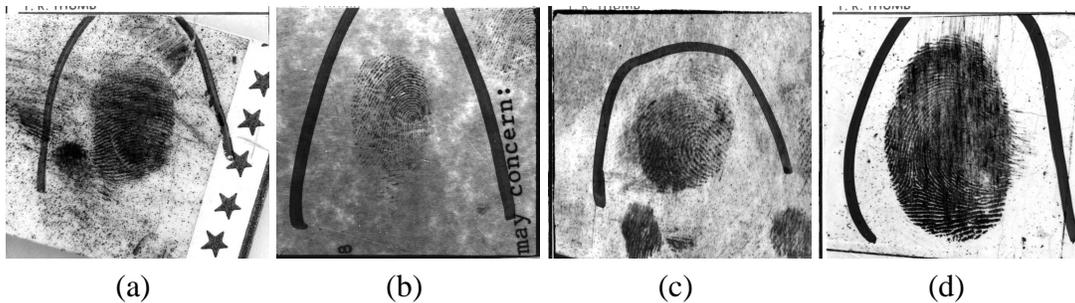


Figure 1.10: Examples of latent fingerprints in NIST SD 27 database [4].

Due to the aforementioned factors, latent fingerprint images are generally of very poor quality compared to rolled and plain fingerprint, which makes it very difficult to automatically extract reliable features from them. Thus, features in latents are usually manually marked by latent examiners. Features that are commonly marked by latent examiners are minutiae, singular points and region of interest (ROI). Furthermore, because latents capture only a partial impression of the finger, the number of minutiae in latents compared to plain or rolled prints is small. For example, in NIST SD 27 [4], the average number of minutiae is 21 in the latents and 106 in the mated rolled prints. Since

latent prints can be lifted from a variety of object surfaces in uncontrolled environments, latents are also more likely to contain large amount of distortion. The feature extraction in latents and the matching of latent prints to reference prints are challenging and they require more specialized algorithms than the ones available for matching reference fingerprints (rolled to rolled or slap to slap).

NIST evaluation of reference fingerprint matching technology reports excellent performance. In the Fingerprint Vendor Technology Evaluation (FpVTE) [20] multiple fingerprint recognition systems were evaluated on different types of data and scenarios. The experiments were divided into three main groups: large-scale test (LST), medium-scale test (MST) and small-scale test (SST). The best performing fingerprint recognition system, evaluated on a medium-scale test and using single fingerprints of operational quality, reached a true accept rate of 99.4% at 0.01% false acceptance rate. Latent fingerprint matching technology has also been evaluated. In Phase I of the Evaluation of Latent Fingerprint Technologies (ELFT), the average reported rank-1 identification accuracy (100 latents against 10,000 rolled prints) was 67%. In Phase II, the best reported rank-1 identification accuracy (835 latents against 100,000 rolled prints) was 97.2%. It is important to clarify that in Phase II, the latent images selected were of very good quality compared to Phase I, which explains the impressive performance of Phase II results over Phase I results.

A tentative benchmark was established for fingerprint indexing, but no systematic evaluation is yet available [36]. Several researchers have reported the performance of their indexing approaches on publicly available databases such as NIST SD4 [2] and FVC databases [37]. For example, at 20% penetration rate, the hit rate is around 98% for NIST SD4 and as high as 100% for FVC databases [38]. The indexing performance for latents is not as good as for the reference prints. The reported hit rates for latents in NIST SD27 (the only public domain latent database) are: 97.3% at 39% penetration rate [11], and 92.7% at 40% penetration rate [10]. The results reported in [10] could not be verified by us since by applying the algorithm in [10] (code was provided to us by the authors) to our background database of similar size, the hit rate was only around 82% at 40% penetration rate.

1.2.2 Face Representation and Matching

Face is composed of the skull characteristics and the musculature and associated soft tissue [22]. These structures influence the variation among human faces, along with gender and age. Some challenges to face recognition systems include matching in the presence of variations in pose, lighting, expression, occlusion, weight changes, hair style, etc.

1.2.2.1 Representation

The first step in many face-related applications is face detection, in which the face location in the sensed image is determined. This process is usually done by placing masks or filters of varying size in the image and solving the two-class classification problem: face vs non-face for each patch. After the face is detected and localized inside the image, facial features can be extracted. Face images are usually aligned based on the eye positions and normalized (w.r.t. size and illumination) prior to matching.

Similar to fingerprint features, facial features can be divided into three levels [39]. Level 1 features consist of gross facial characteristics such as general geometry of the face. These coarse features are usually the global face features. They can be easily obtained even from low resolution images, and they can be used, for example, to quickly distinguish between an elongated and a round face. Level 2 features consist of more localized face characteristics. Some examples of level 2 features include the structure of facial components (e.g. mouth), the spatial relationship between facial components, etc. As in fingerprints, level 2 features are the most important features for face recognition. Level 3 features consist of the finest details of the face, which include scars, moles, freckles, etc. [5]. These different levels of facial features are illustrated in Fig. 1.11.

1.2.2.2 Matching

There are three main approaches to match face images: (i) appearance-based, (ii) model-based and (iii) texture-based approaches [5].

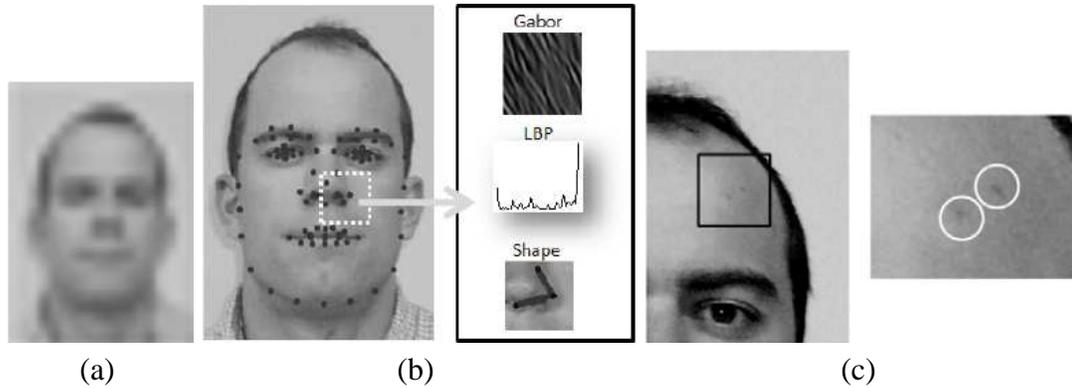


Figure 1.11: Examples of facial features at the three levels: (a) facial geometry (Level 1), (b) facial landmarks, Gabor filter, local binary pattern (LBP), and shape (Level 2) and (c) moles (Level 3) [5].

Appearance-based techniques refer to mapping a face image into a low dimensional sub-space. The set of representative vectors is learned based on a training set of face images. Examples of approaches in this category includes Principal Component Analysis (PCA) [40] and Linear Discriminant Analysis (LDA) [41]. In both of these approaches, a new face image can be represented in terms of the learned basis vectors (as a weighted sum of the basis vectors). The weights associated with a test image can be compared to the weights of reference images in a database by using the Euclidean distance, a measure of the dissimilarity between the two face images. The difference between PCA and LDA representation is that LDA incorporates class information in the training stage (supervised technique) while PCA does not (unsupervised technique). PCA projects the data so that the overall variance is maximized while the LDA projects the data so that the ratio of the inter-class variance to the intra-class variance is minimized.

Model-based approaches refer to the the use of face models. One example is graph matching, where the face is represented based on a model graph. A model graph, in which fiducial points (landmarks) of the face are associated with the nodes in the graph (see Fig. 1.12), is fitted to a face to generate a representation of that face. The model graph contains several local descriptors (bunch) at each fiducial point to account for the variations in the local neighborhoods in the face image. At each fiducial point of a query face image, a local descriptor is extracted and compared to

the descriptors in the stored model; the descriptor of the fiducial point in the query image is chosen as the most similar descriptor in the bunch.

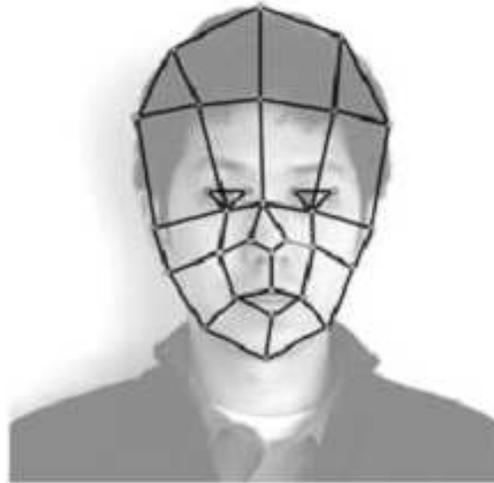


Figure 1.12: An illustration of a graph model fitted to a face [5].

Texture-based approaches use local features such as Local Binary Patterns (LBP) and Scale Invariant Feature Transform (SIFT). Both these local descriptors can be extracted at pre-specified points in the image (fixed if we assume the faces are roughly pre-aligned using the eyes) and feature vectors can be generated for both descriptors. The feature vectors can then be compared to generate a similarity or dissimilarity. SIFT is a histogram of gradient orientations in a neighborhood, whereas LBP is a representation of the relationship among the intensities of neighboring pixels. Fig. 1.13 shows schematic diagrams of SIFT and LBP features.

In the last decade, advances in face recognition technology have been reported by several third party evaluations, for example, Face Recognition Grand Challenge (FRGC) [42] and Multiple Biometrics Evaluation (MBE) [43]. In FRGC, the performances of several face recognition systems were measured in the verification mode. The best verification rate at a fixed false acceptance rate (FAR) of 0.001 was 99% for frontal face images in a controlled environment (studio lighting). However, the verification rate dropped to $\sim 80\%$ at the same FAR of 0.001 when frontal face images from uncontrolled environment were used as queries. MBE evaluated face recognition systems in an identification mode. The results from this evaluation showed a rank-1 accuracy of

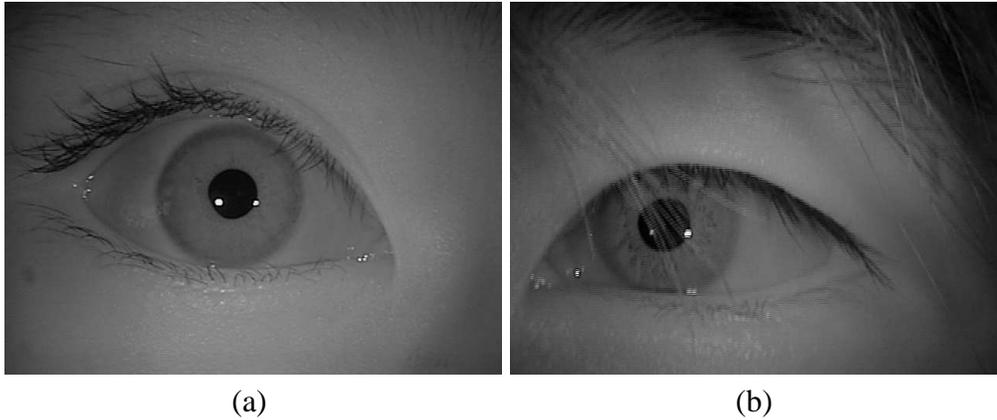


Figure 1.14: Examples of two irises of two different subjects from the CASIA twins database. The image on the right (b) is not of very high quality since the iris is partially occluded by hair.

In the next section, we present a brief overview of iris features and matching¹.

1.2.3.1 Representation

The first method to match two irises was designed and presented by John Daugman [44]. This method is well known and many of the commercial iris recognition systems still use its algorithms [5].

The first step in the extraction of iris features is the segmentation of the iris. Segmentation aims to detect the inner and outer boundaries of the iris. These boundaries are usually represented as circles or ellipses. The segmentation also generates a mask of the valid regions of the iris. The second step is the normalization of the segmented iris region. This step is accomplished by using a transformation that maps the cartesian coordinates of the iris to a pseudo-polar coordinate system. Each concentric region of the iris is mapped to a row in a rectangular image of the iris. In other words, the distance of a point in the iris image to the center of the iris is associated with a row in the rectangular representation, and the angles formed with the horizontal axis are associated with the columns. This mapping is also called “unwrapping of the iris”. This normalization process is necessary to account for the variations in the size of the iris texture. For example, if a pupil is

¹For a more detailed description, please refer to Chapter 4 of [5]

dilated, the captured texture area will be small, while if the pupil is contracted, the captured texture area will be large.

After the normalization step, two dimensional Gabor wavelets are convolved with the normalized image. Then, a demodulation is performed by taking the sign of the real and imaginary parts of the convolution output, thus yielding a two-bit representation for each point in the normalized image. This two-bit representation is basically a representation of which quadrant the phase falls into. The size of the normalized image is usually 1024 bits, so the resulting binary output is a vector of size 2048 and this binary vector is called the iris code.

1.2.3.2 Matching

Two iris codes can be compared by using the Hamming Distance [5]. Basically, each valid bit in the iris code is compared to its corresponding valid bit in the other iris, and the number of bits that differ normalized by the total number of valid bits in both codes is considered as the distance between the two iris codes. Fig. 1.15 shows an example of an iris recognition system.

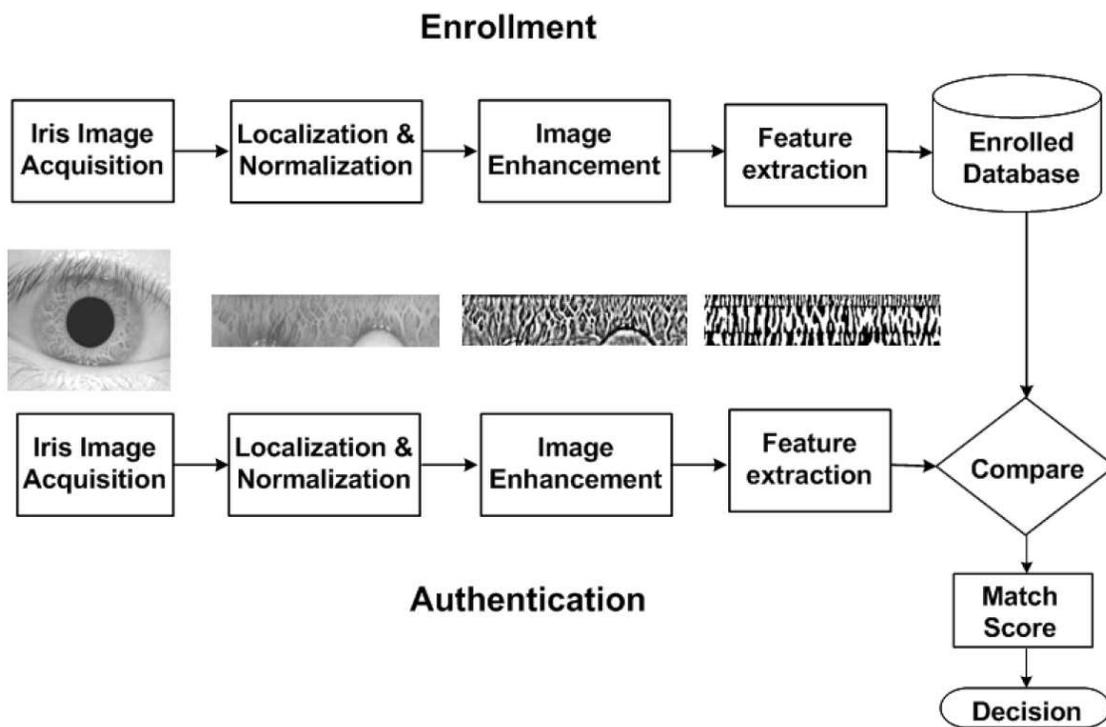


Figure 1.15: An overview of an iris recognition system [5].

NIST has evaluated 95 iris recognition algorithms in the identification scenario; iris images from over 2 million subjects were used [21]. When only one eye was used, the false negative error rates (the percentage of cases in which the true mate was not identified) are at 1.5% or higher; when both eyes were considered, this number dropped to 0.7%. These error rates are extremely low and are mostly due to poor quality images, abnormal irises, ground truth errors, etc.. When a threshold was set to produce no more than 25 false matches (false positive errors) in every 10^{13} comparisons, the false positive rate was still low (below 2.5%).

1.3 Applications of Biometric Systems

The increasing concern with security and identity fraud has supported the growth in the use of biometric recognition technology for applications other than forensic ones. Also, due to the great advancements in technology, nowadays it is more reliable and much less expensive to incorporate biometric recognition into several applications, such as border control, ATM access, national identification, voting access, health care access, criminal and victim identification, etc. Biometric recognition can be used in small and large scale applications. For example, one can use fingerprint to protect one's laptop or other device; on the other hand, in the Unique Identification Authority of India (UIDAI) program, biometric data (fingerprint, face and iris) were collected from hundreds of millions of subjects and this number keeps growing [45].

Other applications of biometric systems include:

- Health care access

According to [46], estimates indicate that, in 2006, approximately 20% of the cost incurred by health insurance providers involves some type of fraud. Also, the cost of health insurance could go down as much as 3.5% if the persons who have the health insurance and are authorized to receive the health insurance benefits were the only ones to actually benefit from it. To avoid fraud by an unauthorized person receiving health care under another person's health insurance, automated fingerprint verification is used at the health care providers.

- Elections

In Brazil, every adult must vote in all elections (federal, state and city levels). In 2000, the election was fully automated. The next step was to include fingerprint technology to efficiently and correctly verify whether the person who was trying to vote was the actual person registered to vote. It is expected that all the ballot boxes will have biometric recognition systems by 2020 [47].

- Finger vein recognition

A finger vein recognition system was commercialized by Fujitsu Ltd. to be placed in ATMs in Japan. The bank client still has an ATM card, but instead of inputting a personal identification number (PIN), the client will place his/her hand over the finger vein sensor [48]; the collected data is compared to the data stored in the card, and the access is granted based on the verification result. The entire verification process (insert the ATM card and verify finger veins) takes about the same time as if someone inserts the ATM card and input a PIN.

1.4 Contributions

The contributions of this dissertation are as follows:

1. Biometric-based identification of identical twins

- An analysis of the discriminative ability of the three most common biometric modalities (face, fingerprint and iris) of identical twins is conducted, including results of combining different biometric traits.
- An analysis of using multiple biometrics to quantitatively determine whether two subjects are identical twins is performed.

2. Latent fingerprint matching

- To improve latent fingerprint matching performance by the enhancement of the latent image.
- To improve latent fingerprint matching performance by using a descriptor-based Hough Transform.

3. Latent fingerprint indexing

- An indexing approach for latent fingerprints that reduces the total search computational time by a factor of 5 while still maintaining the matching accuracy.

1.4.1 Distinguishing Identical Twins

One of the challenges in biometric identification is related to distinguishing twins, especially identical twins. There are two types of twins: monozygotic (or identical) and dizygotic (or non-identical) [49]. Monozygotic twins are a result of a single fertilized egg that splits into two cells, each one giving origin to one individual. Dizygotic twins are a result of two different fertilized eggs. Monozygotic twins have the same deoxyribonucleic acid (DNA), thus their genotypic features (features influenced by the genetic material) are the same since they share the same genetic material, while some phenotypic features (features influenced by the environment) may be different. Therefore, identical twins are more likely to have biometric features with somewhat higher degree of similarity compared to non-identical twins or unrelated persons. Because of this, an analysis of the ability of a biometric system to distinguish identical twins is essential. In recent years, the Federal Bureau of Investigation (FBI) has shown an interest in further investigating the biometrics of twins by supporting the collection of twins' biometric data. A pair of face images of identical twins is shown in Fig. 1.16.

In this dissertation, we study the performance of biometric systems in the presence of identical twins for the three most commonly used biometric modalities, namely fingerprint, face and iris, as well as an analysis of the performance of multimodal systems, which use a combination of biometric traits to make a decision. Face recognition of identical twins is shown to be a chal-



(a)

Figure 1.16: A pair of identical twins from CASIA Twins database Version 2.

length to current face recognition systems, while fingerprint and iris matching results for identical twins show performance comparable to unrelated persons. The fusion of different instances from the same modality (e.g., multiple fingers and left and right iris) yields the best performance for distinguishing identical twins as well as for unrelated persons. We also provide an analysis of how multiple biometrics can be used to determine whether two subjects are identical twins. The similarities between the biometrics of identical twins can be very useful to advance the state-of-the-art biometric systems.

1.4.2 Latent Fingerprint Matching

Fingerprints have been successfully used for person recognition for over a century. Latent fingerprints constitute a valuable source of evidence in law enforcement agencies to help solve crimes. However, a majority of the processing (e.g. feature marking) involving latent prints is done manually by forensic (latent) experts. While progress has been made to automate this process, according to the latest evaluations by NIST [50], the state-of-the-art technology for fully automated latent matching (lights out processing) is not mature. The current practice involves latent examiners to manually mark the features in the latent and then input it to the system for automatic matching with reference prints. The matcher returns a list of candidates that are manually checked by a latent examiner who then makes the final identification. The practice of manually marking features in latents is labor intensive, but it is nevertheless currently more reliable than state-of-the-art

systems.

Latents are a particularly challenging to match because they are usually of poor quality, contain a small friction ridge area, present large non-linear distortion, can be blurred or smudged, and usually contain complex background noise. All these challenges make it very difficult to reliably extract features in latents; unreliable features lead to low matching performance. Given that the minutiae are the most commonly used features in fingerprint matching, latents contain significantly fewer number of minutiae than a reference fingerprint. This makes rolled-to-rolled fingerprint matching algorithms unsuitable for latent matching. In this dissertation, our efforts are focused on improving latent fingerprint matching performance, given reliable manually marked features in latents.

The first method consists of enhancing the latent image and fusing the matching score of the enhanced latents with the score based on manually marked minutiae. This first method outperforms a commercial fingerprint matcher on the publicly available latent database NIST SD 27.

The second method consists of developing a new latent fingerprint matcher based on manually marked minutiae that uses minutiae and orientation field information in the matching. The proposed matching algorithm outperforms three fingerprint matching algorithms on two different latent fingerprint databases (NIST SD 27 and WVU Latent databases).

1.4.3 Latent Fingerprint Indexing

The identification of a person requires that his/her fingerprint be compared to all the fingerprints in the database [1]. If the database is very large, matching a query fingerprint to the entire database might become computationally unfeasible. Thus, strategies to quickly filter out a large portion of the database without degradation in the matching performance are very useful.

Given manually marked features in latents, our indexing approach consists of using singular points, minutiae, orientation field and frequency to significantly reduce the background database size. By applying our indexing approach, we are able to filter out 80% of the reference database while maintaining the latent matching accuracy.

1.5 Dissertation Organization

This dissertation is organized as follows. In Chapter 2, we study the problem of distinguishing identical twins based on unimodal (face, fingerprint and iris trait alone) and multimodal (based on multiple traits) biometric systems. also analyze the possibility of using multibiometrics to provide a quantitative measure to determine whether two subjects enrolled in the database are identical twins. In Chapter 3, two different approaches to improve latent fingerprint matching performance using minimal amount of manually marked features in latents are presented and discussed. The proposed techniques perform better than some of the commercial fingerprint matchers on two different latent fingerprint databases. In Chapter 4, an indexing approach is presented to filter out a large portion of the background database, thus greatly speeding up the search while maintaining comparable matching accuracy. Finally, summary of our research and some ideas for future research are presented in Chapter 5.

CHAPTER 2

BIOMETRIC TRAITS OF IDENTICAL TWINS

2.1 Introduction

A twin is “one of two offspring produced in the same pregnancy” [49]. Twins can be categorized into two types with respect to the number of eggs fertilized: monozygotic (or identical) and dizygotic (or non-identical) [49]. Monozygotic twins are a result of a single fertilized egg that splits into two cells, each one giving origin to one individual; thus, monozygotic twins have the same deoxyribonucleic acid (DNA). Dizygotic twins are a result of two different fertilized eggs, and therefore they do not have the same DNA. Face images of a pair of identical twins and of a pair of non-identical twins are shown in Fig. 2.1. Three offspring produced in the same pregnancy are called triplets, four are called quadruplets, and so on. In this dissertation we are primarily interested in the ability of biometric traits to distinguish identical twins.

In the year 2009, approximately 153 of every 100,000 births in the United States were triplets or more [51]. This birth rate increased 400% during the 1980s and 1990s [52], and it has been fluctuating since 1999, with a slight downward trend [51]. However, the rate in the year 2009 represented a significant increase from 2008 (4%) compared to the fluctuations in this rate in previous



Figure 2.1: Face images of (a) a pair of identical (monozygotic) twins and (b) a pair of non-identical (dizygotic) twins from the University of Notre Dame ND-Twins database [7].

years [51]. The rapid increase of multiple births in the 1980s and 1990s has been associated with the increase in the use of fertility therapies and “the older age at childbearing”, the latter because “women in their 30s are more likely than younger women to conceive multiples spontaneously [52].”

The twin birth rate (33.2 per 1000 births in the United States in 2009) increased at an average rate of 3% per year between 1990 and 2004 [51]. Then, the average rate of increase slowed to less than 1% per year over the period 2005-2009 and increased to 2% from 2008 to 2009 [51]. Although the average rate of increase has decreased, the overall twin birth rate in the United States has increased by 76% since 1980 [51]. The number of dizygotic twin births varies depending on the ethnic group [53], while the number of monozygotic twin births is believed to be constant worldwide. According to Nora *et al.* [54], the frequency of identical twins ranges from 0.35% to 0.4% among all births. This, in turn, has created a requirement for biometric identification systems to accurately determine the identity of a person who has an identical twin.

Accurately distinguishing identical twins has important legal ramifications. In 1985, a woman was accused of making fraudulent money transfers from several banks by simply making telephone calls. She was acquitted of those charges after claiming that her sister — who was her identical twin but disappeared before the trial — committed the fraud. In the end, both twins were convicted.

Based on an analysis of their voice samples, a speech scientist and voice examiner testified that indeed both women had placed the fraudulent telephone calls [55]. In this case, an ability to distinguish identical twins based on the voice biometric was essential to the convictions.

In another incident, a rape suspect won two mistrials because there was no way of determining whether the key DNA evidence came from him or his identical twin. However, in the third trial held in 2006, the accused was convicted after the prosecutors provided evidence that he had committed a series of sexual assaults over time and had attempted sexual assaults with characteristics similar to the rape for which he was being tried [56]. A failure to distinguish the twins was the reason for the charges being dropped in the first two mistrials.

The similarity in facial appearance of identical twins may also give them a greater incentive to commit fraud than an average person. Imagine a scenario where you have an identical twin sibling, who is covered by health insurance while you are not. If you fall sick and need medical treatment, you need health insurance. While the health care providers need to establish your identity using a photo identification, you could use your twin's health insurance without being caught. But, suppose the health insurance company requires that you need to be identified using fingerprints (this is the case for some of the health insurance companies in Brazil) [57]; if the fingerprint recognition system can distinguish identical twin fingerprints, then you will not be able to get health care using your twin's photo ID and insurance policy.

The above facts and scenario indicate that the ability of biometric systems to identify identical twins is necessary. Since monozygotic twins have the same DNA, they cannot be distinguished using DNA alone [58]. Thus, it is necessary to use other forms of identification for monozygotic twins.

Recognition using biometric traits is now a well accepted and proven method. A biometric system relies on the distinctiveness of the biometric characteristics to perform the recognition. While many biometric techniques are extremely accurate, some variations in sensing data, noise, etc. can cause the matching performance to drop significantly. One could argue that it is more difficult to discriminate identical twins than unrelated persons because of their genetic similarity. Although

identical twins cannot be distinguished from each other using DNA, some of the biometric modalities, such as fingerprints, iris, and palmprints, can still be used to distinguish them [59]. Some studies have shown that face and voice can also be used to distinguish identical twins [60, 61]. Due to the difficulty in obtaining a large biometric database of identical twins, most experiments reported in the literature are based on small databases.

This chapter explores the ability of unimodal and multimodal biometric systems to distinguish identical twins. In Section 2.2 of this chapter, we present studies on biometric twin data reported in the literature. In Section 2.3, we present the underlying characteristics of the three biometric traits that will be used in our experiments, namely fingerprint, face, and iris. In Section 2.4, we analyze the experimental results on matching individual biometric traits as well as various combinations of modalities. In Section 2.5, we end this chapter by presenting our conclusions.

2.2 Related Work

In order to design a robust and efficient biometric system, the system must be able to handle a variety of situations like noisy data, limitations of the sensors, environmental conditions, and the presence of identical twins. Due to the similarity of their biometric characteristics, identical twins are more likely to pose a challenge to a biometric system. Therefore, it is important to address this problem when designing a biometric system. Table 2.1 summarizes the studies on discrimination of identical twins that have been reported in the literature; these studies are also discussed below.

Daugman and Downing [62] assessed the distinctiveness of iris pattern as a biometric identifier. The authors compared unrelated irises (irises from different persons), and genetically identical irises (irises that came from the same DNA), for example pair of irises of the same person or the irises of identical twins. They observed that the iris patterns of genetically identical eyes were as uncorrelated as the patterns of unrelated eyes. For example, the similarity of the six pairwise comparisons they performed between identical twins showed the same mean as for eyes that were not genetically related.

¹It is not clear in this work how many pairs of identical twins were used.

Table 2.1: Summary of studies on the biometrics of identical twins. Sets can include identical twin pairs as well as non-identical twin pairs.

STUDY	YEAR	BIOMETRIC TRAIT	DATABASE SIZE
Daugman and Downing [62]	2001	Iris	1 set
Jain <i>et al.</i> [58]	2002	Fingerprints	94 sets
Kodate <i>et al.</i> [60]	2002	Face	10 sets
Han <i>et al.</i> [63]	2004	Fingerprints	66 sets
Patil and Basu [64]	2004	Voice	12 sets
Bronstein <i>et al.</i> [65]	2005	Face (3D)	1 set
Kong <i>et al.</i> [59]	2006	Palmprints	53 sets
Srihari <i>et al.</i> [66]	2008	Fingerprints	298 sets of twins
Ariyaeinia <i>et al.</i> [61]	2008	Speech	49 sets
Sun <i>et al.</i> [67]	2010	Face, Fingerprints, Iris	51 sets
Hollingsworth <i>et al.</i> [68]	2011	Iris	76 sets
Phillips <i>et al.</i> [7]	2011	Face	126 sets
Pruitt <i>et al.</i> [69]	2011	Face	126 sets
Biswas <i>et al.</i> [70]	2011	Face	186 subjects ¹
Klare <i>et al.</i> [71]	2011	Face	126 sets
Vijayan <i>et al.</i> [72]	2011	Face (3D)	107 sets
Srinivas <i>et al.</i> [73]	2012	Face	126 sets
Tao <i>et al.</i> [74]	2012	Fingerprints	83 sets

In a study of identical twins' palmprints, Kong *et al.* [59] used 1,028 palmprint images from 53 pairs of identical twins. They performed two different twin matches. In the first experiment, they matched the palmprints of identical twins, which they called real twin match. In the second experiment, they matched the left and right palmprints of the same person, which they called virtual twin match. Note that in both the experiments the palmprints shared the same genetic information. The matching algorithm was based on the angular differences of orientation fields (of the palmprint ridge pattern) in the two palmprints being compared. The authors observed that while these genetically equivalent palmprints have correlated features, they can still be distinguished based on features extracted from palmprint patterns.

The above observation concerning palmprints is also true for fingerprints, as a study based on 94 pairs of identical twin fingerprints showed [58]. Jain *et al.* showed that fingerprint verification systems can be used to distinguish identical twins, even though their fingerprints are generally more correlated than fingerprints coming from two genetically unrelated persons. As an example, at a False Rejection Rate (FRR) of 1.05%, the twin-twin matching had a False Acceptance Rate (FAR) of 8.51%, while the twin-nontwin matching had a FAR of 2.20%. For another threshold value (FRR of 3.49%), the twin-twin FAR dropped to 1.06% and the twin-nontwin FAR dropped to 0.29%. In another analysis of fingerprints from 66 pairs of twins [63], it was also concluded that fingerprints can be used to identify identical twins with an insignificant drop in the matching performance: the Equal Error Rate (EER) reportedly increased by only 1-2% compared to nontwin impostor matching. The authors also extended their studies to assess the similarities among families of nontwins (52 families of four persons — parents and two children). They observed that the largest similarity occurred between identical twins, followed by between two siblings, between parents and their children, and between unrelated persons.

Srihari *et al.* [66] analyzed the similarity between twins' fingerprints in a study using fingerprint images from 298 pairs of twins (74 fraternal twins and 224 identical twins). The authors analyzed this similarity based on Level 1 and Level 2 features (that is, pattern of the ridge flow, and minutiae, respectively). With the level 1 features, they observed that twin fingers with the same label occurred approximately 55% of the time and approximately 32% for non-twins, which means twins' fingers are much more likely to have the same pattern type than non-twins' fingers. With the level 2 features, they concluded that the similarity between twin fingers is higher than between two arbitrary fingers (with identical twin fingers similarity being not significantly different from fraternal one), but twins can still be distinguished using fingerprints.

Although it is believed that it is difficult for face and voice, along with hand geometry, to distinguish between identical twins [58], some researchers have obtained encouraging results using face and voice to distinguish monozygotic twins [60, 65, 61]. A verification experiment with 10 sets of identical twins was performed using a 2D face recognition system; the experiment consisted

of enrolling one of the twins, and asking the other to try to log into the system. On this small database, the rejection threshold was always satisfied, leading to a rejection of all the impostor twins [60].

Another face recognition experiment for twins was based on 3D face images [65]. The recognition task was to distinguish between two identical twins; the authors tested three different algorithms: a 2D algorithm based on eigenfaces, a 3D algorithm based on rigid surface, and another 3D algorithm based on canonical forms. The test consisted of enrolling one of the twins and matching the enrolled subject to the other twin. All the other subjects served as impostors. For the first algorithm (2D eigenfaces), the rank-1 accuracy was 70.59% when enrolling the first twin and 75% when enrolling the other twin. In the second experiment, the rank-1 accuracy was 82.36% and 100.0%, respectively, while the third algorithm achieved perfect matching performance. According to the Web site Digital World Tokyo [75], the Japanese company Sagawa Advance has invented an infra-red based face recognition technology that is able to distinguish identical twins.

Ariyaeinia *et al.* [61] performed recognition experiments using voice data from 49 pairs of identical twins. The authors performed basically two different experiments: a general experiment, in which any two persons in the dataset were considered impostors, and the twin experiment, in which the impostor tests consisted of comparing a person and his/her twin. The Equal Error Rate reported was 1.0% for the twin experiment using short interval voice pattern (each person saying his/her date of birth), and 0.5% for the general configuration. Other authors have tried to distinguish identical twins based on voice in a multilingual environment [64]. Using a database of 12 twins, Patil and Basu reported the highest success rate as being 100% for a particular size (60 seconds) of the training speech, and particular size (in seconds) of the test speech. They also observed that the majority of errors were due to matching the actual speaker with his/her twin brother/sister.

However, in most of the previous studies, the identical twin biometric database is very small and an in-house biometric matcher instead of COTS matchers were used. This affects the reliability of conclusions. Furthermore, no previous study was conducted to compare the performance of

different biometric traits and fuse multiple biometric traits of identical twins.

A number of studies and methods related to biometrics of twins has been published since our work in multibiometrics of twins appeared [67]. They are summarized below.

Hollingsworth *et al.* [68] suggested that genetically identical irises (from identical twins or left and right irises of a person) might be more similar than genetically unrelated irises. They conducted experiments on iris texture similarity in which humans viewed pairs of iris images. There were basically two experiments: in the first experiment, volunteers had to provide their opinion on whether the pair of irises they were viewing was from the same person (left and right) or from different persons; in the second experiment, volunteers had to provide their opinion on whether the pair of irises they were viewing was from identical twin pair or from unrelated persons. The results showed that humans were successful in these two tasks in more than 80% of the cases. These results imply that, although iris biometrics can distinguish between identical twins, some similarity is still present. It should be noted that humans are not necessarily good at matching iris textures.

Tao *et al.* [74] analyzed the performance of two state-of-the-art fingerprint matchers on identical twins by using a database of 83 identical twin pairs, 4 fingers per person, and 6 impressions per finger. Their conclusion that fingerprint matchers can distinguish between identical twins with a small drop in performance agrees with previous studies. They also analyzed and compared the probability of fingerprints from identical twins having the same pattern. This probability was found (as in [58]) to be much higher compared to the probability of fingerprints from unrelated persons having the same pattern – 0.7440 vs. 0.3215. Another conclusion of this study is that all fingers have the same probability of being of the same fingerprint type in identical twins.

Phillips *et al.* [7] conducted matching experiments using face images of 126 pairs of identical twins. They used three state-of-the-art face recognition systems, and the experiments were conducted on face images captured under two different lighting conditions, indoor (or studio lighting) and outdoor (or ambient lighting), and with two facial expressions, neutral and smile. Their experiments led to the conclusion that the changes in ambient lighting and expression largely affect the

performance of the systems; the studio to studio lighting matching yielded the best performance, while the ambient light conditions matching yielded the worst performance. As an example, at a False Acceptance Rate of 0.01, the verification rate of studio to studio lighting with neutral expression is about 0.85, while for the ambient lighting also with neutral expression is about 0.05. Their experiments also showed that, if we fix the lighting to studio conditions, the Equal Error Rate (EER) in the case of expression changes is significantly higher for the three algorithms. In summary, recognition of identical twins in studio lighting with neutral expression is promising, while when the lighting or the expression change, the performance is drastically reduced.

Pruitt *et al.* [69] conducted experiments similar to [7]. They used the same database used in [7], and also analyzed the influence of different expressions and lighting using four face recognition systems, and reaching the same conclusions as in [7]. In addition, they analyzed the influence of the presence of eyeglasses, and they found that the performance is about the same.

Biswas *et al.* [70] also concluded that uncontrolled environment affects the ability to distinguish between identical twins, but their experiments analyzed the human ability. Their experiments were conducted using face images from a database of 186 subjects; 23 volunteers were asked to look at pairs of face images for a limited time and to decide whether the pair of face images came from the same person or from identical twins. Humans could correctly distinguish between identical twins with an average accuracy of almost 79%. Another interesting result was that humans used moles/scars/freckles as the most important features in their correct decisions.

Park and Jain [76] applied their facial marks detection and matching algorithm to a few identical twins face images, and showed that facial marks helped in distinguishing between twins. More recently, Srinivas *et al.* [73] presented an analysis and algorithm to distinguish between identical twins by using facial marks. In their work, they analyzed and matched manually marked and automatically extracted facial marks. Their results showed that facial marks can help in distinguishing between identical twins; however, it is more difficult to distinguish identical twins than distinguishing between unrelated persons. Their results also suggested a correlation in the distribution of facial marks in twins.

Klare *et al.* [71] analyzed the effect of different features in the identical twin recognition based on face images in the database used in [7]. They analyzed the importance of different components of the face, the improvement using facial marks, and the differences in performance when the systems were trained on images of identical twins compared to unrelated persons. Their conclusions include: (i) all facial components seem to have the same importance in identical twins matching scenario as in the standard matching one, (ii) the use of facial marks improves the discriminability of identical twins without reducing the accuracy on unrelated persons, and (iii) regarding discriminative learning, it appears that training on twins does not help the overall system, but this might be only because the number of twins available for training is not sufficiently large.

Vijayan *et al.* [72] presented the results of applying four state-of-the-art 3D face recognition algorithms to an identical twins dataset (3D Twins Expression Challenge) containing 3D scans of neutral and smiling expressions of 107 pairs of identical twins. They conducted experiments on different cases based on the expressions of the probe and gallery scans of each of the twins. In the two scenarios in which the gallery had one expression and the probe had a different expression, the best performing algorithm presented an Equal Error Rate (EER) of 0.2% and 0.5% and a rank-1 identification rate of 98.1% for both scenarios. In the scenarios with uncontrolled expressions (one twin in the gallery has the same expression as the other twin in the probe, and vice versa), the EER was 0.8% for both the cases, and the rank-1 identification rate was 91.6% and 93.5%, respectively. Their results argue that 3D face recognition remains an open problem if different facial expressions are present.

2.3 Multibiometrics

A multibiometric system uses multiple sources of biometric information in order to recognize an individual. For example, a multibiometric system may use fingerprint and face, the left and right iris, or a fusion of different fingerprint matching algorithms to recognize a person. In the next subsections, we focus our attention on the distinctiveness of fingerprints, face, and iris for identical twins. These three modalities will then be used in our multibiometric experiments.

2.3.1 Fingerprint Discriminability

A fingerprint is the impression of the friction skin on a finger. The individual characteristics of friction ridge skin are determined during fetal development. Their formation is similar to the formation of blood vessels or capillaries during the growth of the fetus in the uterus. The fingerprint formation starts at approximately 6 or 7 weeks of gestational age and it is due to the flow of amniotic fluids in a micro-environment [22]. A minor change in this flow and in the position of the fetus in the uterus cause the minute skin structures around palm, finger tips and soles to differentiate. Friction ridge skin can be distinguished from the skin of the rest of the body due to a variety of factors, such as the presence of raised ridges, increased sensory abilities, absence of hair and sebaceous glands, and a thicker and more complex epidermis. Friction ridges are useful for grasping and gripping objects, which explains their presence in our hands and feet.

As mentioned in Chapter 1, a fingerprint pattern has some details that are present in each individual fingerprint, like ridge endings, the point where a ridge ends abruptly, or a ridge bifurcation or trifurcation, where the ridges are divided into different branches. However, collectively, these details (minutiae) are supposedly different in every fingerprint, even in prints of identical twins, since a very small difference in micro-environment is sufficient to change the process of cell formation, causing minutiae points to be different. As a result, fingerprint is considered very reliable in terms of biometric identification because of its distinctiveness. Another property of fingerprint that makes them useful in biometric identification is that fingerprints do not change significantly over time, an essential characteristic of a biometric modality since a biometric system is typically meant to be used to identify a person over a long period of time.

2.3.2 Face Discriminability

Face is composed of the skull characteristics and the musculature and associated soft tissue. To study the variation among human faces, it is necessary to study these structures. The facial skeleton serves as the bony framework for the mimetic musculature. Since these muscles are stretched across the facial skeleton like a mask, the variation in facial appearance is caused mostly because

of the form of the facial bones. A majority of individuals can be divided in three categories in terms of facial appearance: (i) a long, narrow head, (ii) a proportional length to width head and (iii) a short, wide head; these differences in facial shapes are due to the form of the cranial base [22].

Facial form is influenced by gender. Males are more likely to have larger faces because of their usually larger bodies and their need for more air in order to support larger muscles and viscera, causing them to also have a larger nose. Males also usually have a more protrusive forehead. Besides the gender influences in the face of an individual, a face also changes when a person ages. Infant faces tend to be wide and short because of the development of the brain; over time, the face develops and this wide and short face of the baby tends to change. Other effects of aging are dehydration of the dermis and reabsorption of subcutaneous fat deposits, which result in a decrease in the facial volume and wrinkling [22].

The muscles may vary in their presence (not everyone has all the muscles that could be in a face), form, location, and control. These factors influence the kind of facial movement that an individual can create. Furthermore, the facial movements of an individual change his/her face as he/she ages. With aging, the elasticity of the skin decreases and the face then is marked with the expressions that occur frequently, becoming relatively permanent features. This fact may explain why identical twin faces are more likely to be distinguished as they are older than when they are infants [22].

There are a number of factors that influence the performance of a facial recognition system. Besides those already cited above, there are differences in pose, illumination, expression, occlusion, accessories like glasses, weight changes, hair style changes, etc. All these variations make facial recognition systems not as accurate as some other biometrics, like fingerprint and iris. Identical twins present a particularly difficult situation for face recognition systems, since they are usually extremely similar in facial appearance.

2.3.3 Iris Discriminability

The iris is an annular shaped part of the eye between the pupil and the sclera (see Fig. 2.2) that regulates the amount of light entering the eye through the pupil. The iris is physically small in size (about 11 mm) but well-designed optical systems can magnify a human iris into a high-resolution image that is 200 to 300 pixels in diameter. There are many minute features such as freckles, coronas, stripes, furrows and crypts, etc. randomly distributed in the iris region, which constitute the unique iris texture for each eye.

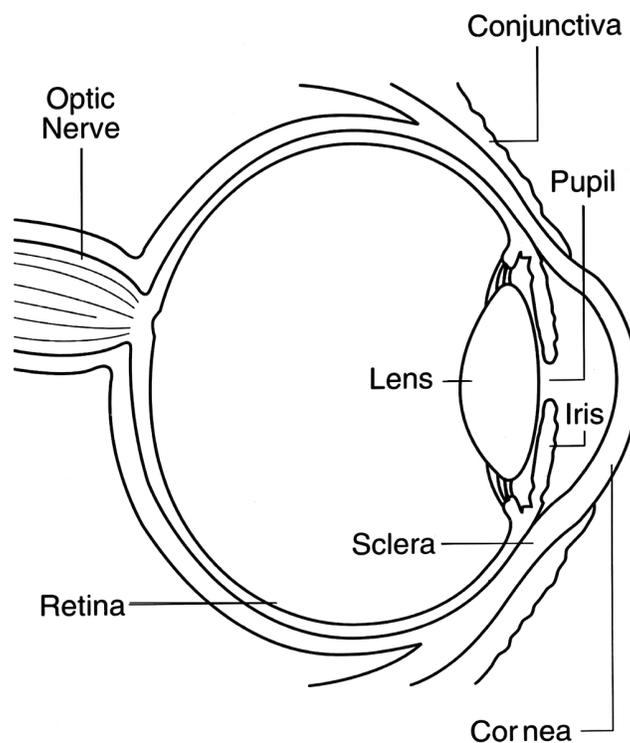


Figure 2.2: A diagram of the human eye ².

The iris texture pattern is formed and becomes stable after the eighth month of gestation. It is commonly believed that the formation of iris pattern is determined by the gestation environment, i.e. iris is a phenotypic biometric trait [62]. So even identical twins can be discriminated using suitable iris features; even the irises of left and right eyes of the same person are different.

²Figure from <http://www.nei.nih.gov/photo/>.

Random processes involved in iris development determine the irregular shape and random spatial layout of micro anatomical structures (MAS). The MAS in the iris surface may exhibit different reflectance properties in the near infrared light, leading to sharp intensity variations across iris image region. We can regard each iris region as a piece of 2D surface in a 3D coordinate system. So the surface shape from valley to peak resembles an odd Gabor filter and the shape from valley to peak then valley resembles an even Gabor filter, and vice versa. Daugman [62] proposed to match each iris region with even or odd Gabor filters and then encode positive correlation as 1 and negative correlation as 0. Thereby, an iris image can be represented by 256 bytes, called the iris code. The iris codes of twins (or non-twins) have only 50% chance of being matched because their corresponding regions independently have an equal probability to be either 1 or 0 in iris feature coding. In contrast, multiple iris images of the same eye have a much higher probability than 50% to be matched in their iris codes, even though noise may perturb some parts of iris codes. A more general explanation of the effectiveness of iris recognition for twins is based on ordinal measures [77].

2.4 Experimental Results and Analysis

In this section, we first describe the databases used in our experiments, followed by an analysis of the ability of biometric systems in distinguishing between identical twins and their ability in determining whether two subjects are identical twins based on their biometric traits, namely face, fingerprint and iris.

2.4.1 Databases

The first version of the CASIA Multimodal Biometrics Database of Twins (CASIA-TwinsV1) was collected in October 2, 2007 at the Beijing Chaoyang Park during the Fourth Annual Festival of Beijing Twins Day. Figure 2.3a shows the kiosk where the biometric acquisition was performed and Figure 2.3b shows the face acquisition device. This database includes face, iris, and fingerprint

images from 92 pairs of twins and 2 sets of triplets. All the images were captured indoors (inside a tent) on the same day (i. e., single data capture session).



Figure 2.3: The kiosk for biometric acquisition (a) and the face acquisition device (b).

Not all of the 94 sets of twins or triplets provided images for all the modalities. Since some of our experiments involved combinations of units/modalities, we considered only those individuals who have a complete set of images (face, two irises and four fingerprints) in the database, and whose twin's images were also present in the database as a complete set. As a result, the total number of subjects used in our experiments consisted of 134 subjects (66 families, including two families of triplets – 51 pairs of identical twins and 15 pairs of non-identical twins). For all biometric modalities (four fingers, two irises, and face), the number of genuine matches performed was 134, the number of identical twin impostor matches was 102, and the number of general impostor matches was 17,720.

The second version of the CASIA Multimodal Biometrics Database of Twins (CASIA-TwinsV2) was collected in 2009. This second version also includes face, iris, and fingerprint images from 59 pairs of twins and 1 set of triplets. Not all of the 60 families provided images for all the modalities. This version is more challenging than the first one since the images were captured indoor and outdoor, and the face images contain variations such as lighting, pose and expression. By themselves, these variations already pose a great challenge to face recognition.

Most of the twins are identical (or monozygotic twins), but some are non-identical (or dizygotic) twins. This information was not recorded, so we derived this information based on observing

whether the face images of a set of twins were very similar or not³. In the FBI's collection of biometric data from twins, the information whether the twins were identical or not was revealed by the subjects, i.e., the subjects were asked whether they were identical or non-identical twins. Both methods are not perfect and subject to errors. For example, non-identical twins can have very similar appearance which would cause them to be classified as identical twins based on facial similarity; or twins might report identical or non-identical twinning based on their guesses or similar appearance rather than based on DNA testing. We divided our database into two groups, identical twins and non-identical twins. Most of the subjects in the database are children, but there are some adults as well. The subject age ranges from 5 to 65 (5 to 63), with the average age being 16.8 (16.2) in the first (second) version of the database. There are 12 families for which their biometrics were collected in both versions. We used matching scores from iris and face images to find the subjects who were present in both collections.

In the following sections, we present more detailed descriptions of the two databases separated by the biometric modality, the experiments performed, and experimental results on the two topics: distinguishing identical twins and finding similarities between identical twins.

2.4.2 Distinguishing Identical Twins

2.4.2.1 Fingerprint

The fingerprint images were captured using Symwave sw6888 [78], a sweep sensor, with a resolution of 500 dpi. This dataset contains images from four different fingers and the number of images per finger is not fixed, but varies from 6 to 31. Because of this large variability and the poor quality of many images, we selected one image per finger as the template and another image from the same finger as the query image that contain the largest number of minutiae.

Figure 2.4 shows some fingerprint images from the twin dataset. Figure 2.4a shows images of the four fingers 1, 2, 3, and 4 for the first twin of an identical twin pair; Figure 2.4b shows the

³The ideal method to distinguish identical or non-identical twins should be based on DNA.

fingerprints of the corresponding fingers for the second twin of the pair. Figures 2.4c and 2.4d follow the same scheme for a non-identical twin pair. The similarity between ridge flow patterns of corresponding fingers of identical twins is evident. While all four pairs of corresponding fingers of identical twins have the same fingerprint pattern type (left loop, left loop, right loop, right loop for fingers 1 to 4, respectively), only two pairs of corresponding fingers of non-identical twins have the same pattern type (left loop, right loop for fingers 1 and 3, respectively).

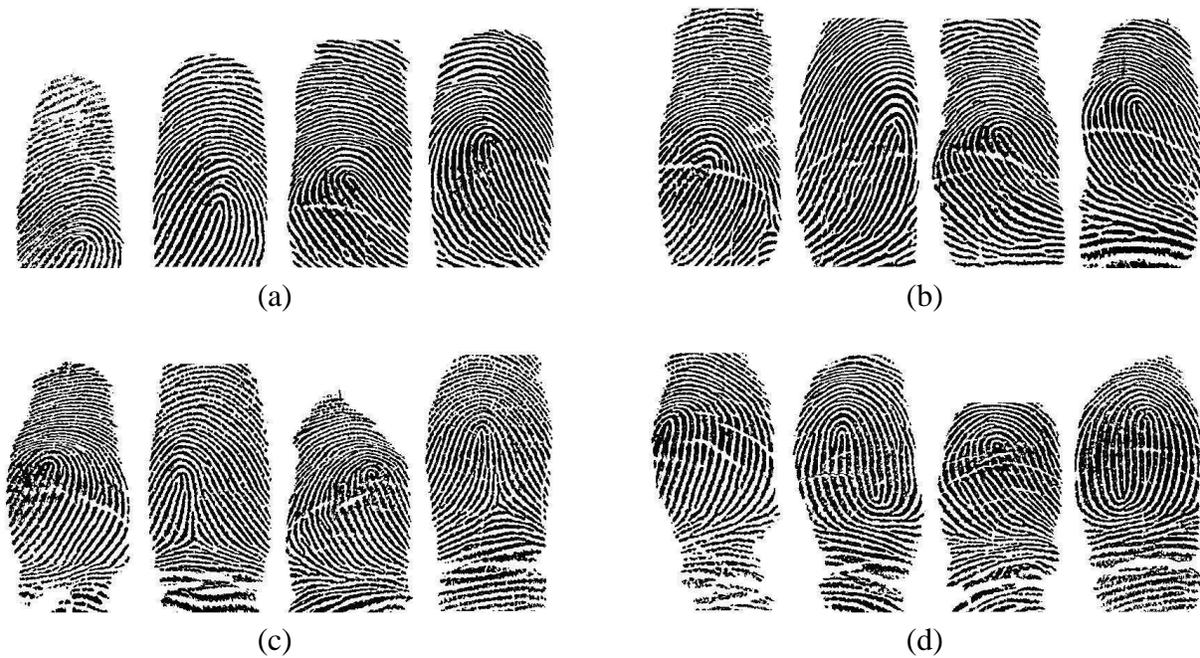


Figure 2.4: Fingerprint images of fingers 1, 2, 3, and 4 of the first twin (a), and the four images of the corresponding fingers of the second twin in an identical twin pair (b); similarly, (c) and (d) show fingerprint images of a non-identical twin pair. Note the similarity in ridge flow pattern between identical twins. All four corresponding fingers of identical twins in (a) and (b) have the same pattern type. But for non-identical twins in (c) and (d), only two corresponding fingers (no. 1 and 3) have the same pattern type.

A minutiae based commercial matcher, VeriFinger 4.2 [79], was used to obtain the matching scores in the fingerprint experiments. In our experience with Verifinger 4.2 match scores, there are no impostor scores greater than 300, so scores greater than 300 were set to this number, and then the scores were normalized from $[0, 300]$ to the range $[0, 1]$.

Three different match score distributions were generated in order to analyze the results of the experiments. The genuine distribution was obtained by matching the gallery image of one modality

of one individual to the probe image of the same modality and the same individual. Identical twin impostor distribution was generated by matching each image of one person to his/her identical twin. The general impostor distribution was generated by all the impostor matches except the identical twin impostor matches described above. This scheme was used for all the experiments.

Figures 2.7a and 2.7b show the genuine, identical twin, and the general impostor distributions of finger no. 3, which had the worst performance among the four fingers, and finger no. 4, the best performing finger, respectively. We can see that, although identical twin impostor distribution of finger no. 3 is similar to the general impostor distribution of this same finger, there are some peaks in the right tail of the identical twin impostor distribution that differentiates them. This is an indicator of the larger similarity between fingerprints of identical twins compared to the similarity between fingerprints of unrelated persons. This is also the case for finger no. 4. The Kolmogorov-Smirnov 2-sample test [80] indicates that the two samples (identical twin and general impostors) do not come from the same population with a significance of 0.05.

Another indicator of this large similarity between identical twin fingerprints is the number of matched minutiae. For example for finger no. 4, the mean of the number of matched minutiae for genuine pairs is 22.00 (± 9.50), for identical twin pairs, 6.44 (± 3.97) and for general impostor pairs, 4.22 (± 2.82). For the other three fingers, these numbers of matched minutiae are about the same. Although the number of matched minutiae for identical twin pairs is much smaller than for genuine pairs, it is still larger than the number of matched minutiae for general impostor pairs, which indicates the similarity between identical twins' fingerprints is larger than between unrelated persons. The distribution of the number of matched minutiae between fingerprints of identical twins is significantly different from the distribution of the number of matched minutiae between general impostors (based on the Kolmogorov-Smirnov test).

2.4.2.2 Face

The face images were captured in color, all of them from a USB camera. The image size is 480×640 pixels, but the face area in the image varies from 280×300 to 300×400 pixels. The images

contain non-uniform background, and some variations in illumination. The images were captured in a sequence and over a short time interval. There are about 20 face images per subject. Some face image examples are shown in Figure 2.5. The first two images show face images of an identical twin pair (Figure 2.5a), while the other two are face images of a non-identical twin pair (Figure 2.5b).

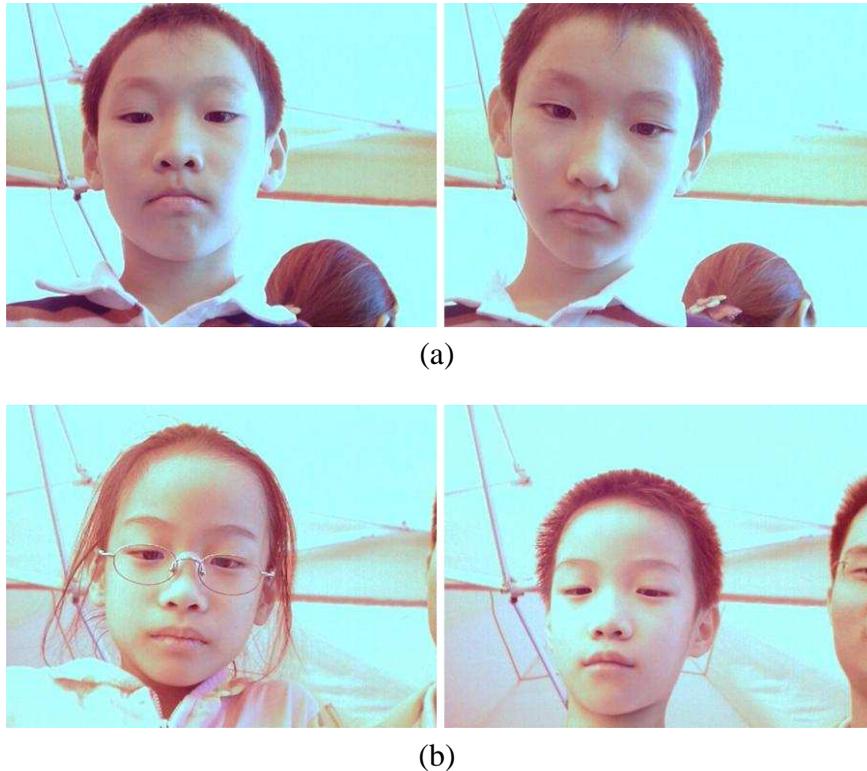


Figure 2.5: Face images of the first and second twin in (a) an identical twin pair, and (b) a non-identical twin pair.

The face subset used in the experiments contained 134 subjects, each one having around 20 images. The commercial matcher used to perform the facial matching was FaceVACS 7.1 [81], and the scores from this matcher are in the range $[0, 1]$. We considered two face images per person (the template and the query) in our experiment, mainly because the pictures were taken over a very short time interval, which makes them very similar. The first and the last images of each person were used, since they are expected to be the least similar. Figure 2.7c shows identical twin and general impostor distributions, along with the genuine distribution for the face experiments. In

the genuine distribution, we can see that there is almost no tail and that about 90% of the genuine match scores are of more than 0.95. This is due to the high degree of similarity between the two face images of the same person, since they were taken in sequence and in a very short time interval. Also, we can observe that the identical twin impostor distribution is more similar to the genuine distribution than to the general impostor distributions, meaning identical twins are a real challenge to face recognition systems.

2.4.2.3 Iris

The iris images were captured using an IKEMB-100 camera produced by IrisKing [82]. The size of the images is 480×640 pixels, but the approximate iris diameter is 200 pixels. IKEMB-100 is an embedded dual-eye iris camera and has an LCD displaying real time iris images for user convenience. The iris images were captured in sequence and over a short time interval. There are 10 images for most of the subjects, with a few individuals having a smaller number of images. Some examples of iris images are shown in Figure 2.6. Figures 2.6a and 2.6b show iris images of an identical twin pair, where the two images in 2.6a are the left and right iris images of the first twin, and the two images in 2.6b are the left and right iris images of the second twin. Similarly, Figures 2.6c and 2.6d show iris images of a non-identical twin pair.

The iris feature representation method based on ordinal measures [77] is used to test the performance of iris recognition for twins. The match scores range from 0 to 1. Two iris images were randomly selected as probe and gallery for each eye. Figure 2.7d shows the genuine, identical twin, and general impostor distributions of the right iris, which performed slightly better than the left iris. The identical twin impostor distribution is very similar to the general impostor distribution. However, the peaks that are present in the identical twin impostor distribution tail may indicate that the irises of identical twins have some correlation. The Kolmogorov-Smirnov statistical test indicates that the identical twin and general impostor distributions are significantly different (significance of 0.05).

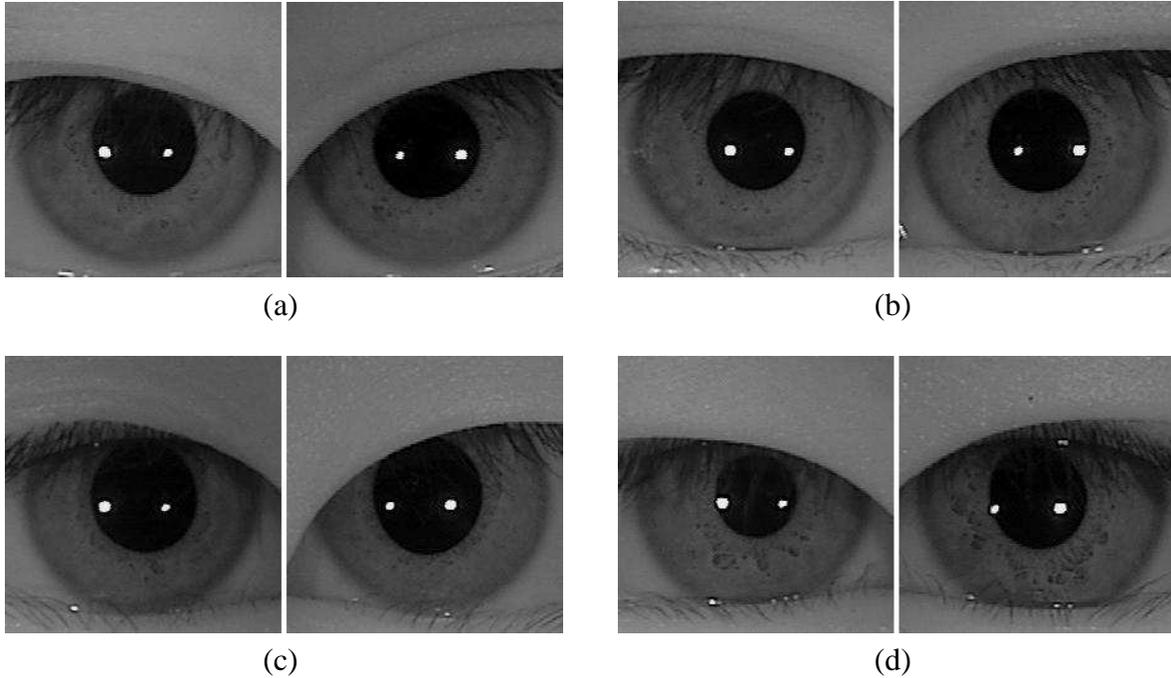


Figure 2.6: The left and right iris images of identical ((a) and (b)) and non-identical twin pairs ((c) and (d)).

2.4.2.4 Multibiometric Experimental Results

In our multibiometric experiments, we first combined units of the same modality, like the two irises or all the four fingers of a person. We performed a simple fusion using scores from the four fingers by summing them. The identical twin and general impostor distributions, along with the genuine distributions, are shown in Figure 2.8a. A fusion using both the irises was also performed, and the distributions are shown in Figure 2.8b. We can observe from these figures that the matching results are extremely good; genuine and impostor distributions are well-separated. This indicates that multimodal biometric systems can work very well even in the presence of identical twins in the biometric database.

Two different Receiver Operating Characteristic (ROC) curves were generated for each experiment. Identical twin impostor ROC curve means the impostor matches considered were just identical twin impostor matches, while a general impostor ROC curve means we considered all the impostor matches, except the identical twin impostor matches. Figure 2.9a shows the ROC curves for fingers nos. 3 and 4, the best and worst performing fingers, respectively, and the ROC curves

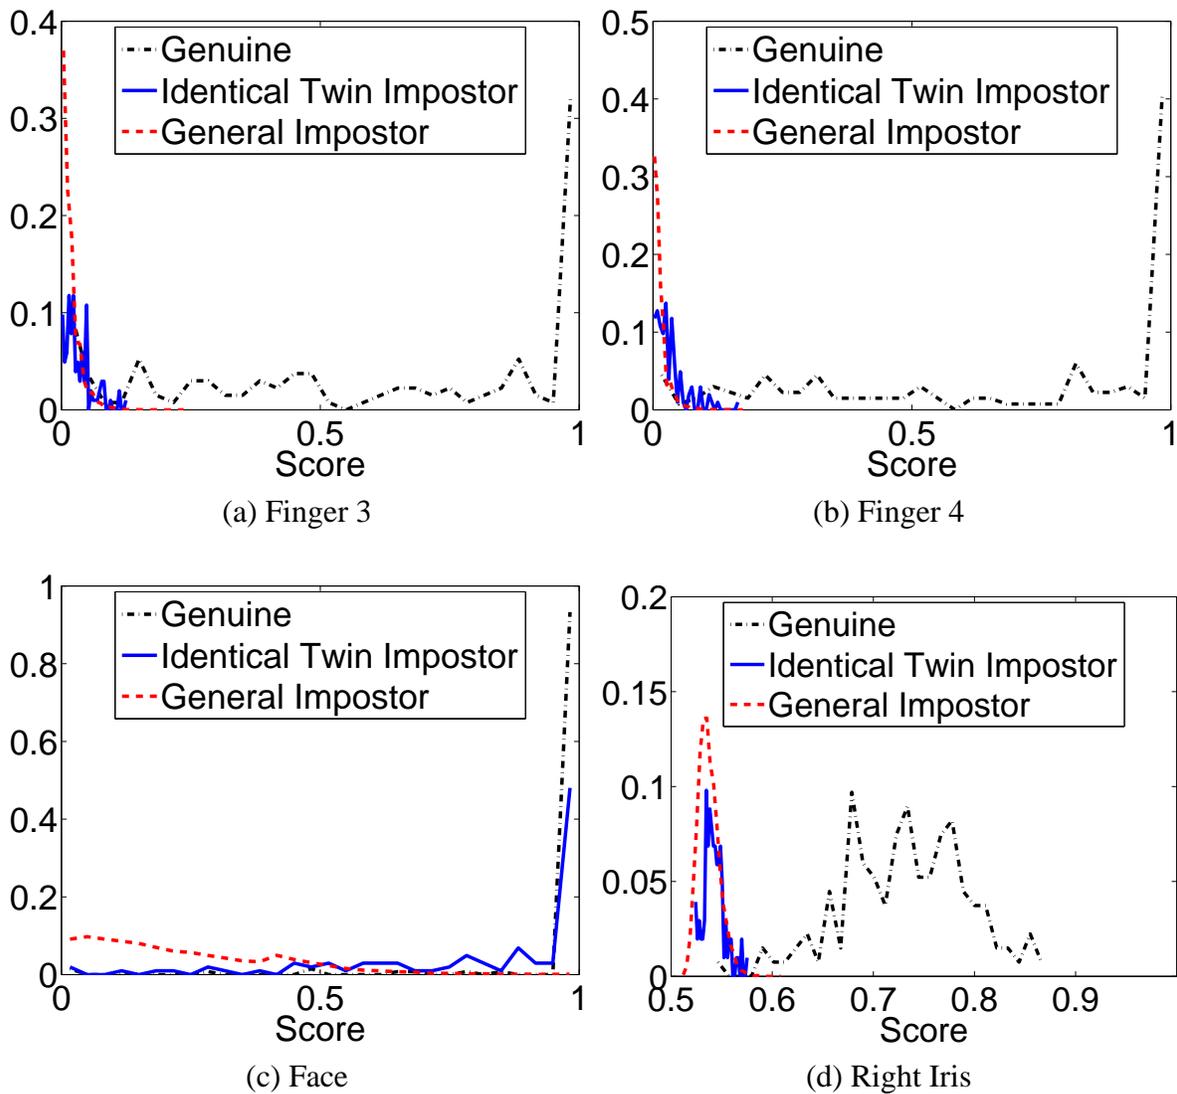


Figure 2.7: Identical twin, general impostor distributions and genuine distributions for fingers 3 and 4, face, and right iris.

of the 4-finger fusion. The performance of the fingerprint identification system in distinguishing genuine matches from general impostor matches is always better than the performance in distinguishing genuine matches from identical twin impostor matches. It is important to note that the 4-finger fusion had the highest performance among the experiments involving fingerprints.

The iris matching experiment results show that the iris biometric system can distinguish identical twins as much as it can distinguish any two different persons who are not identical twins, as shown in Figure 2.9b. Iris experiments presented the best performance among all the experiments

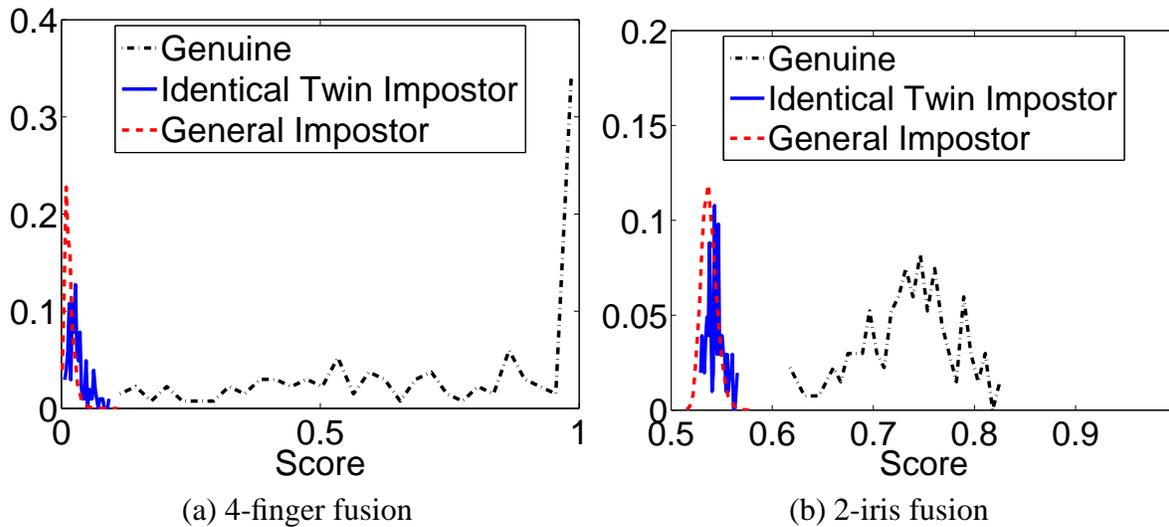


Figure 2.8: Identical twin, general impostor, and genuine distributions of the 4-finger fusion and 2-iris fusion, respectively.

that used just one biometric characteristic (one finger, one iris, or face). This may be because irises from identical twins are more uncorrelated than fingerprints from identical twins, or may be due to the fact that iris images in our dataset have a very good quality due to automatic rejection of poor quality images by quality control software, while the quality of many fingerprint images is very poor due to improper sweep operation by some of the children, and the image area is very small. Also, the 2-iris fusion showed improvements compared to the performance of each iris alone.

We also fused face and finger no. 4, where the fingerprint scores were appropriately normalized, and the results are shown in Figure 2.9d. We did not combine iris with another biometric modality because iris results are already really good and the database is not large enough for measuring lower error rate. The ROC curves of the face experiments are shown in Figure 2.9c. The face experiment shows that the presence of identical twins in the face database causes the face recognition performance to drop. Although the performance was good based on the general impostors (no identical twins in the data involved), we can only get a True Acceptance Rate (TAR) for the identical twin data greater than zero at a false acceptance rate of over 10% for identical twins. Compared to face matcher, the fusion of face and finger no. 4 has a better performance in dealing with identical twin and general impostors. Compared to fingerprint matcher, this fusion degraded

the performance on identical twin impostors, but improved the performance on general impostors. It should be noted that the inferior performance of fingerprint matcher compared to face matcher in our experiments is due to the specific fingerprint sensor (swipe sensor) used for data collection and the difficulty of obtaining good quality fingerprint images of younger subjects using the swipe sensor. On the other hand, there is very small intra-class variation for face images.

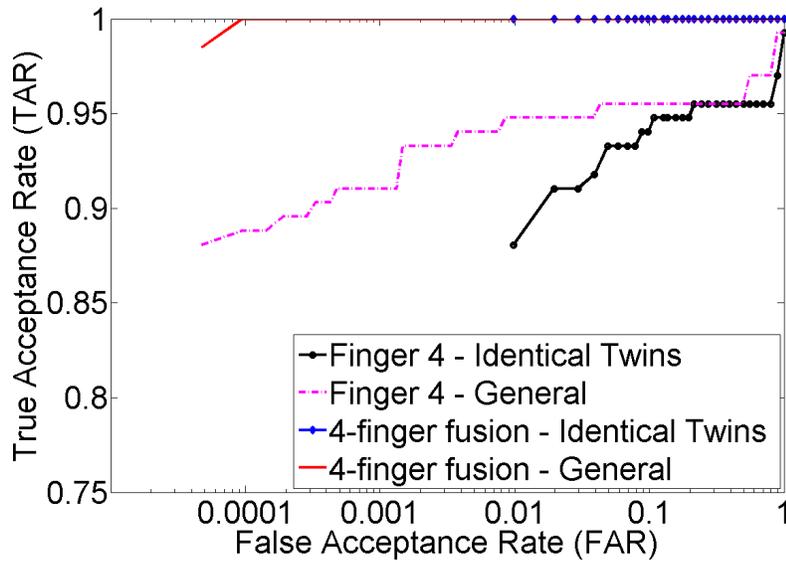
Table 2.2 shows the equal error rates for the worst and best performing fingers (finger numbers 3 and 4), 4-finger fusion, left and right irises, iris fusion, face, and face and finger 4 fusion; the equal error rates are shown based on the identical twin impostors and general impostors separately. It can be observed that all the modalities have a better performance on the general population than on the identical twin population.

Table 2.2: Equal Error Rate (%) for distinguishing (i) genuine vs. impostor identical twins and (ii) genuine vs. general impostors based on different biometric features.

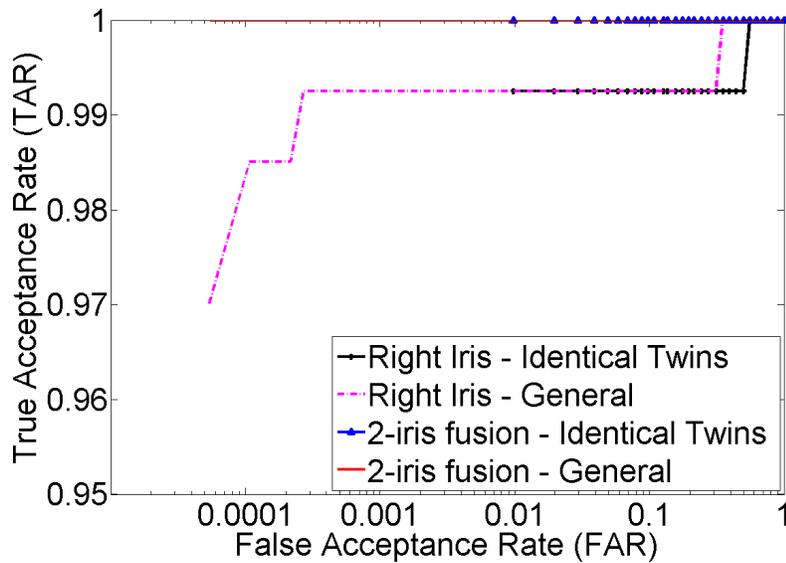
Modality	Finger 3	Finger 4	4-finger fusion	Left Iris	Right Iris	2-iris fusion	Face	Face+ Finger 4
Identical twins	13.95	6.79	0.49	1.35	0.86	0.49	13.67	7.65
General impostors	10.71	4.40	0.00	0.75	0.75	0.00	3.79	2.48

2.4.2.5 Discussion

In the previous sections, we presented experimental results based on state-of-the-art biometric recognition systems available in 2009. In this section, we discuss more recent results to indicate the development of the systems overtime. In Fig. 2.11, we show ROC curves for face recognition using FaceVacs 8.6, applied to both versions of CASIA Twins database (see Fig. 2.10 for examples of face images in CASIA Twins V2). It can be noted that face recognition performance shows an overall improvement given the controlled environment of CASIA Twins V1. However, the challenge of the illumination, pose and expression variations is evidenced in Fig. 2.11b, which



(a) Fingerprint



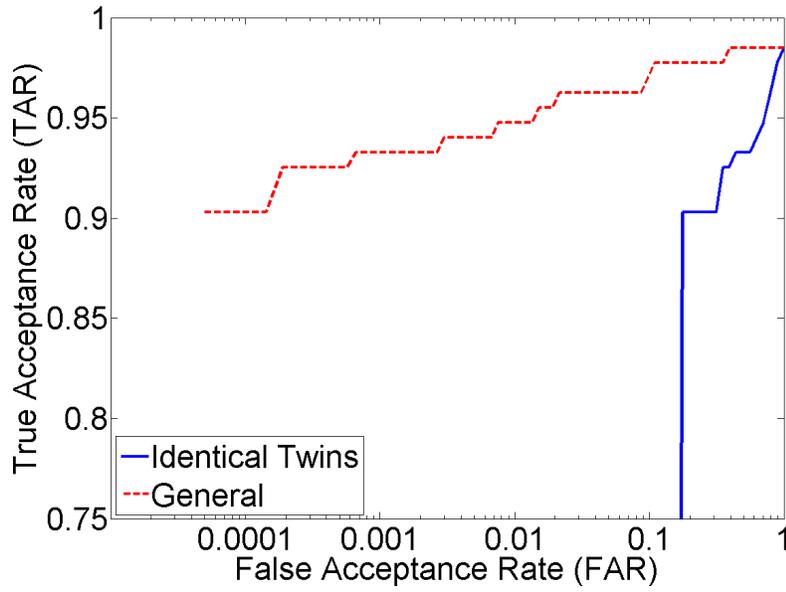
(b) Iris

Figure 2.9: ROC curves for fingerprint, iris and face, and a fusion example (face + finger 4). Due to the small number of identical twin impostor (102), FAR less than 1/102 cannot be estimated.

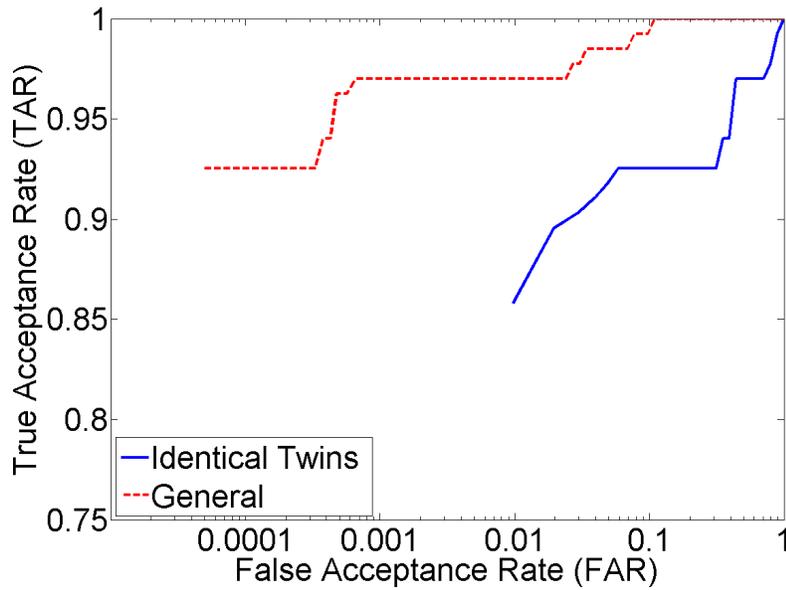
shows the true acceptance rates at various false acceptance rates in the CASIA Twins V2 database. In this unconstrained environment, it is even more difficult to distinguish between identical twins.

In the iris and fingerprint cases, we repeated the experiments using VeriEye [83] and Verifinger 6.3 [79]. The ROC curves considering identical twin impostors and general impostors are very

Fig. 2.9 (cont'd)



(c) Face



(d) Multimodal Fusion

similar for iris, as it was the case in our previous experiments; for fingerprints, the small drop in the performance is still present when identical twins impostors are considered compared to general impostors.



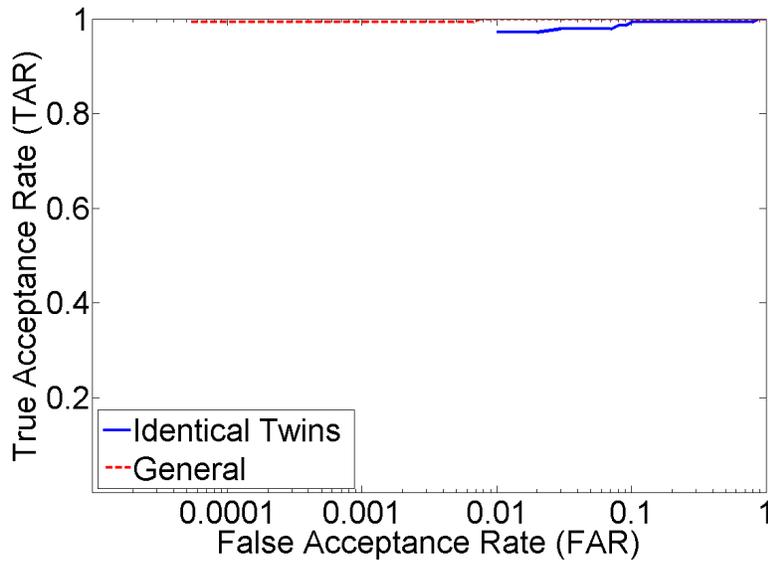
Figure 2.10: Examples of face images in CASIA Twins V2 database.

2.4.3 Finding Similarities between Identical Twins

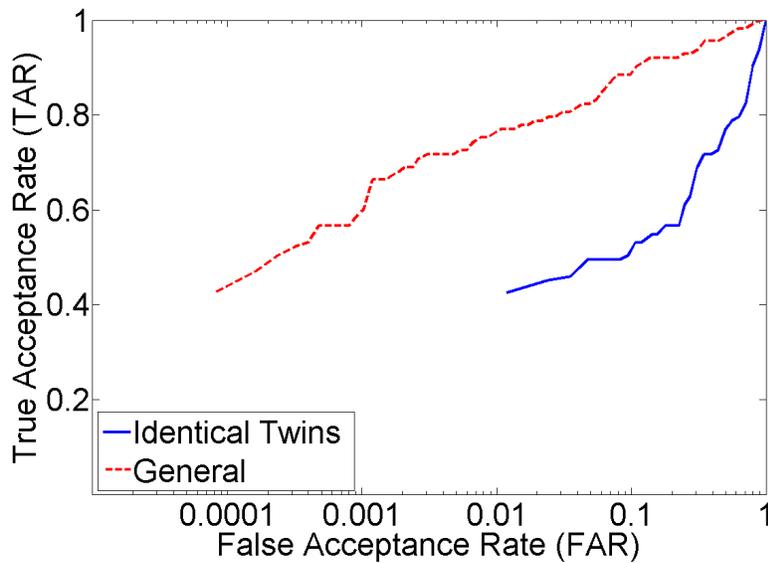
In the previous section, we discussed the ability of biometric systems to distinguish between identical twins to answer the following question: can biometrics be used to distinguish identical twins? A complementary question of interest is: can biometrics be used to detect identical twins? The answer to this question is not known. The goal of this section is to provide quantitative measures that will help in answering this second question.

Identical twins are the most genetically similar persons since they share the same DNA. Studies on the characteristics of identical twins can generally lead to improvements in the overall knowledge about unrelated persons. For example, the knowledge about the influence of the genetics in the occurrence of a certain disease can lead to the discovery of environmental factors in addition to genetic factors. This type of knowledge can also be useful in biometric recognition. In [68], the authors discussed the similarity between irises of identical twins. In their paper, they found that humans can envision some type of texture similarity between irises of identical twins, and therefore to determine genetic relations from iris images might be possible. This, in turn, could help in identifying victims without any form of identification, by finding genetically related persons.

In general, studies on the similarities between biometrics of twins can lead to a broader understanding of the similarities between biometrics of the general population. Starting with twin similarities, it can possibly be extended to find other blood (familial) relations. In addition, the protection of biometric data should be taken seriously since by stealing a template, one might be



(a) CASIA Twins V1



(b) CASIA Twins V2

Figure 2.11: ROC curves for face using FaceVacs 8.6.

able to find genetically related persons. In the general research scenario, identical twins are of great interest. If biometrics can be used to correctly determine whether two subjects are identical twins, it might be much less expensive and quicker than DNA comparison.

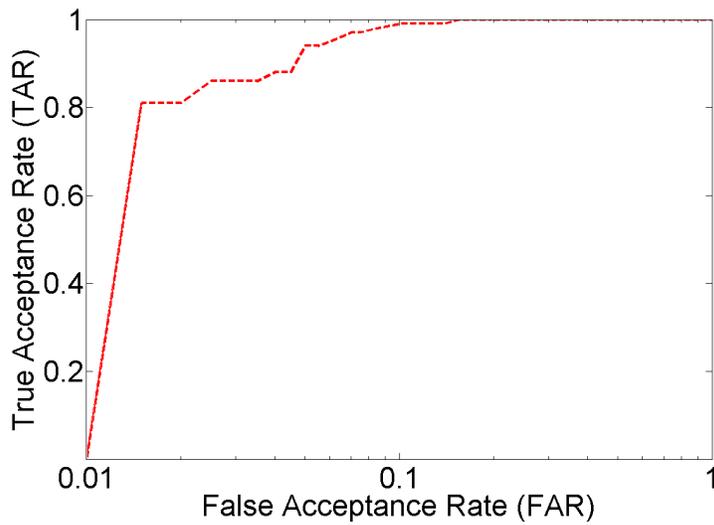
Suppose we have N genuine biometric sample pairs, N identical twin impostor pairs, and N general impostor pairs. Among these pairs, how well can we determine the identical twin pairs?

By using face and iris modalities together, we are able to correctly detect 80% of the identical twins while falsely accepting 2% and 10% of genuine or general impostor pairs as identical twins in database CASIA Twins V1 and V2, respectively (see Fig. 2.12). Due to the different characteristics of the databases, it can be said that, in conditions closer to frontal face images and indoor lighting, we can correctly determine 80% of the identical twin pairs, while falsely accepting 2% of general subject pairs as identical twins. When the lighting, pose and expression conditions vary, the false acceptance rate goes up to 10%.

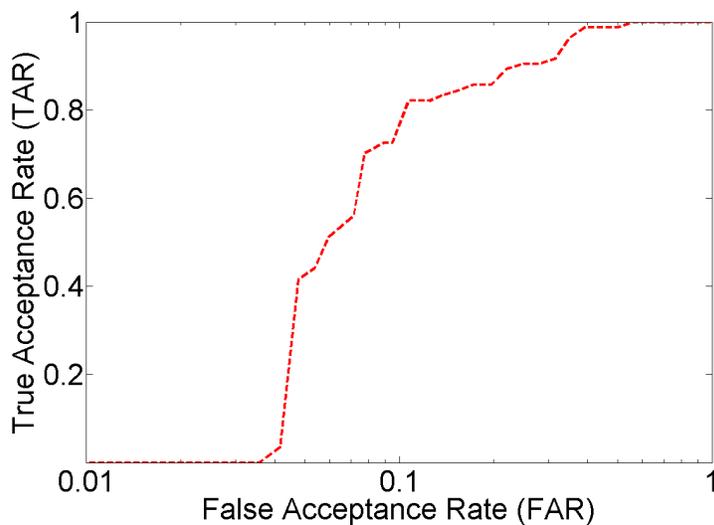
We considered the following scheme to detect identical twins: if iris score between the subjects is more than a threshold, the pair is considered a genuine pair, and the measure for identical twin determination is set to zero. If iris score is less than some threshold, then the pair is considered an impostor pair (note that identical twins fall under this category). Because of the high similarity between faces of identical twins, the measure for identical twin determination is then set to the similarity between the faces of the subjects. For this experiment, we randomly choose two face images and two images of the same iris (right) for each subject. When using fingerprints to determine genuine pairs instead of iris, the performance is about the same for the CASIA Twins V2, while the performance drops for CASIA Twins V1 mainly because the quality of the fingerprint images in V1 is worse than in V2 (see Fig. 2.13).

Fig. 2.14(a) shows face images of an identical twin pair for which the identical twin evidence is extremely high due to the similarity in their facial appearance and dissimilarity in their iris textures. However, the twin pair shown in Fig. 2.14(b) who was classified by us as a non-identical twin pair also shows an extremely high evidence of being identical twins according to face and iris.

Our work focused on answering the question: given a pair of subjects, can we determine whether they are identical twins by means of biometrics? Another more general way of approaching the problem is to come up with a likelihood that a given pair of subjects is an identical twin pair based on their biometrics.

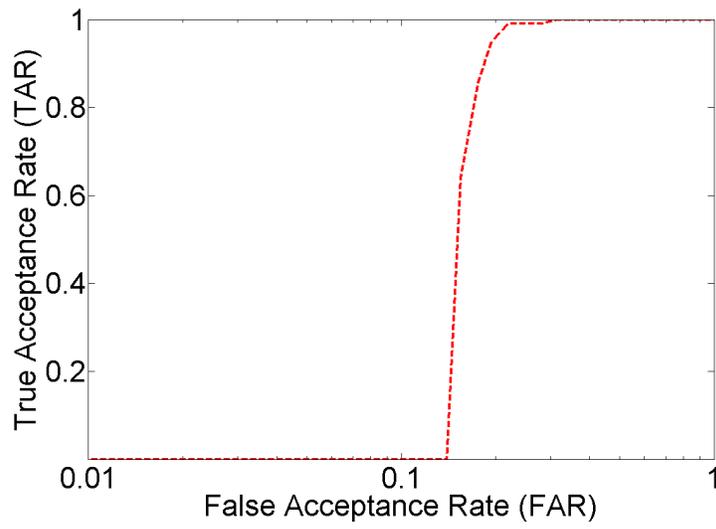


(a) CASIA Twins V1

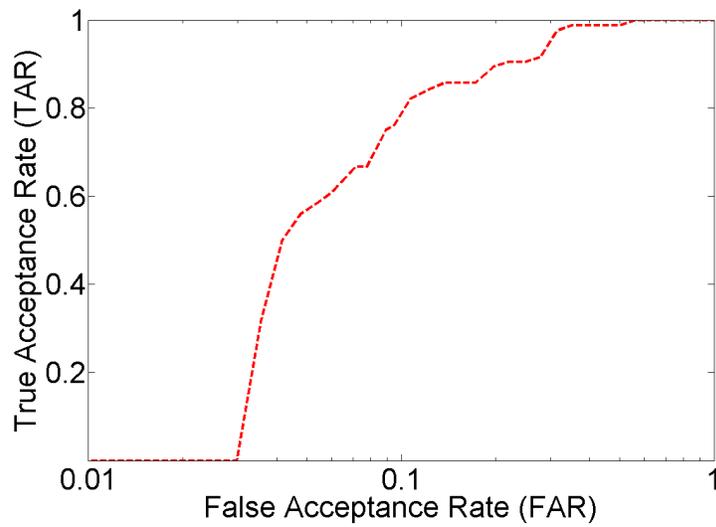


(b) CASIA Twins V2

Figure 2.12: ROC curves for twin detection using iris and face.



(a) CASIA Twins V1



(b) CASIA Twins V2

Figure 2.13: ROC curves for twin detection using fingerprint and face.



(a) Match score: 0.99954



(b) Match score: 0.99951

Figure 2.14: Face images of pairs of (a) identical and (b) non-identical twins with extremely high face match scores considering the range of scores is $[0, 1]$ using FaceVacs 8.6.)

2.5 Summary and Conclusions

We have presented a study of the distinctiveness of biometric characteristics of identical twins (fingerprint, iris and face). The discriminability of these three biometric traits is supported by anatomy and the formation process of the biometric characteristics, as discussed in Section 2.3. We have assessed the capacity of state-of-the-art commercial biometric matchers in distinguishing identical twins based on fingerprint, iris, and face.

The unimodal face biometric system can distinguish two different persons who are not identical twins much better than it can distinguish identical twins. More recent studies support this conclusion, and it becomes even more challenging when facial expression changes or different lighting conditions are involved during the image capture process. Some efforts have been made to improve the face recognition performance for identical twins, for example, by using facial marks. As described in Section 2.3, face undergoes change over time based on usual facial expressions, environment, diet, etc. At the same time, face recognition systems are evolving, so, in future, they are expected to have a better performance.

Although the unimodal fingerprint biometric system also can discriminate two different persons who are not identical twins better than it can discriminate identical twins, this difference is not as large as for the face biometric system. In the fingerprint experiments, the identical twin impostor distribution is shifted to the right, getting closer to the genuine distribution. This suggests a higher correlation between fingerprints of identical twins compared to fingerprints of unrelated persons. Previous studies have shown that the fingerprint type is much more likely to be the same in twins than in unrelated persons, and more recent studies confirm this.

The iris matching experiment results show that the performance of the biometric system for the identical twin data and for the general data are very similar, which means the iris biometric system can distinguish identical twins to the same extent as it can distinguish any two persons who are not identical twins. However, the shift in the identical twin impostor distribution for iris also suggests a higher level of similarity between irises from identical twins than from unrelated persons.

Among all the unimodal biometric systems considered here, the iris performed the best. Again,

this may be due to the fact that irises from identical twins are more uncorrelated than fingerprints from identical twins, or may be due to the fact that iris images in our dataset have a very good quality, while the quality of many fingerprint images is relatively poor and the friction ridge area is small due to the age of the subjects.

Multimodal biometric systems that combine different units of the same biometric modality (fusion of 4 fingers or 2 irises) lead to an almost perfect separation between genuine and impostor distributions. For both the general population and identical twins, multimodal biometric systems that combine different modalities, one being face, show improvements in the performance on the general population compared to individual traits. However, the performance of multibiometric systems drop on the identical twin data compared to the highest performance modality that is being combined. This is because in the fusion of finger and face matchers, the performance of the face matcher for identical twins is extremely poor.

Biometric traits can also be used to determine whether two subjects are identical twins. By using face and iris modalities together, for example, we can correctly determine 80% of the identical twin pairs as such, while only 2% of subject pairs who are not identical twins will be incorrectly considered identical twins.

To our knowledge, although multibiometric databases for twins have been collected before (the ten fingers and/or the two palmprints), no previous study has combined different units/modalities to study the multibiometric matching performance. We have performed the multibiometric experiments and showed that the performance of a multibiometric system that uses different units of the same modality is significantly better compared to unimodal systems, approaching almost perfect accuracy on our database. Also, there have not been any previous studies of identical twin irises on a database this large. In addition, we have used commercial matchers for the face and fingerprint experiments, which are usually more accurate than in-house matchers – used in previous studies. Based on our experiments on this relatively small twin database, we can conclude that the presence of identical twin data poses a real challenge to commercial face recognition systems.

CHAPTER 3

LATENT FINGERPRINT MATCHING

3.1 Introduction

Fingerprint identification was a completely manual approach until 1970s. Due to growing demands on fingerprint matching, and the large size of the fingerprint database of criminals, research was initiated to automate fingerprint recognition, which led to the development of *Automated Fingerprint Identification Systems* (AFIS). These systems are now used worldwide not only by law enforcement agencies but also in many other government applications, including background check of certain employees (e.g., those serving in the military), border control and national ID cards. The use of fingerprint recognition in civilian applications (e.g. logical/physical access control) is also gaining more widespread acceptance [84].

There are essentially three types of fingerprints encountered in law enforcement applications (see Fig. 3.1): rolled, plain and latents. Rolled prints are obtained by rolling the finger from nail-to-nail, while plain (or slap) fingerprints are obtained by placing the finger flat on a surface. It is common practice in law enforcement agencies to collect fingerprints from all the ten fingers of a subject. An example of a tenprint card is shown in Fig. 3.2, in which the first two rows show rolled prints of the ten fingers of a subject, while the last row shows plain fingerprints also from the ten fingers of a subject. Plain fingerprints are collected by first placing the four fingers of one hand, then the four fingers of the other hand, followed by the capture of the impressions of

the two thumbs (4-4-2 process). Latents are impressions of fingers lifted from surfaces of objects that are inadvertently touched or handled by a person typically at crime scenes. Lifting of latents may involve a complicated process, and it can range from simply photographing the print to more complex dusting or chemical processing [35].

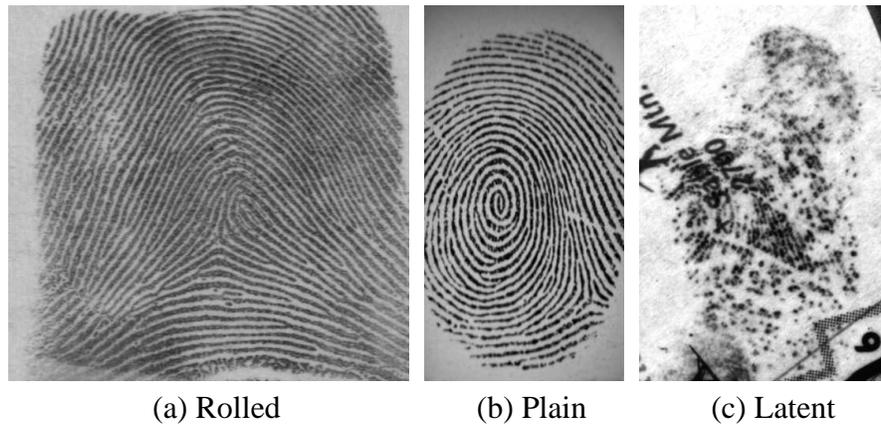


Figure 3.1: Three types of fingerprint impressions. Rolled and plain fingerprints are also called reference fingerprints.

Two types of matching are of interest to law enforcement agencies: tenprint-to-tenprint and latent-to-tenprint. Tenprint-to-tenprint matching is used in border control, background checks, etc., while latent-to-tenprint matching is used to identify suspects from impressions lifted from crime scenes.

Rolled prints contain the largest amount of information about the ridge structure on a fingerprint since they capture the largest finger surface area; latents usually contain the least amount of information for matching or identification because of their small size (they capture only a subset of the complete friction ridge pattern) and inherent noise. Compared to rolled or plain fingerprints, latents are smudgy and blurred and have large nonlinear distortion due to pressure variations. Due to their poor quality and small area, latents have a significantly smaller number of minutiae compared to rolled or plain prints (the average number of minutiae in NIST Special Database 27 (NIST SD27) [4] images is 21 for latents versus 106 for their mated rolled prints). These characteristics make the latent fingerprint matching problem very challenging.

APPLICANT		LEAVE BLANK <i>Leave Blank</i>		TYPE OR PRINT ALL INFORMATION IN BLACK LAST NAME FIRST NAME MIDDLE NAME Teacher, Theresa C.				FBI LEAVE BLANK <i>Leave Blank</i>	
SIGNATURE OF PERSON FINGERPRINTED		ALIAS AKA Formerly: Theresa Smith		ID NY9219402 NYSTED Dept-FPU ALBANY, NY				DATE OF BIRTH DOB 12/31/70	
RESIDENCE OF PERSON FINGERPRINTED 318 School Street Hometown, NY 11111		OFFICIAL US		SEX F		RACE W		HEIGHT 5'7"	
DATE 5/01/02		SIGNATURE OF OFFICIAL TAKING FINGERPRINTS		WEIGHT 155		HAIR Gr		EYES Bro	
EMPLOYER AND ADDRESS <i>(if applicable)</i> Smart Falls Central School Dist Smart Falls, NY 11111		VIEWER <i>Leave Blank</i>		LEAVE BLANK <i>Leave Blank</i>					
READY FINGERPRINTED <i>Leave Blank</i>		ARREST RECORD NO. <i>Leave Blank</i>		CLASS <i>Leave Blank</i>					
		SOCIAL SECURITY NO. 000-10-1111		REF <i>Leave Blank</i>					
		IDENTIFICATION NO. <i>Leave Blank</i>							

Figure 3.2: An example of a tenprint card containing two types of impressions (rolled and plain) of the ten fingers. The first portion of the card contains personal information related to the person being fingerprinted, such as name, date of birth, place of birth, gender, race, height, weight, etc. The second portion of the card contains impressions of the ten fingers. The first and second rows of impressions consist of rolled impressions of the fingers from the left and right hands, respectively (from left to right: thumb, index, middle, ring, little). The third row contains plain impressions of the left four fingers taken simultaneously, left thumb, right thumb, and right four fingers taken simultaneously.

Fingerprint examiners who perform manual latent fingerprint identification follow a procedure referred to as ACE-V (analysis, comparison, evaluation and verification) [85]. Because the ACE-V

procedure is quite tedious and time consuming for latent examiners, latents are usually matched against reference prints of a small number of suspects identified by other means, such as eye witness description or M.O. (mode of operation). With the availability of AFIS, fingerprint examiners are now able to match latents against a large fingerprint database using a semi-automatic procedure that consists of following stages: (i) manually mark the features (region of interest, minutiae and singular points) in the latent, (ii) launch an AFIS search, and (iii) visually verify the top- N (N is typically 50) candidate fingerprints returned by AFIS. The accuracy and throughput of this latent matching procedure is still not satisfactory. It certainly does not meet the “lights-out mode”¹ of operation desired by the FBI and which is the goal of the Next Generation Identification [86].

It is our opinion that research efforts in latent fingerprint identification should focus on improving the matching accuracy based on existing mark-ups by latent experts, rather than on completely eliminating manual input, or asking examiners for too much input. This opinion is supported by the following facts: (i) latent matching accuracy is still the major concern of law enforcement agencies, (ii) manually marking extended features is very labor extensive, and (iii) state of the art “lights-out” latent identification systems cannot yet offer satisfactory accuracy for most latents of casework quality. We believe latent matching will not reach the same level of performance of tenprint-to-tenprint matching due to some of the characteristics of the latents: (i) in tenprint acquisition, if the quality of the fingerprint image is not good, the fingerprint can be re-captured; in the latents case, this option is not available; (ii) for some latent prints, there might not be enough information to make an identification, for example, if the friction ridge area is extremely small and the ridges are not prominent or faded. Our goal of improving latent fingerprint matching performance aims to reduce the manual labor involved in identifying a latent. This is achieved by providing higher retrieval accuracy at top ranks, thus reducing the number of latent and rolled print pairs that need to be manually verified.

¹Lights-out identification refers to a system requiring minimal or zero human assistance in which an image is presented as input, and the output consists of a short candidate list (definition from http://biometrics.nist.gov/cs_links/latent/workshop09/proc/DefineLPlightsout.pdf).

Our efforts to improve latent fingerprint matching performance are concentrated on two different directions: (i) enhance the latent image and fuse the matching score of the enhanced latent with the score based on manually marked minutiae [87]; (ii) develop a new latent fingerprint matcher based on manually marked minutiae [88]. In the current practice, latent examiners are required to mark minutiae and optionally mark singular points (core/delta). In the enhancement approach, manually marked features in the latents include minutiae, region of interest, and singular points; in the latent matching approach, we reduce the manual input to only the manually marked minutiae.

Orientation field information is critical in the fingerprint enhancement process. Orientation field can usually be reliably estimated from the fingerprint image itself in the case of rolled or plain fingerprints of good quality. In latent fingerprint images, orientation field estimation is not very reliable because of their poor quality. Fig. 3.3 shows the estimated orientation field of a latent and its mated rolled print in NIST SD27 latent database using gradient-based approach [8]. Manually marked orientation field, although reliable, is not easy to obtain; it requires a lot of effort and prior training. Therefore, for the enhancement process, we reconstruct the orientation field based on manually marked features (minutiae, singular points, and region of interest) [89]. Gabor filters are then used to enhance latent images, which are automatically matched to the reference print database. We show that the performance of manually marked minutiae matching can be improved by combining scores from both the matching of manually marked minutiae and of automatically extracted minutiae from enhanced latents. It is important to point out that the performance of fully automated minutiae extraction and matching based on the input image is very poor. The experiments were conducted on a public domain fingerprint database, NIST SD27, consisting of 258 latents along with their rolled mates. To make the conclusion more reliable, the gallery size was increased by including 27,000 rolled images from NIST SD14 [90].

For fingerprint matching, there are two major problems which need to be solved. The first is to align the two fingerprints to be compared and the second is to compute a match score between the two fingerprints. Alignment between a latent and a rolled print is a challenging problem because latents often contain a small number of minutiae and undergo large skin distortion. To deal with these

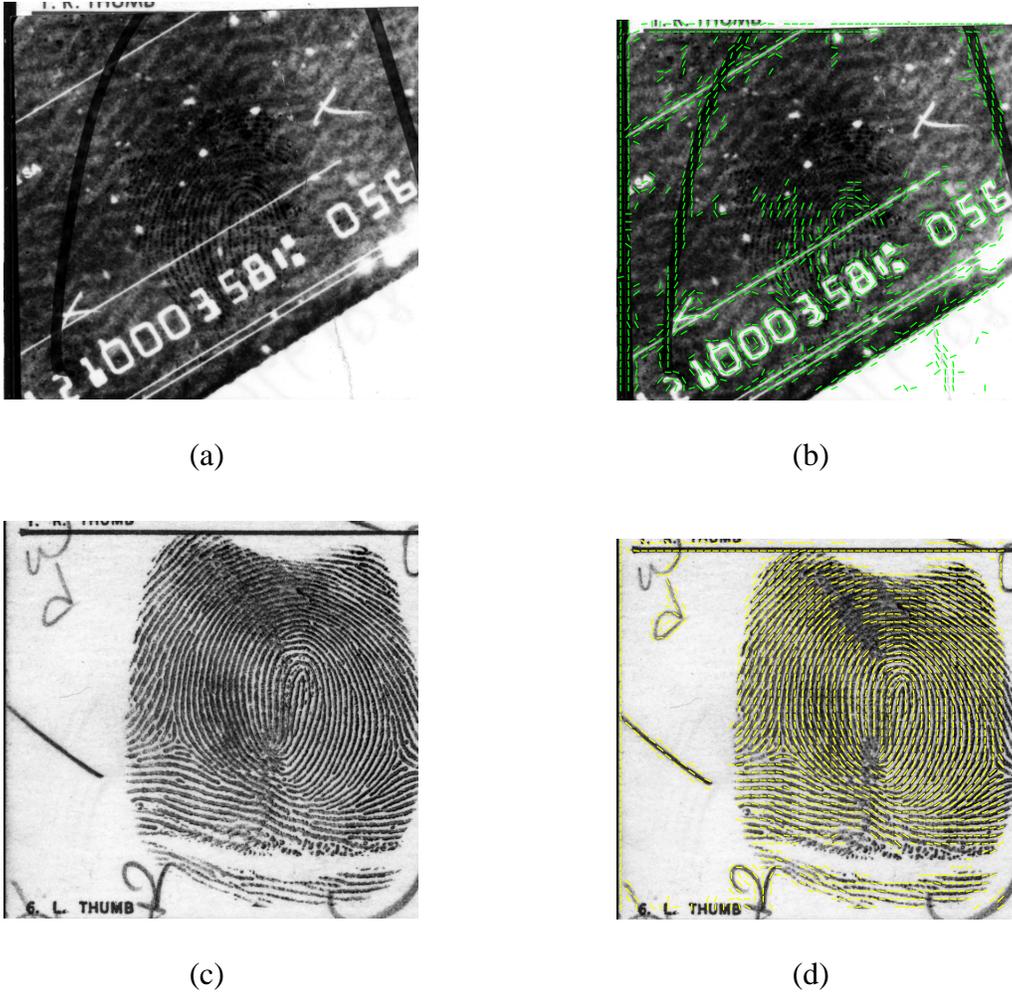


Figure 3.3: Examples of estimated orientation field ((b) and (d)) by gradient-based approach [8] for a (a) latent and (c) its mated rolled print in NIST SD27.

two problems, we propose the descriptor-based Hough transform (DBHT), which is a combination of the generalized Hough transform and a local minutiae descriptor, called the Minutia Cylinder Code (MCC) [91]. The MCC descriptor improves the distinctiveness of minutiae while the Hough transform method can accumulate evidence as well as improve the robustness against distortion. Match score computation between a latent and a rolled print is also challenging because the number of mated minutiae is usually small. To address this issue, we further consider orientation field as a factor in computing the match score. Since we only consider manually marked minutiae for latents, a reconstruction algorithm based on minutiae alone is used to estimate the orientation field [89]. The proposed matcher was tested on two latent fingerprint databases, NIST SD27 database

and West Virginia University latent fingerprint database (WVU LFD). Two COTS matchers and a state-of-the-art non-commercial fingerprint matching algorithm (MCC SDK) were evaluated on these two databases. Our algorithm was found to perform better than these three matchers on both the databases. Extensive experiments on fusion of matchers and effect of fingerprint quality were also conducted and reported here.

The rest of this chapter is organized as follows: in Section 3.2, related work is reviewed; in Section 3.3, the features used in the matching experiments are described; in Section 3.4, the two approaches for improving latent matching performance are presented; in Section 3.5, the databases used are described; in Section 3.6, experimental results based on the two proposed approaches are presented and discussed; in Section 3.7, our work on latent fingerprint matching is summarized.

3.2 Related Work

In this section, we review related work in four areas: (i) published research on rolled fingerprint matching², (ii) published research on latent fingerprint matching, (iii) NIST evaluation of latent fingerprint technologies (ELFT), and (iv) evaluation of latent examiners. Rolled fingerprint matching technology is very advanced, and there exist commercial matchers with excellent matching performance; however, there is very little information available about the details of these algorithms and, therefore, we do not discuss them here.

3.2.1 Rolled Fingerprint Matching

The majority of the algorithms developed for fingerprint matching are based on minutiae. Although minutiae carry a great amount of discriminatory information, additional features may be necessary to achieve the desired level of accuracy. Most proposed algorithms for fingerprint matching that use non-minutiae features also use minutiae. For example, some algorithms combine ridge orientation with minutiae information either at feature level by including ridge orientation information in local

²See Chapter 4 in [1] for a more comprehensive review of this topic.

minutiae descriptors [27, 92] or at score level by combining scores from minutiae matching and global orientation field matching [92, 93].

Several recent studies on fingerprint matching have focused on the use of local minutiae descriptors [91, 27, 92, 94, 95, 96, 97]. In most of these studies, the initial step consists of using local minutiae descriptors to obtain the alignment between two fingerprints by considering the most similar minutiae pair; then, a global consolidation step is performed to obtain a better matching performance. Since these algorithms are usually tuned and evaluated using FVC databases (plain fingerprints) or NIST Special Database 4 (rolled fingerprints), their performances on latent fingerprints are unknown.

3.2.2 Latent Fingerprint Matching

Recent research and development efforts on latent fingerprints can be classified into three streams according to the manual input required from fingerprint examiners: consistent with existing practice, increasing the amount of manual input, or reducing the amount of manual input. Because of large variations in latent fingerprint quality and specific requirements of practical applications (crime scenes, border crossing points, battle fields), each of the three streams has its value.

Improved latent matching accuracy has been reported by using extended features, which are manually marked for latents [28, 98, 99, 11]. However, marking extended features (orientation field, ridge skeleton, etc.) in poor quality latents is very time-consuming and might be only feasible in rare cases when there are none or very few minutiae. Thus, some studies have concentrated on latent matching using a reduced amount of manual input, such as manually marked region of interest (ROI) and singular points [100, 101], or no manual input [102]. However, the reported performance of these approaches is not very good. Hence our proposed matcher takes manually marked minutiae as input and, therefore, it is consistent with existing practice in forensics. There have also been some studies on fusion of multiple matchers [103] and use of multiple latent prints [104].

3.2.3 NIST Evaluation of Latent Fingerprint Technologies

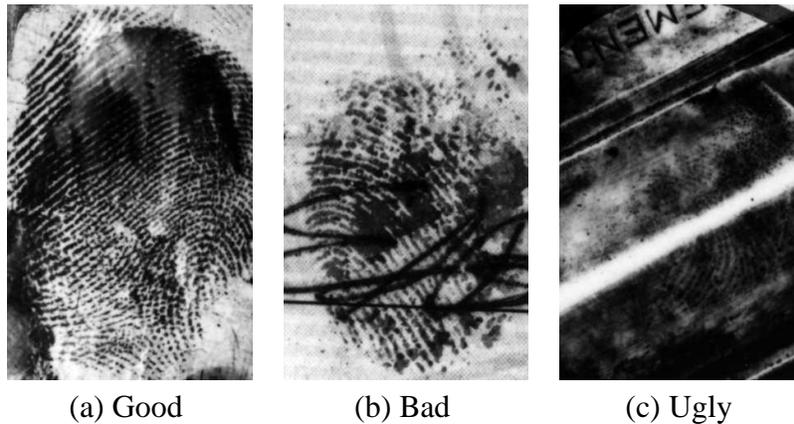


Figure 3.4: Latent fingerprints of three different quality levels in NIST SD27.

NIST has been conducting a multi-phase project on Evaluation of Latent Fingerprint Technologies (ELFT) to evaluate latent feature extraction and matching techniques [105]. Since all participating algorithms in ELFT are proprietary, no detailed information about these algorithms is available. The purpose of ELFT-Phase I was to assess the feasibility of latent fingerprint identification systems using Automated Feature Extraction and Matching (AFEM), while the purpose of ELFT-Phase II was to actually measure the performance of state-of-the-art AFEM technology and evaluate whether it was viable to have those systems in operational use to reduce the amount of time needed by latent examiners to manually mark latents thereby increasing the throughput.

In Phase I, latent images were selected from both operational and non-operational scenarios. The average reported accuracy at rank-1 was of 67% (100 latents matched with 10,000 rolled prints) [106]. In Phase II, latent images were selected from only operational environments. The rank-1 accuracy of the most accurate system was 97.2% (835 latents matched with 100,000 rolled prints) [50]. The Phase I and Phase II accuracies cannot be directly compared since the Phase I and Phase II evaluations used different latent databases. Furthermore, the quality of latents used in Phase II is much better compared to Phase I. As shown in Fig. 3.4, the quality of operational latents can vary significantly.

The impressive matching accuracy reported in ELFT, specially Phase II, does not support that

the current practice of manually marking minutiae in latents should be changed. Although latents in Phase II were selected from operational scenarios, they represent successful identifications in actual case examinations using existing AFIS technology. In the ACE-V process, when the examiner analyzes the latent image he decides whether the latent has (i) value for exclusion only, (ii) value for individualization or (iii) no value. If a latent is classified as of “no value”, no comparison is performed. If the latent is classified in one of the other two categories, then a comparison is carried out and the examiners can make a decision on individualization and exclusion, or determine the comparison to be inconclusive. So the latents which are successfully identified constitute only a small part of all latents, which are of reasonable quality. For this reason, in the ELFT-Phase II report [50] the authors concluded that only a limited class of latents can benefit from the AFEM technology.

NIST has conducted another evaluation of latent fingerprint technologies using extended feature sets (EFS) manually marked by latent examiners, called ELFT-EFS [107]. In this evaluation, the purpose was to investigate the matching accuracy when (i) latent images and/or (ii) sets of manually marked latent features were provided to the matcher. This evaluation suggested that the highest accuracy was obtained when the input included both the latent image and manually marked latent features.

3.2.4 Evaluation of Latent Examiners

A latent examiner can be viewed as a very accurate “matcher”. But, there is the issue of subjectivity in the human decision making process along with low throughput. Because of the time consuming process of manual decision making, quantitatively estimating the accuracy of latent examiners is not easy. Hence the numbers of fingerprint pairs used in several “black box” studies of latent examiners are not large [108, 109, 110]. Although the exact numbers reported in these studies may not reflect the real practice, the qualitative conclusions are very useful. It was found that latent examiners’ conclusions are not always in agreement, especially in the case of poor quality latents [108]. In addition, the same examiner can change his/her conclusions on the same fingerprint pair

at a later time [109]. These inconsistencies may even increase under bias [110].

These issues associated with including latent examiners in the latent identification process will only be resolved when the automatic matcher can outperform latent examiners in accuracy. No matter how successful the application of automatic fingerprint recognition technology might be, we cannot say fingerprint matching is a “solved problem” before we can reach the goal of outperforming latent examiners.

3.3 Latent Fingerprint Features

In both of our approaches to improve latent matching performance, minutiae and orientation field are extracted from latent as well as rolled prints. Minutiae are marked by latent examiners in the latent, and automatically extracted using commercial matchers in the rolled print. Orientation field in the latents is reconstructed from minutiae location and direction, as proposed in [89], and orientation field is automatically extracted from the rolled print images by using a gradient-based method [8]. In the matching using descriptor-based Hough Transform, local minutia descriptors are used. In this section, the orientation field reconstruction algorithm is described, followed by the description of the local minutia descriptors used in this work.

3.3.1 Orientation Field Estimation

The value of an orientation field at a given pixel is the angle that the fingerprint ridges form with the horizontal axis in a small neighborhood around that pixel (see Fig. 3.5). The simplest and the most natural approach for computing orientation field is based on the gradient values. Because of the computational effort and the presence of noise, the orientation field is usually computed in a small neighborhood instead of at each pixel [1].

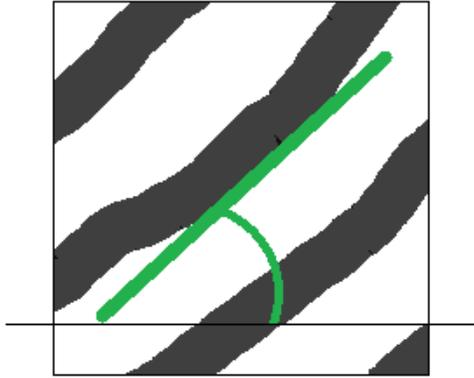


Figure 3.5: Illustration of the orientation field computed at a small block.

Another approach to compute the orientation field that differs from this image-based approach is the model-based approach (e.g. zero-pole model [111, 112, 113]), which uses the location of singular points to estimate the orientation field based on a model.

In the latent fingerprint case, since the images are usually of poor quality, it is very difficult to estimate the orientation field based only on the image itself. Also, in many cases, the friction ridge area of the latent fingerprint is small and does not contain singular points, which makes the model-based orientation field unreliable. A combination of image-based and model-based approaches can be very useful for latents. In [100], such a combination was used to enhance the latent fingerprint images and it was shown to improve the matching performance compared to the performance based on latent image alone.

In [89], the authors proposed a method to reconstruct the orientation field assuming (i) only the minutiae information (location and direction) is given, and (ii) both minutiae and singular point information is given. The advantages of this approach over the approaches described earlier [111, 112, 113] are: it can be used in cases where no singular points are available (small area latents) and it uses minutiae information that is reliable. Our experiments show that when this reconstructed orientation field is used for latent enhancement, the matching performance of the enhanced image is comparable to the matching performance of images enhanced by manually marked orientation field. This reconstructed orientation field is also useful for orientation field

matching. Therefore, we adopt this approach which is described in detail below.

3.3.1.1 Minutiae based Orientation Field Reconstruction

In this section, we describe the method proposed in [89] to reconstruct the orientation field based on minutiae and singular points, if available. Here we assume that a latent expert has marked the available features (minutiae, singular points, region of interest). Also, in cases where an odd number of singular points were marked in the region of interest, paired singular points outside of the region of interest are guessed (referred to as virtual singular points [100]) to satisfy the constraints of singularity number, which states the total number of singular points is even, and the cores and deltas appear in pairs.

Because of the reasons mentioned in Section 3.3.1, the orientation field is also computed in non-overlapping blocks of a predefined size (e.g., 8×8 or 16×16). Now, consider lines passing through the non-overlapping blocks, that divides the image into 8 equally spaced sectors as shown in Fig. 3.6. Then, a local orientation field estimate is obtained, for each block, based on the direction of the nearest minutia in each of the 8 sectors.

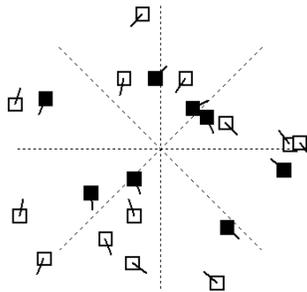


Figure 3.6: Local ridge orientation estimation based on the nearest minutiae in each sector.

Let $\{x_n, y_n, \alpha_n\}$, $1 \leq n \leq N$, be a set of N fingerprint minutiae marked by a latent expert, where (x_n, y_n) is the location and α_n is the direction of the n^{th} minutia. Then, by doubling the minutia direction α_n , which means taking $2\alpha_n$ instead of α_n as the minutia direction, it becomes equivalent to $\alpha_n + \pi$ (this is necessary since fingerprint ridges are not oriented). For the K minutiae selected in eight sectors (in our experiments, one minutia per sector), cosine and sine components can be

computed and summed as follows:

$$u = \sum_{k=1}^K \cos(2\alpha_k)w_k, \quad (3.1)$$

$$v = \sum_{k=1}^K \sin(2\alpha_k)w_k, \quad (3.2)$$

where w_k is a weighting function based on the Euclidean distance between the block center and the k^{th} minutia; this makes the closest minutia direction dominates the ridge orientation of neighboring blocks.

The orientation at block (m, n) is then computed as:

$$D(m, n) = \frac{1}{2} \arctan \frac{v}{u}. \quad (3.3)$$

If we consider that marked singular points are also available in the latent fingerprints, then the direction field of N_s singular points is given by the Zero-Pole model [111], [112]:

$$D_s(m, n) = \sum_{i=1}^{N_s} t_{s_i} \arctan \frac{n - n_{s_i}}{m - m_{s_i}}, \quad (3.4)$$

where m_{s_i} , n_{s_i} , and t_{s_i} (core: 1, delta: -1) are the location and type of the i^{th} singular point.

Orientation field is first estimated using minutiae whose direction is subtracted by D_s . The reconstructed orientation field is then given by

$$O(m, n) = \frac{2D(m, n) + D_s(m, n)}{2}. \quad (3.5)$$

The orientation field is reconstructed simply as $D(m, n)$ (Eq. 3.3) in the enhancement case in which there are no marked singular points, and in the matching using Descriptor-Based Hough Transform in which we consider only minutiae are marked. The orientation field is only reconstructed in the region of interest (ROI). In the first approach (enhancement) we assume region of interest is manually marked. In the second approach, we estimate the region of interest as the convex hull of minutiae.

Figs. 3.7 and 3.8 show some examples of reconstructed orientation field. In both figures, (a) shows the original latent image, (b) shows the orientation field estimated directly from the

gray scale image, (c) shows the reconstructed orientation field and (d) shows manually marked orientation field. We can observe that it is not easy to estimate the orientation field from the latent image. But, the reconstructed orientation field from minutiae and singular points is quite reliable, although it is not as smooth as manually marked orientation field.

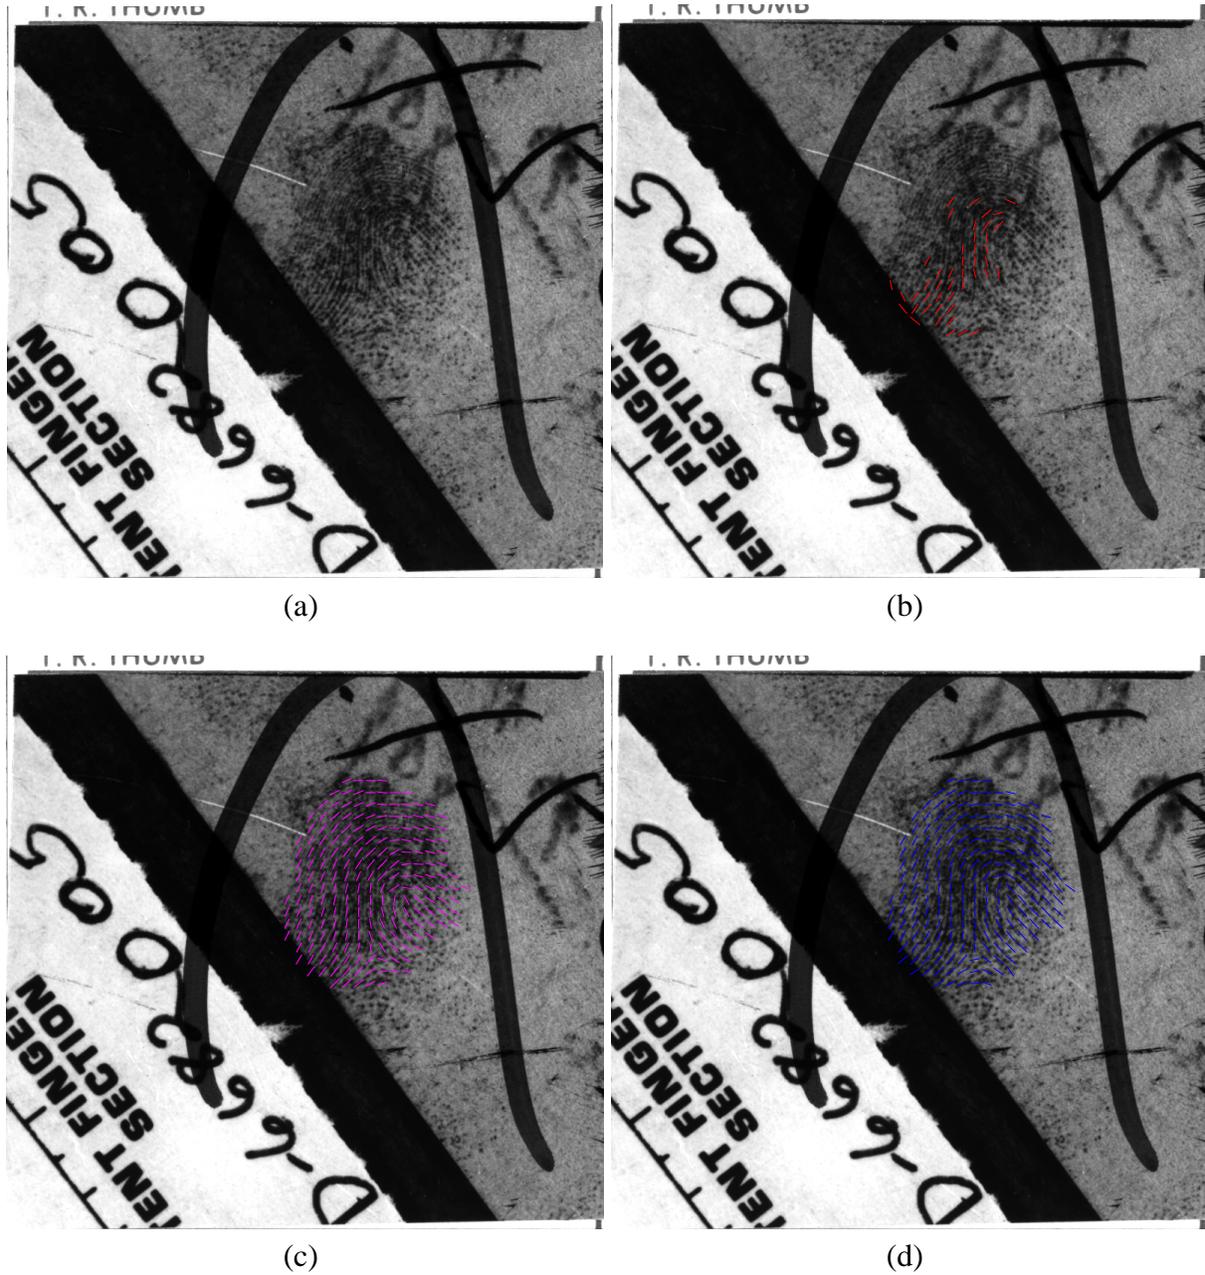


Figure 3.7: Comparison of orientation field estimation methods: (a) Original latent fingerprint image, (b) orientation field estimated directly from the given image, (c) orientation field estimated from minutiae and singular points, and (d) manually marked orientation field.

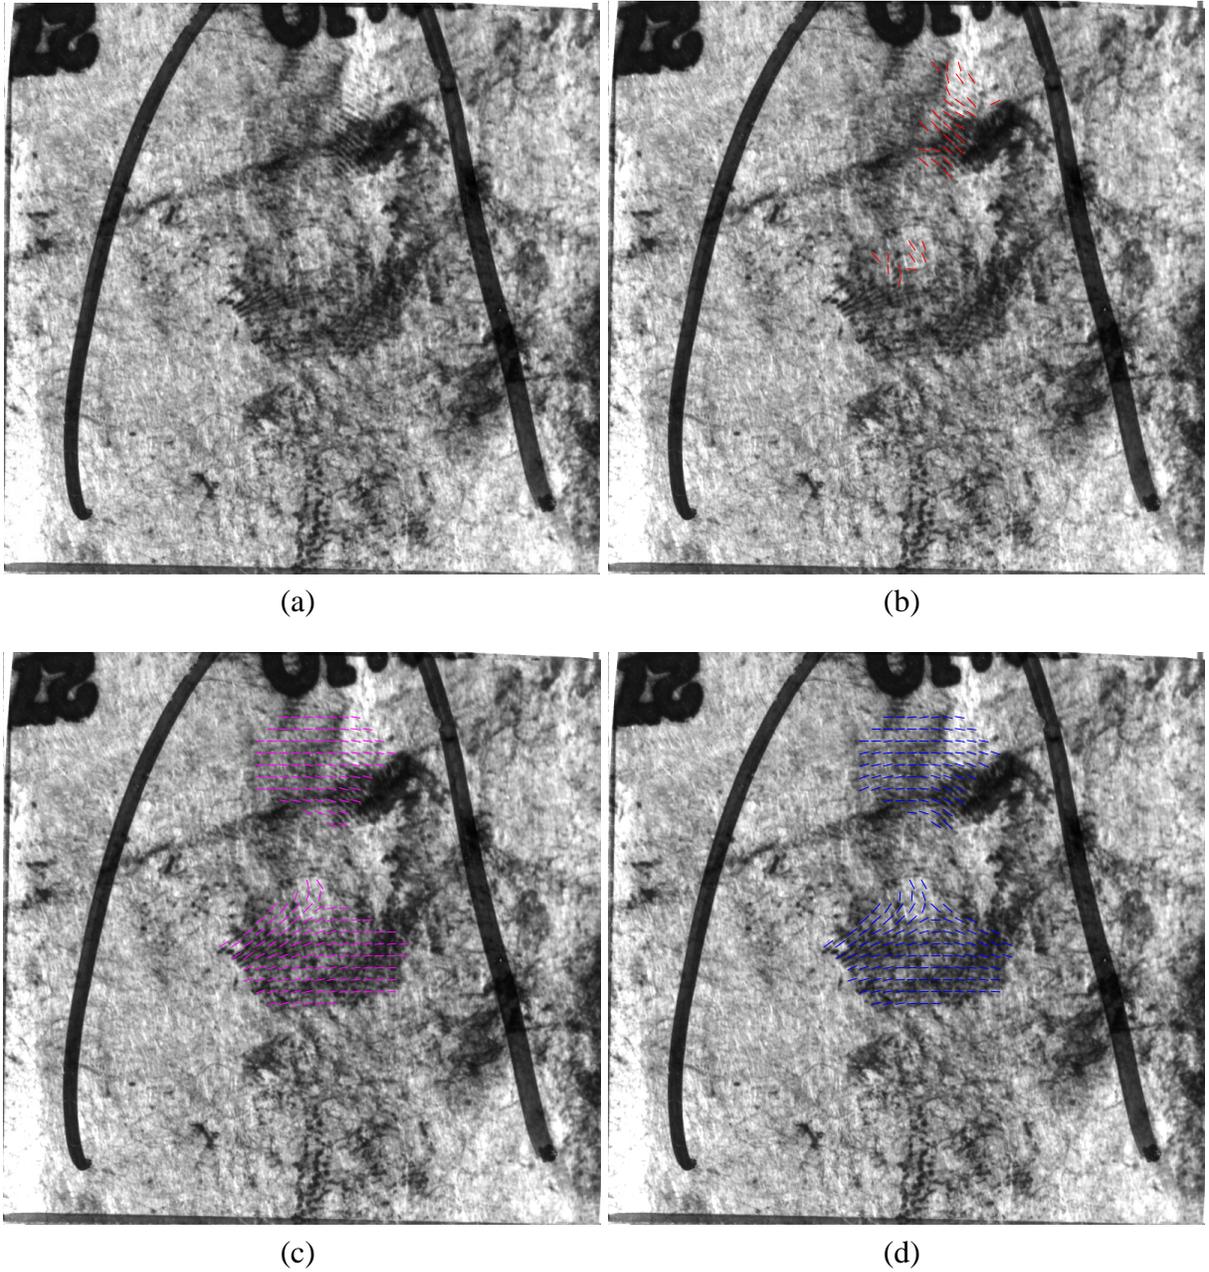


Figure 3.8: Comparison of orientation field estimation methods: (a) Original latent fingerprint image, (b) orientation field estimated directly from the given image, (c) orientation field estimated from minutiae and singular points, and (d) manually marked orientation field.

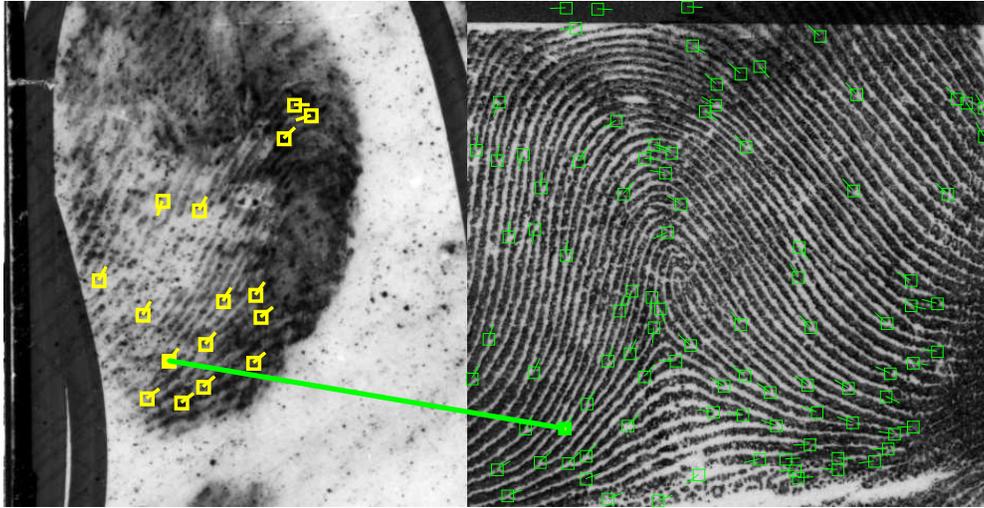
3.3.2 Local Minutia Descriptor

Local descriptors have been widely used in fingerprint matching (e.g. [94, 27, 95, 11, 91]). Feng and Zhou [114] evaluated the performance of local descriptors associated with fingerprint matching in four categories of fingerprints: good quality, poor quality, small common region, and large plastic distortion. They also coarsely classified the local descriptors as image-based, texture-based, and minutiae-based descriptors. Their results show that the minutiae-based descriptor, Minutia Cylinder Code (MCC) [91], performs better in three of the four categories (good quality, poor quality and large plastic distortion), and texture-based descriptor performs better for the small common region category.

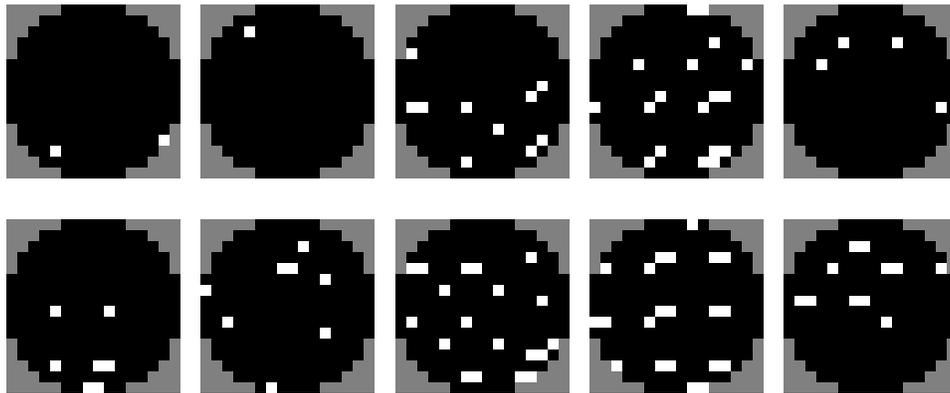
3.3.2.1 Minutia Cylinder Code (MCC)

A minutia cylinder records the neighborhood information of a minutia as a 3D function. A cylinder contains several layers and each layer represents the density of neighboring minutiae along the corresponding direction. The cylinder can be concatenated as a vector, and therefore the similarity between two minutia cylinders can be efficiently computed. Fig. 3.9(b) shows the sections of two valid cylinders associated with the two corresponding minutiae (in the latent and in the rolled print) indicated in Fig. 3.9(a). A more detailed description of the cylinder generation and of the similarity between two cylinders can be found in [91].

A cylinder is divided into sections (height), and each section is divided into cells (base). Each cell in the cylinder receives contribution from two different sources: the first is related to the distance of neighboring minutiae to the center of the cell, and the second is related to the directional difference between the neighboring minutiae and the center minutiae (the main minutiae for which the cylinder is being built). Sections correspond to directional differences, and a cell receives a contribution only from those neighboring minutiae that have a directional difference to the center minutia similar to the directional difference corresponding to the section where the cell is located. A cell is only valid if it is inside the cylinder and if it is inside the convex hull of minutiae (with some tolerance for the last condition). Smoothing functions (e.g. sigmoid functions) are used



(a) Latent and corresponding rolled print with a mated minutiae pair indicated.



(b) Sections of the cylinder corresponding to the minutia indicated in the latent (first row) and in the rolled print (second row).

Figure 3.9: Sections of two cylinders associated with the two corresponding minutiae, one in latent and other in rolled print.

to compute directional differences between minutiae and the center minutia, and to compute the distance among neighboring minutiae and cell center points. The contributions of all neighboring minutiae to a cell are combined under another smoothing function to limit the contribution of noisy regions, which usually contain a large number of minutiae. This contribution is compared to a threshold so that the final cylinder is a binary representation (bit-based implementation).

Only valid cylinders and cells are used in the similarity computation. If c_a and c_b are two cylinders represented as vectors, \hat{c}_{ab} indicates which cells are valid in both cylinders, $c_{a|b} = c_a \text{ AND } \hat{c}_{ab}$, and $c_{b|a} = c_b \text{ AND } \hat{c}_{ab}$, then the similarity between the two cylinders in the bit-based implementa-

tion case is given by

$$s(c_a, c_b) = \begin{cases} 1 - \frac{\|c_a|b \text{ XOR } c_b|a\|}{\|c_a|b\| + \|c_b|a\|} & \text{if } c_a \text{ and } c_b \text{ are matchable} \\ 0 & \text{otherwise,} \end{cases} \quad (3.6)$$

where c_a and c_b are matchable if (i) the directional difference between the two minutiae is less than some predefined threshold (maximum rotation allowed), (ii) there is a minimum number of corresponding elements in the two cylinder (minimum size of the valid intersection), and (iii) the denominator is different from zero.

3.4 Latent Matching Approaches

In this section, we present our two proposed approaches to improve latent fingerprint matching performance. The first approach accomplishes that by using enhancement of the latent images, while the second approach improves the performance by using a descriptor-based Hough Transform alignment in the matching algorithm.

3.4.1 Latent Matching with Enhanced Image

The feature extraction process in latent fingerprints usually does not yield reliable features because of the poor quality of the latents. One way of improving the latent feature extraction process is by enhancing the latent image so that more reliable features can be extracted and a better matching performance can be achieved. In [115], the authors proposed to enhance fingerprint images using Gabor filters to improve the clarity of ridges and valleys. This enhancement is performed in local blocks of the image, and, for each block, an orientation is required. In this approach, orientation field reconstructed from minutiae and singular points (if available) as described in Section 3.3.1 is used to enhance the latent image. Then, a commercial matcher is used to extract minutiae and perform the matching. The scores from manually marked minutiae and from the automatically extracted minutiae from the enhanced images are then combined to improve the latent matching

performance. The overall scheme of the latent matching using enhanced image is illustrated in Fig. 3.10.

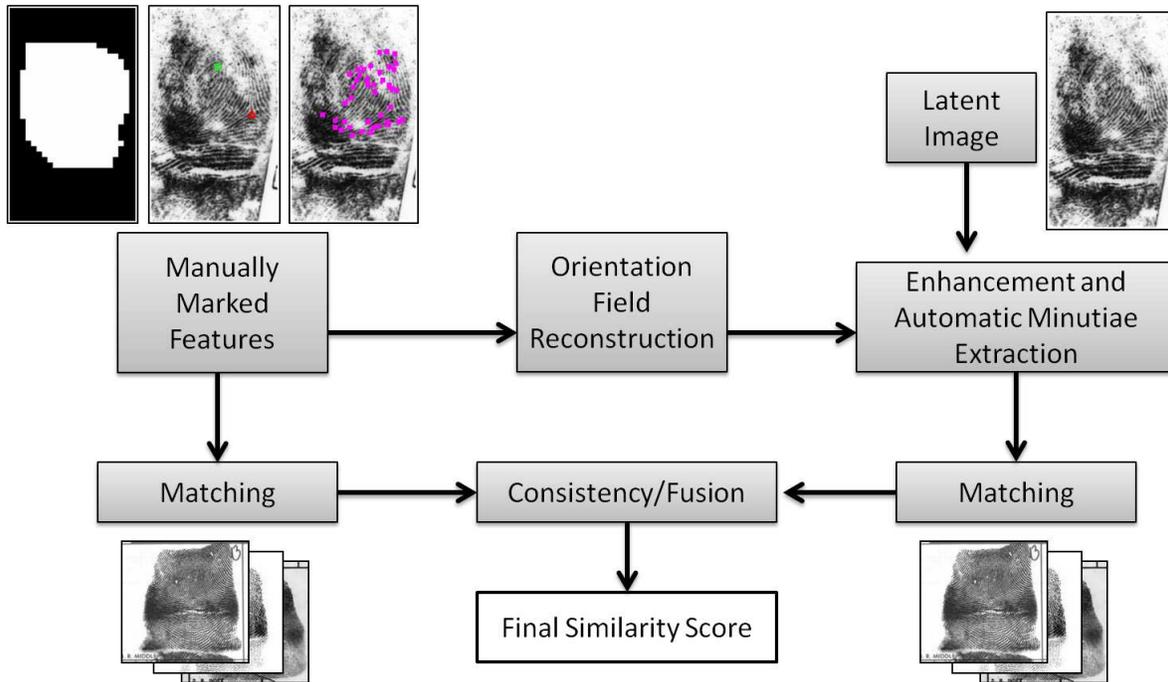


Figure 3.10: Overview of the latent matching with enhanced image approach.

Fig. 3.11 shows some examples of enhanced latent images, along with their skeletons, the corresponding enhanced images and their enhanced image skeletons. It can be noted from the figure that the clarity of the ridges improves and the noise is greatly reduced.

Another approach that is often used to improve the matching accuracy consists of combining different sources of information. Information fusion can be implemented at a variety of levels, such as feature level, score level, and rank level. Feature level fusion in our context would involve both manually marked minutiae and automatically extracted minutiae for matching. In our preliminary experiments, feature level fusion did not show promising results, so we focus on the other two fusion levels, rank and score.

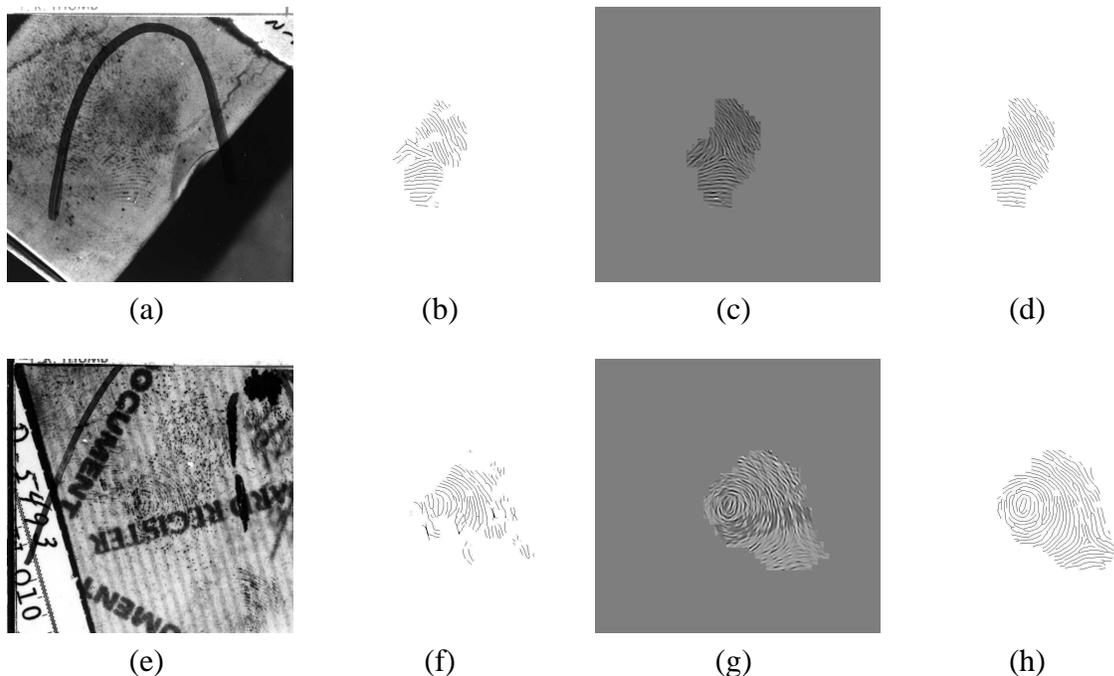


Figure 3.11: Examples of enhanced latent images. (a) and (e) are original latent images, (b) and (f) are the skeletons of the original latent images extracted by VeriFinger 4.2 SDK, (c) and (g) are enhanced images using reconstructed orientation field, and (d) and (h) are the skeletons of the enhanced images extracted by VeriFinger.

3.4.1.1 Rank-level Fusion

For rank-level fusion, we considered two techniques: Borda count [116] and the highest rank method.

Borda count is a generalization of the majority vote. In our case where we have two different matchers (classifiers), one matching manually marked latent to rolled and the other matching enhanced latent image to rolled, the Borda count for a given rolled print in the database will be the sum of the number of rolled prints that are ranked below the true mate by each matcher. Then, the ranking is performed by sorting the rolled prints in descending order based on their Borda count. The magnitude of the Borda count for each rolled print in the database measures how much the two matchers agree on whether the input comes from the true mate of the latent.

In the highest rank approach, each rolled print in the database is assigned the highest rank as computed by the two matchers [117]. If there are any ties, the rolled prints are sorted by their lower

rank. If both higher and lower ranks are the same, the tie is broken randomly.

Our experiments show that while Borda count does not improve the matching performance of manually marked minutiae, the highest rank fusion shows some improvement. These results will be discussed in Section 3.6.1.1.

3.4.1.2 Score-level Fusion

Min, max, product and weighted sum rules are some of the well-known score-level fusion rules. Let s_1 and s_2 be the two match scores obtained for the same (latent, rolled) fingerprint pair by using two different matchers, called matcher 1 and matcher 2. Then, for the match score pair s_1 and s_2 we can compute fused scores based on the score level fusion rules mentioned above as follows:

$$\begin{aligned}
 S_{min} &= \min(s_1, s_2) \\
 S_{max} &= \max(s_1, s_2) \\
 S_{prod} &= s_1 s_2 \\
 S_{wsum} &= w_1 s_1 + w_2 s_2,
 \end{aligned} \tag{3.7}$$

where $w_1 + w_2 = 1$.

In addition to the above score fusion rules, we also considered combining the two match scores by applying a slight modification of the boosted max fusion approach proposed in [98]. The main purpose of boosted max is to boost the scores of genuine matches. It relies on the assumption that the spatial transformations among the given latent, enhanced latent and rolled fingerprints are consistent for genuine matches and inconsistent for impostor matches.

In general, given a transformation matrix T and a two dimensional vector $v = (x, y)$, the coordinates for v , $v' = (x', y')$, after the transformation is applied is given by $v' = Tv$:

$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta & \Delta x \\ -\sin \theta & \cos \theta & \Delta y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}, \tag{3.8}$$

where θ is the rotation and Δx and Δy are translation parameters in x and y , respectively.

The parameters θ , Δx , and Δy in the transformation matrix can be estimated by using pairs of corresponding minutiae points. Note here that scaling parameter is not considered due to the fact that the fingerprints are all in the same resolution (500 ppi).

Let T_{MR} , T_{ER} , and T_{ME} be the estimated transformations between manually marked minutiae and rolled fingerprints, enhanced latent images to rolled fingerprints, and manually marked minutiae to enhanced latent images, respectively. Given T_{ER} and T_{ME} , we can compute $T_{MER} = T_{ER} \times T_{ME}$, which denotes the spatial transformation of the manually marked latent to the automatically extracted minutiae in the enhanced latent, which is then matched to the rolled fingerprint image. It is expected that the transformations T_{MR} and T_{MER} should be similar for genuine pairs and different for impostor pairs. This would be likely the case since the two minutiae sets (manually marked and extracted from enhanced image) are different and are expected to match to different parts of the impostor rolled prints, while the two sets are expected to be correctly matched to their true corresponding minutiae in the true mate rolled print. The similarity between the two transformation matrices is measured in terms of the rotation angle and the Euclidean distance between the translations in x and y directions of each transformation. If the difference between the two rotation angles in the two matrices T_{MR} and T_{MER} is less than some threshold and the Euclidean distance between the two translations in those same matrices is less than another threshold, then the pair is considered consistent. In our study, the rotation and translation thresholds were set to 30 degrees and 140 pixels, respectively. The boosted max score for a given pair of match scores s_1 and s_2 is

$$S_b = \begin{cases} w_1 \max(s_1, s_2) + w_2 \min(s_1, s_2), & \text{if consistent} \\ \max(s_1, s_2), & \text{otherwise.} \end{cases} \quad (3.9)$$

In Section 3.6 we discuss the results obtained by using enhanced images and different fusion schemes.

3.4.2 Latent Matching with Descriptor-Based Hough Transform

Another way of improving matching performance is to design a matcher especially for latent fingerprint matching, which takes into account the special characteristics of latents such as small

number of minutiae. Our goal is to develop an algorithm for latent fingerprint matching that uses as few manually marked features as possible.

Given a latent fingerprint (with manually marked minutiae) and a rolled fingerprint, we extract additional features (besides minutiae coordinates) from both prints, align them in the same coordinate system, and compute a match score between them. These three steps are described in the following subsections. An overview of the proposed algorithm is shown in Fig. 3.12.

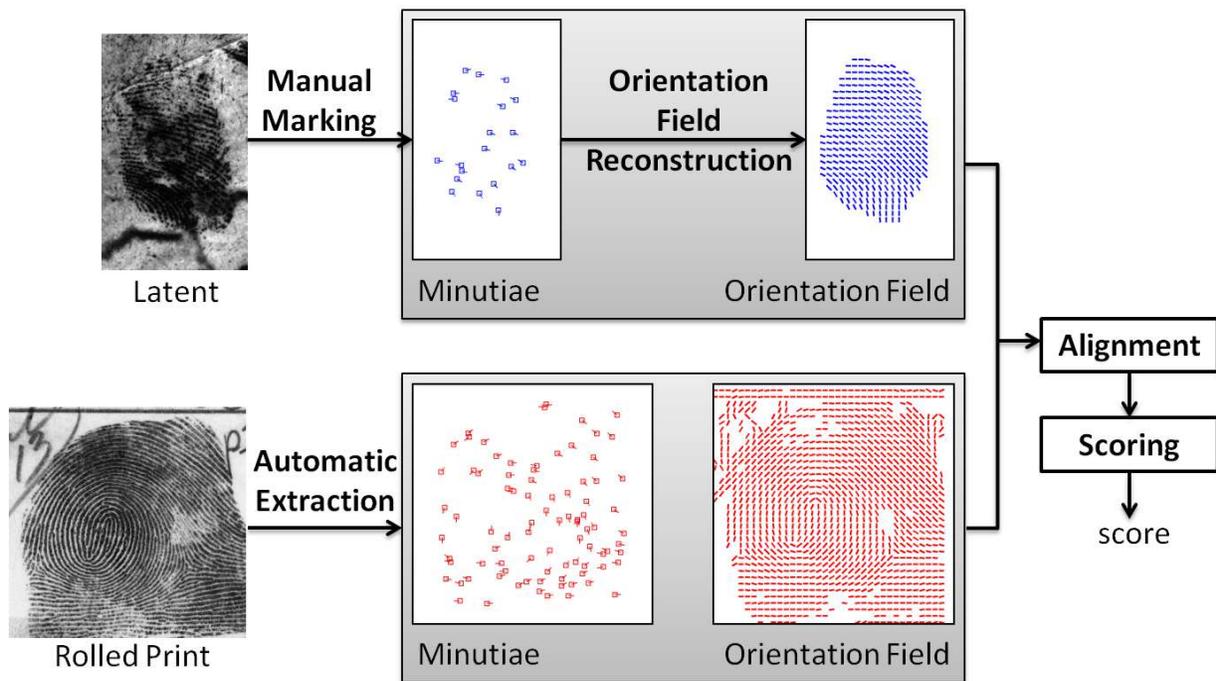


Figure 3.12: Overview of the proposed latent matching approach.

3.4.2.1 Alignment

Fingerprint alignment or registration consists of estimating the parameters (rotation and translation) that align two fingerprints. A scale parameter can be ignored here because all the images have the same resolution. There are a number of features that may be used to estimate alignment parameters between two fingerprints, including singular points, orientation field, ridges, and minutiae. There are also a number of methods to align two fingerprints: Generalized Hough Transform

(e.g., [29]), local descriptors (e.g., [91]), energy minimization, etc³.

In the latent matching case, singular points are not always present in latents, making it difficult to base the alignment of the fingerprint on singular points alone. As state earlier, obtaining manually marked orientation field is expensive, and to automatically extract reliable orientation field from a latent image is difficult. Since manually marking the minutiae in a latent is a common practice for latent matching, our approach to align two fingerprints is based on minutiae.

Local descriptors can also be used to align two fingerprints. In this case, usually the most similar minutiae pair is used as a base for estimating the initial transformation parameters (rotation and translation); the most similar minutiae pair is chosen based on a measure of similarity between the local descriptors of the minutiae pair.

Ratha *et al.* introduced an alignment method for minutiae matching that estimates rotation, scale, and translation parameters using the Generalized Hough Transform [29]. Given two sets of points (minutiae), a matching score is computed for each transformation in the discretized set of all allowed rigid transformations. For each pair of minutiae, one minutia from each image (latent or reference), and for given scale and rotation parameters, unique translation parameters can be computed. Each parameter receives “a vote” that is proportional to the matching score for the corresponding transformation. The transformation that gives the maximum score is considered to be the best one. In our approach, the alignment is conducted in a similar way, but the evidence for each transformation parameter is accumulated based on the similarity between the local descriptors of the two minutiae being matched, with the similarity and descriptor being the ones described in Section 3.3.2.1.

The descriptor-based Hough transform alignment algorithm takes as input two sets of minutiae, M_L and M_R , and two sets of local descriptors C_L and C_R , one set corresponding to the latent and one to the rolled print. Each set contains a local descriptor for each minutia. A high level algorithm of the proposed approach to align two fingerprints given the sets of minutiae and local descriptors is shown in Algorithm 1, and an illustrative scheme is shown in Fig. 3.13.

³Refer to Chapter 4 of [1] for details and published work.

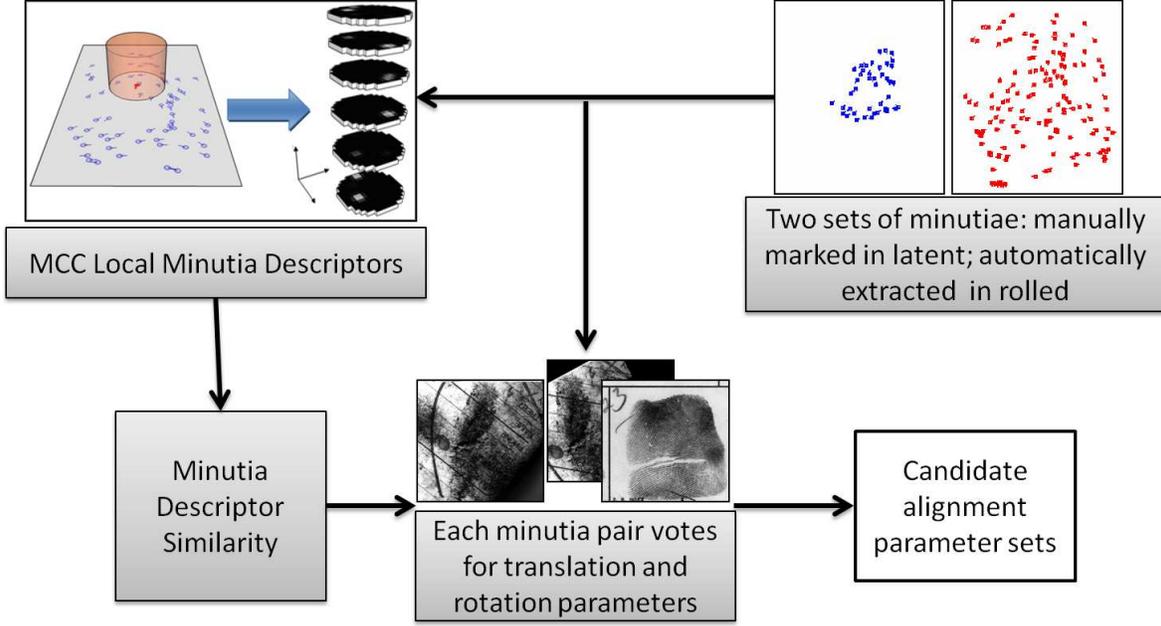


Figure 3.13: Illustration of the MCC descriptor-based Hough Transform alignment.⁴

Given two sets of minutiae, one from the latent and the other from the rolled print being compared, translation and rotation parameters can be obtained for each possible minutiae pair (one minutia from each set). Let $\{(x_l, y_l, \theta_l)\}$ and $\{(x_r, y_r, \theta_r)\}$ be the minutiae sets for latent and rolled prints, respectively, centered at their means. Then, for each pair of minutiae, we have

$$\Delta\theta = \begin{cases} \theta_l - \theta_r, & \text{if } \|\theta_l - \theta_r\| \leq 180 \\ \theta_l - \theta_r - 360, & \text{if } \theta_l - \theta_r > 180 \\ \theta_l - \theta_r + 360, & \text{if } \theta_l - \theta_r < -180 \end{cases} \quad (3.11)$$

$$\begin{pmatrix} \Delta x \\ \Delta y \end{pmatrix} = \begin{pmatrix} x_r \\ y_r \end{pmatrix} - \begin{pmatrix} \cos\Delta\theta & \sin\Delta\theta \\ -\sin\Delta\theta & \cos\Delta\theta \end{pmatrix} \begin{pmatrix} x_l \\ y_l \end{pmatrix}, \quad (3.12)$$

where Δx and Δy are translation parameters in x and y , respectively, and $\Delta\theta$ is the rotation parameter.

Since the scale (resolution) is fixed in fingerprint matching (all images have the same resolution), unique translation parameters can be obtained for each pair based on the rotation difference

⁴MCC Local descriptors figure is extracted from <http://biolab.csr.unibo.it/research.asp?organize=Activities&select=&selObj=81&pathSubj=111%7C%7C8%7C%7C81&Req=&>.

Algorithm 1 Descriptor-based Hough Transform.

Input: $\{m_l\} = \{(x_l, y_l, \theta_l)\} \in M_L$, $\{m_r\} = \{(x_r, y_r, \theta_r)\} \in M_R$, C_L , and C_R

Output: A set of 10 rigid transformation matrices

Initialize the accumulator array A

Compute local minutiae descriptor similarity (W) for every possible minutiae pair using C_L and C_R

for all possible minutiae pairs m_l, m_r **do**

 Compute their direction difference $\Delta_\theta = (\theta_r - \theta_l)$

if $\Delta_\theta < \max_\theta$ **then**

 Compute translation parameters (Δ_x, Δ_y) and increase the voting for this set of alignment parameters:

$$A(\Delta_x, \Delta_y, \Delta_\theta) = A(\Delta_x, \Delta_y, \Delta_\theta) + W(l, r) \quad (3.10)$$

end if

end for

Smooth A using a Gaussian low-pass filter

Find the 10 highest peaks in A

for each peak k **do**

 Compute a rigid transformation between two fingerprints using minutiae pairs that contributed to peak k and its immediate neighborhood

if the estimated rigid transformation is not reliable (not enough non-collinear points) **then**

 Repeat the voting in peak k and its neighborhood using a refined range

 Find the highest peak in the small neighborhood of peak k

end if

end for

between the two minutiae in the pair. The translation and rotation parameters need to be quantized to the closest bins. After the quantization, evidence is accumulated in the corresponding bin based on the similarity between the local minutiae descriptors. The assumption here is that true mated minutiae pairs will vote for very similar sets of alignment parameters, while non-mated minutiae pairs will vote randomly throughout the parameter space. As a result, the set of parameters that presents the highest evidence is considered the best one. For robustness, ten sets of alignment parameters with strong evidence are considered.

In order to make the alignment computationally efficient and also more accurate, we use a two-step approach to compute the alignment parameters for a fingerprint pair. The first step is to perform the voting using the Descriptor-based Hough Transform. If the bins are too small, the true peak in the Hough Transform space is not likely to receive sufficient votes. On the other hand, if

the bins are too large, they will not provide accurate alignment parameters because of the possible non-linear deformation. The strategy we adopted is to keep the bins relatively large (10 pixels), and to include a second step to compute reliable alignment parameters. This second step consists of using the minutiae pairs that vote for a peak to compute a rigid transformation between the latent and rolled fingerprints. The use of voting minutiae pairs to compute the transformation gives more accurate alignment parameters than directly using the peak parameters. In cases where a rigid transformation matrix cannot be reliably obtained, the voting is repeated inside a neighborhood of the corresponding peak, but with a smaller bin size (2 pixels). A peak is chosen from this refined Hough Transform space, and used as the alignment parameters.

3.4.2.2 Similarity Measure

For each of the 10 different alignments, a matching score between the latent and the rolled fingerprints is computed by comparing minutiae and orientation fields. The maximum value of the 10 scores corresponding to the 10 different alignments is chosen as the final matching score between the latent and rolled fingerprints. The details for computing matching scores of minutiae and orientation field are given below, and the final score computation scheme is shown in Fig. 3.14.

To compute minutiae matching score under a given alignment, we first find the corresponding minutiae pairs (one in the latent, one in the rolled print). For this purpose, we align the minutiae sets of the two fingerprints and then find an one-to-one matching⁵ between the two minutiae sets using a greedy algorithm. For each minutia m_l in the latent, a set of candidate minutiae in the rolled print is found. A minutia m_r in the rolled print is called a candidate if it has not yet been matched to any minutia, and both its location and angle are sufficiently close to m_l . The threshold values T_S for spatial distance and T_A for angle distance were determined empirically to measure the proximity between two minutiae. Among all candidates, the one closest to m_l in location is chosen as the matching minutia of m_l .

⁵One-to-one matching means that each minutia in the latent is matched to at most one minutia in the rolled print, and vice versa.

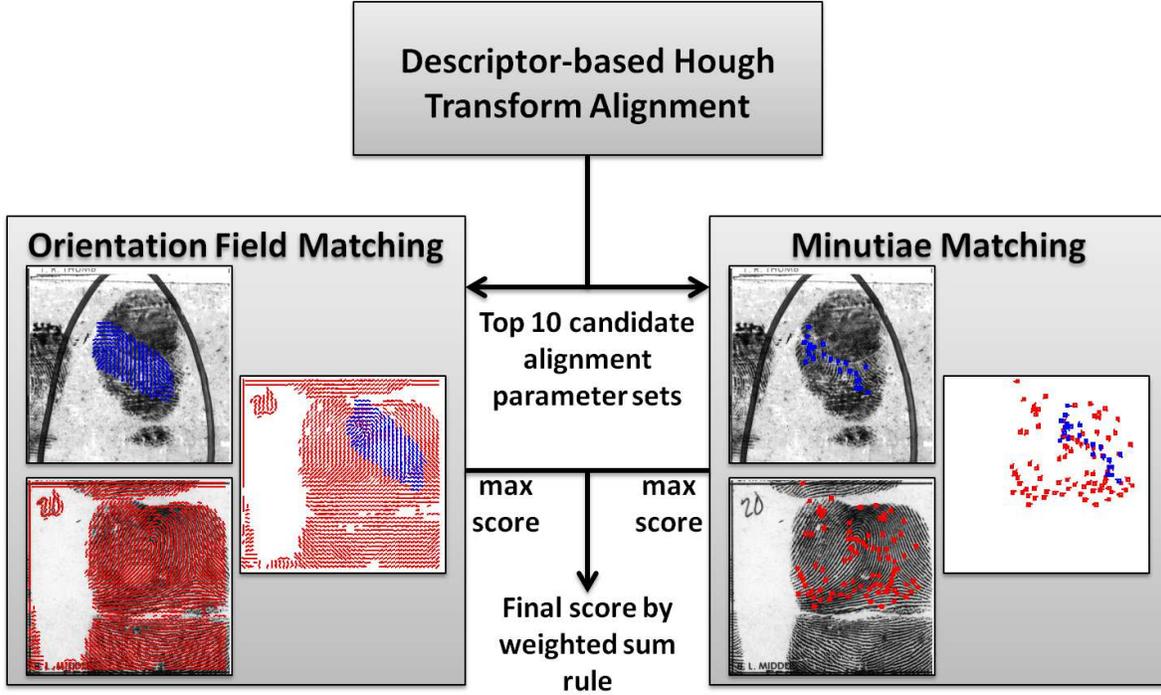


Figure 3.14: Final score computation scheme combining minutiae and orientation field information.

After the corresponding minutiae are found, we compute a matching score between the latent and the rolled print. Suppose that n pairs of matching minutiae between the latent and the rolled print are found. The minutiae matching score S_M between the two fingerprints is given by

$$S_M = \frac{1}{N} \sum_{i=1}^n s_C(i) s_S(i), \quad (3.13)$$

where $s_C(i)$ denotes the similarity between the minutia cylinder codes of the i th pair of matched minutiae, $s_S(i) = 1 - d_S(i)/2T_S$ maps the spatial distance $d_S(i)$ of the i th pair of matched minutiae into a similarity score, and N denotes the number of minutiae in the latent.

According to equation (3.13), the matching score depends on the number of matching minutiae, which itself is affected by the distance threshold T_S . However, due to large distortion present in many latents, it is difficult to choose an appropriate value for T_S . While a large threshold value will lead to more matching minutiae for distorted mated pairs, the number of matching minutiae for non-mated pairs will increase as well. Hence, we use two different values (15 pixels and 25 pixels) for T_S and for each threshold, a set of matching minutiae is found and a matching score

is computed using equation (3.13) at a fixed T_A of 20° . The mean of the two scores is used as the minutiae matching score. Fig. 3.15 shows an example in which the score of the genuine pair is slightly reduced when the smaller threshold is used compared to the larger threshold, while the score of the latent and the rank-1 non-mate⁶ using the larger threshold is greatly reduced when the smaller threshold is used.

We use a simple orientation field matcher that basically measures the consistency of the orientation field differences between the latent and rolled prints. If we use the Euclidean distance, for example, to measure the orientation field differences, a small error in the rotation alignment between the latent and the rolled will contribute a small amount to the orientation field difference for every block being compared, resulting in a large overall difference or small similarity score. In [33], the authors proposed a distance measure for orientation field matching that can handle small rotation alignment errors. Given the aligned latent orientation field O_L and the rolled orientation field O_R , each containing K blocks, namely $O_L(k)$ and $O_R(k)$, the similarity between the two orientation fields is given by

$$S_O = \frac{|\sum_{k=1}^K v_k e^{j2(O_L(k) - O_R(k))}|}{\sum_{k=1}^K v_k}, \quad (3.14)$$

where v_k is 1 if both corresponding blocks k are valid, and 0 otherwise.

The overall matching score between the latent and rolled prints based on minutiae matching and orientation field matching is given by

$$S = (1 - w_0)S_M + w_0S_O, \quad (3.15)$$

where the weight w_0 is empirically set as 0.4. Fig. 3.16 shows an example in which the fusion of minutiae matching and orientation field matching scores helps improve the retrieval rank⁷ of the true mate due to the higher orientation field matching score between the latent and the true mate orientation fields, compared to the score between the latent and the non-mate orientation fields.

⁶The rank-1 non-mate refers to the non-mated rolled print whose match score with the latent ranks first among all the rolled prints in the database.

⁷Retrieval rank of a rolled fingerprint refers to its rank in the whole candidate list which is sorted in the decreasing order of matching scores with the latent.

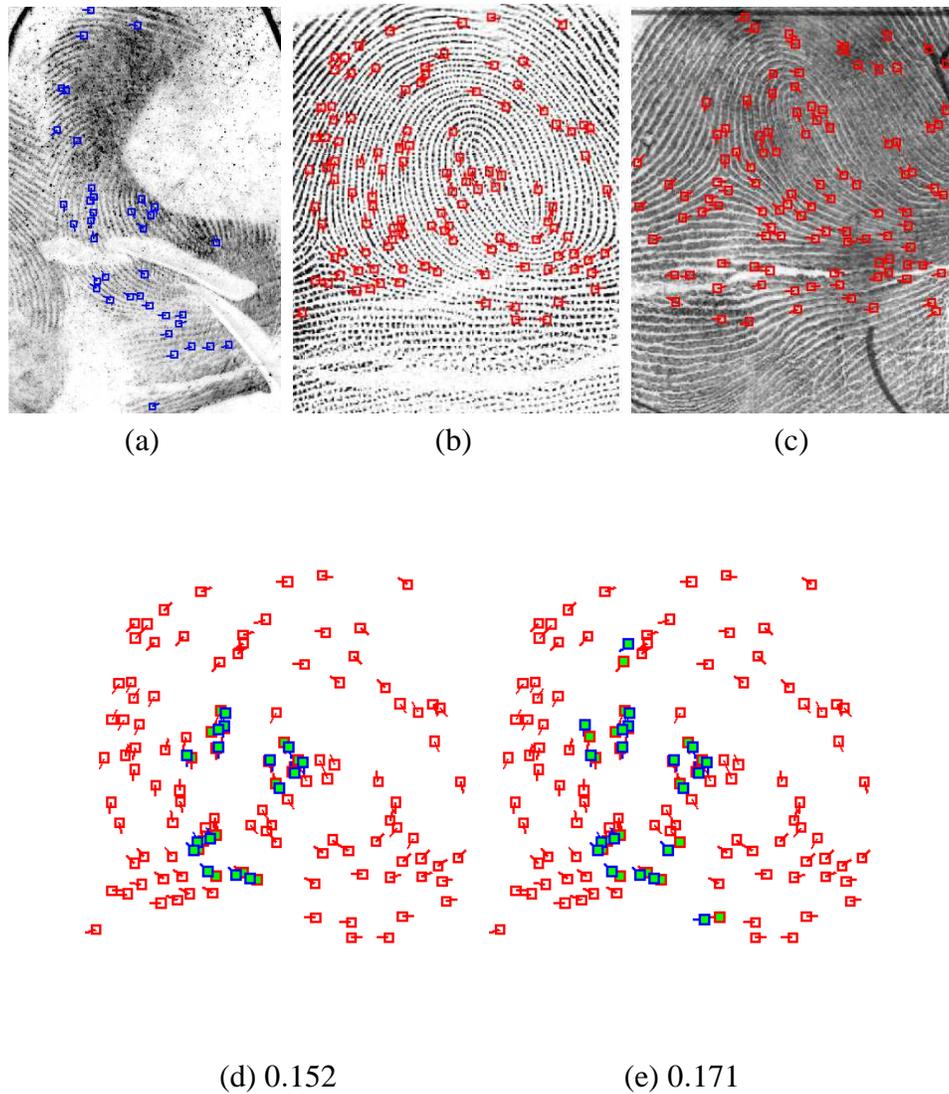
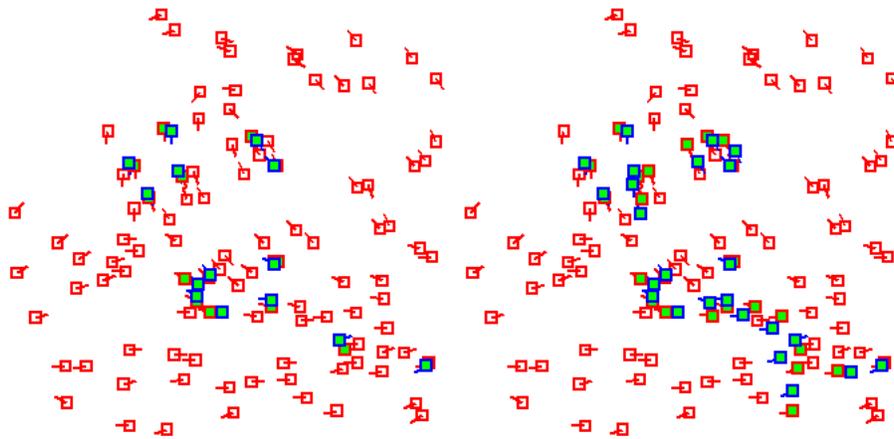


Figure 3.15: Latent print for which the matching score of the genuine pair is slightly reduced when a smaller threshold value is used compared to a larger threshold value, while impostor score is greatly reduced. (a)-(c) show the latent, the true mate, and the rank-1 non-mate according to large threshold, respectively. (d)-(g) show latent minutiae that were matched to rolled print minutiae in the following cases: (d) true mate using small threshold, (e) true mate using large threshold, (f) non-mate using small threshold, and (g) non-mate using large threshold. In (d)-(g), the scores corresponding to each case are included. Filled-in minutiae indicate the matching minutiae.

The retrieval rank of the true mate improved from 2 to 1 after the fusion, while the retrieval rank of the rank-1 non-mate according to minutiae matcher was changed from 1 to 3 after the fusion.

Fig. 3.15 (cont'd)



(f) 0.113

(g) 0.172

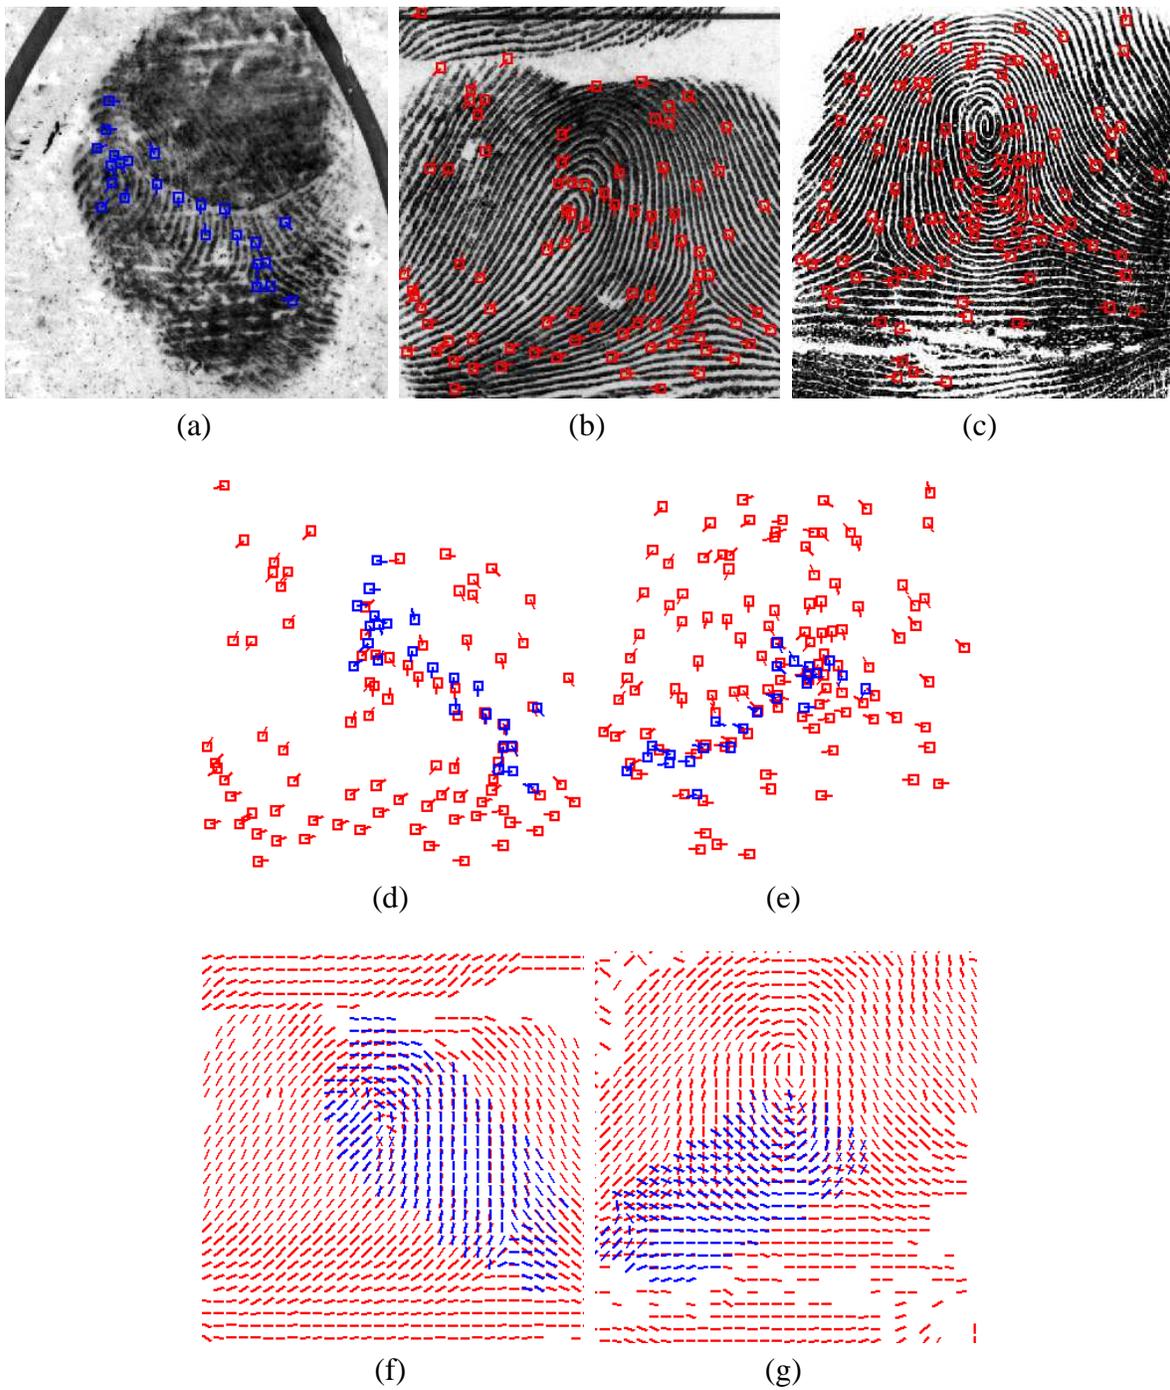


Figure 3.16: Latent print identified at a higher rank after fusing minutiae matching scores with orientation field matching scores. The rank of the true mate was improved from 2 to 1 after the fusion, and the rank of the highest ranked non-mate was 3 after the fusion. (a)-(c) show minutiae and the image of (a) a latent, (b) its true mate, and (c) the highest ranked non-mate according to minutiae matching. (d) and (f) show latent minutiae and orientation field (in blue) aligned with minutiae and orientation field of the true mate. (e) and (g) show latent minutiae and orientation field (in blue) aligned with minutiae and orientation field of the rank-1 non-mate.

3.5 Databases

Matching experiments were conducted on two different latent fingerprint databases: NIST Special Database 27 (NIST SD27) and West Virginia University Latent Fingerprint Database (WVU LFD). In this section, we present some characteristics of these two latent databases.

3.5.1 NIST Special Database 27 (NIST SD27)

NIST Special Database 27 is the only publicly available database comprising latent fingerprints from operational scenarios (latents collected at crime scenes). It consists of 258 latent fingerprint images and 258 corresponding (mated) rolled prints. Both latents and rolled prints are available at 500 and 1000 ppi – we used 500 ppi for consistency with the other databases. The quality of the latents in NIST SD27 varies, reflecting the operational (casework) quality.

NIST SD27 contains latent prints of three different qualities, termed “good”, “bad”, and “ugly”, which were classified by latent examiners. Some examples of latents from those three qualities are shown in Fig. 3.4. Although this classification of latent prints as “good”, “bad”, and “ugly” is subjective, it has been shown that such a classification is correlated with the latent matching performance [28].

Another indicator of fingerprint quality that affects the matching performance is the number of minutiae in the latent print [28]. In other words, while the latent may have clear and sharp edges, it may have relatively few minutiae available for matching. Based on the number of minutiae n in latents in NIST SD27, Jain and Feng [28] classified latents in NIST SD27 into three groups: large number of minutiae ($n \geq 22$), medium number of minutiae ($13 < n \leq 21$), and small number of minutiae ($n \leq 13$), containing 86, 85, and 87 prints, respectively. We present our experimental results for each of the six quality groups based on subjective quality (good, bad and ugly) and the number of minutiae.

We use manually marked minutiae – provided with NIST SD27 – as features in latent fingerprints. For rolled fingerprint images, the minutiae are automatically extracted using the two

commercial matchers.

3.5.2 West Virginia University Latent Database (WVU LFD)

West Virginia University Latent Database⁸ consists of 449 latent fingerprint images collected in a laboratory environment and 4,740 rolled prints, including the 449 mated rolled prints of the 449 latents. The latent images in this database are at 1000 ppi, and they were converted to 500 ppi for our experiments so that all the databases used in our experiments had the same resolution. Fig. 3.17 shows a latent with its corresponding rolled print in the WVU latent database. Similar to NIST SD27 database, manually marked minutiae were also provided with these latents. Minutiae were automatically extracted from the rolled prints using the two commercial matchers.



Figure 3.17: A latent and its corresponding rolled print in the WVU latent database. The NFIQ quality [9] of the rolled print is 4, which is one above the worst NFIQ quality.

There is no subjective quality value assigned to the latents in the WVU database. As mentioned earlier, one way to assign an objective quality measure to a latent is based on the number of minutiae in the latent, so any latent can be assigned an objective quality. If we apply the same objective quality classification scheme as in NIST SD27 to WVU database, we obtain 208, 80, and 161 latent fingerprints in the objective qualities of large, medium, and small number of minutiae, respectively.

⁸To request WVU latent fingerprint database, please contact Dr. Arun A. Ross (<http://www.csee.wvu.edu/~ross/>) at Integrated Pattern Recognition and Biometrics Lab (<http://www.csee.wvu.edu/~ross/i-probe/>).

The two latent databases used in our experiments, NIST SD27 and WVU, have different characteristics: most of the latent images in NIST SD27 contain significant background noise, while in WVU latent images, collected in a laboratory environment, there is a uniform background in most latents. However, overall, based on NFIQ quality [9], the quality of the rolled prints in WVU database is worse than the quality of rolled prints in NIST SD27. Fig. 3.18 shows the histograms of NFIQ values [9], one of the very well recognized measure of fingerprint quality, of the rolled prints which have corresponding latents in NIST SD27 and in WVU databases (258 and 449 rolled prints, respectively). NFIQ defines five quality levels in the range [1, 5] with 1 indicating the highest quality. The overall worse rolled print quality of WVU database compared to NIST SD 27 could be explained because in the operational database such as NIST SD27, rolled prints were captured by experienced law enforcement officers which, to our understanding, is not the case for the WVU database. Further, if the rolled prints corresponding to the latents are of poor quality, the number of mated minutiae will be small and, therefore, it would be much more challenging to identify the true mates of the latents at rank-1.

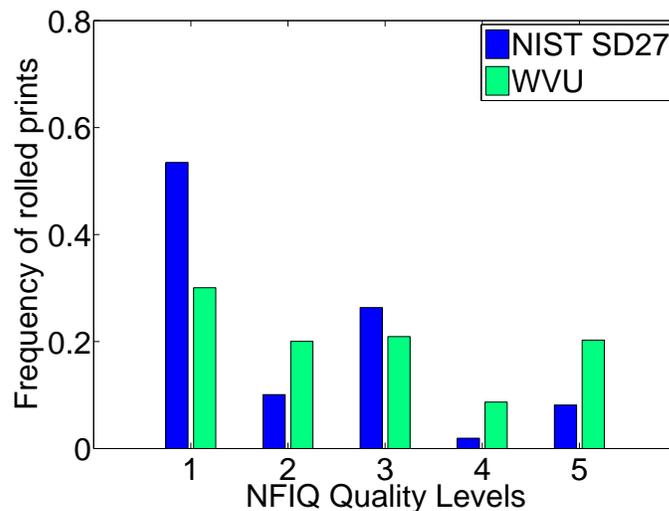


Figure 3.18: Histograms of NFIQ values of rolled prints in NIST SD27 and WVU databases.

3.6 Experimental Results

3.6.1 Latent Matching with Enhanced Images

Matching experiments were conducted on the NIST Special Database 27, which consists of 258 latent fingerprint images and 258 mated rolled fingerprint images. The manually marked features in the latents in this database are region of interest (ROI), minutiae, visible singular points (inside the ROI) and “virtual” singular points (outside the ROI). To make the matching problem more challenging and realistic, the background database (gallery) was increased from 258 mated rolled fingerprints to 27,258 total rolled fingerprints by adding 27,000 fingerprint images from the database NIST SD14 [90]. For the rolled fingerprint images, only minutiae were needed for matching and they were automatically extracted using Verifinger [79]. For boosted max fusion, the transformations were computed based on the matched minutiae output by Verifinger for each pair of fingerprints being matched.

The latent images from NIST SD27 were enhanced using the approach described in [115]. In addition, since latent images are usually of poor quality, ridge frequency values were estimated for each image block [115], and the median of those values was chosen as the ridge frequency for the entire image. We also evaluated the use of a constant ridge frequency value for all the images, and different ridge frequency values in each block of the image. The median ridge frequency value provided the best performance and therefore the results shown here are based on it.

The matching performance of enhanced latent images using reconstructed orientation field against rolled fingerprints is much better than the matching performance using minutiae automatically extracted from the original latent image. While the matching performance of enhanced images is worse than the performance of manually marked minutiae, this performance is comparable to the matching performance of enhanced images using manually marked orientation field, which requires significantly more manual labor. Fig. 3.19 shows the Cumulative Matching Characteristic curves on NIST SD 27 for manually marked minutiae, enhanced images using reconstructed orientation field, automatically extracted minutiae from latent image, and enhanced images using

manually marked orientation field. Minutiae in the latents were either manually marked (MMM) or automatically extracted by Verifinger 4.2 [79].

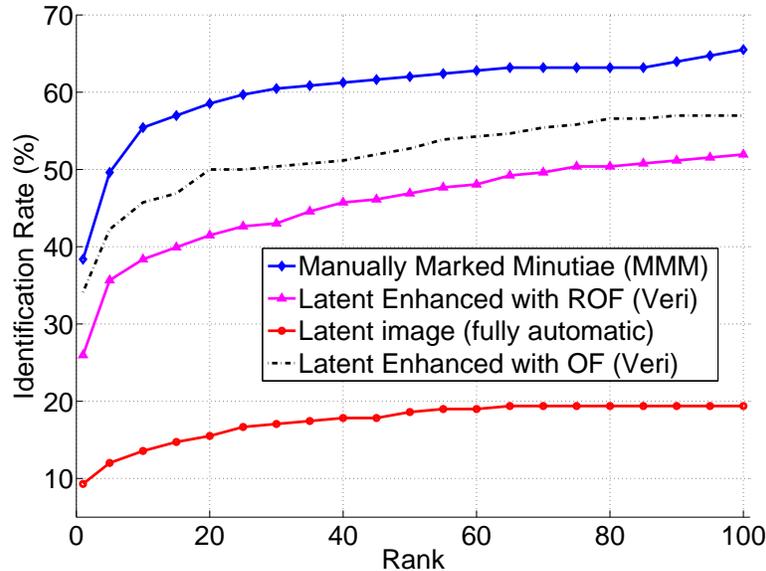


Figure 3.19: CMC curves for manually marked minutiae, enhanced image using reconstructed orientation field, automatically extracted minutiae by Verifinger from latent image and enhanced image using manually marked orientation field.

3.6.1.1 Fusion methods

We performed experiments using different fusion levels and methods. These fusion scenarios always used manually marked minutiae-based information and enhanced images using reconstructed orientation field-based information (scores, ranks, etc). We found that Borda count (rank-level) could not improve the matching performance of manually marked minutiae. This might be because Borda count method assumes all the matchers perform equally well, which is not true in our case involving two matchers. However, the highest rank fusion improved the matching performance most of the time, as can be seen in Fig. 3.20. The highest rank fusion explores the strength of each matcher more effectively, since being assigned a high rank from only one of the matchers increases the likelihood of receiving a high rank after the second sorting.

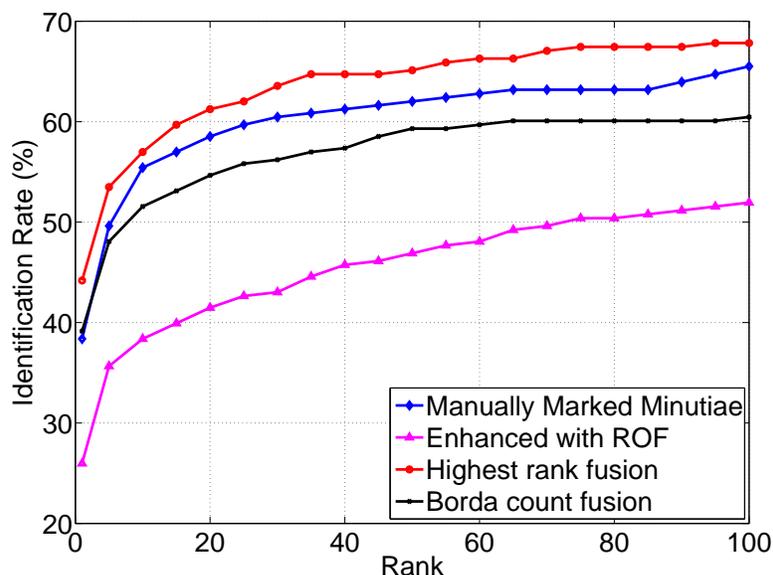
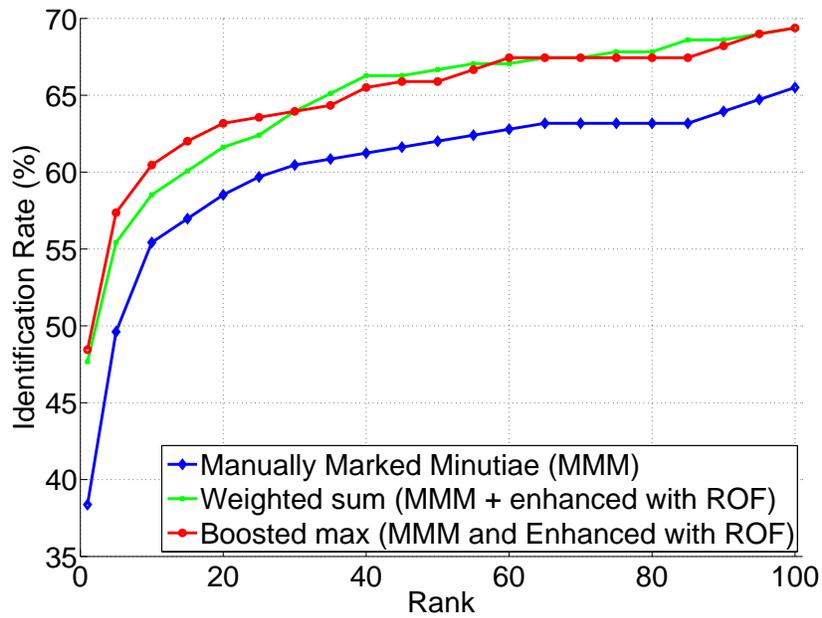


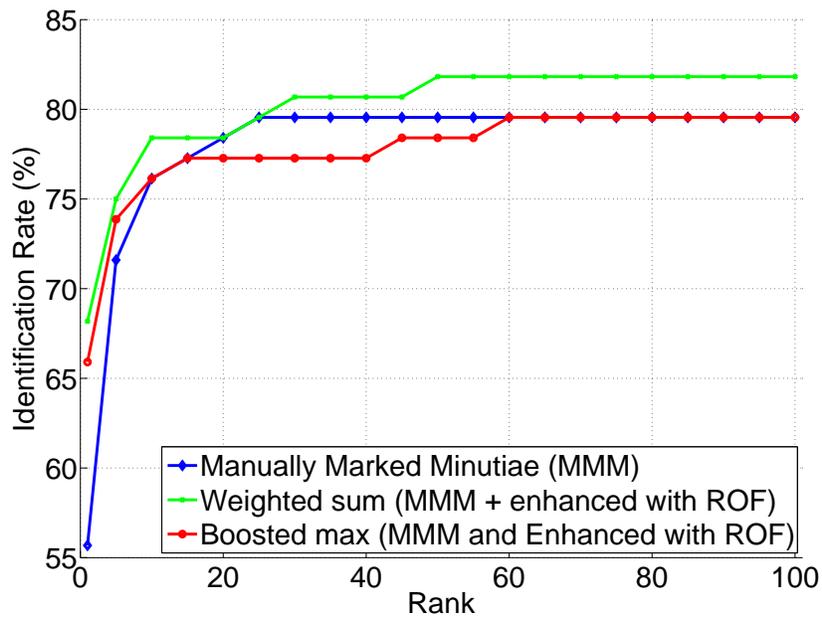
Figure 3.20: CMC curves for manually marked minutiae, enhanced image using reconstructed orientation field, highest score rank-level fusion and Borda count fusion.

Min and product fusion rules do not improve the matching performance while the Max rule provides very minor improvement. The best results were obtained using the weighted sum rule and boosted max fusion (Fig. 3.21). At rank 1, the improvement due to boosted max fusion is approximately 10% compared to the two individual matchers over all latent image quality levels as well as for each specific quality level separately. This means boosted max fusion was able to correctly find the mates of 25 additional latents than using manually marked minutiae alone. In the case of latent search, since the AFIS accuracy is not sufficiently high, the output is a list of N candidates for manual comparison by a latent examiner.

For good quality latent images, although the boosted max method performed better than manually marked minutiae in the top few ranks, its performance dropped for some ranks (Fig. 3.21b). This might indicate that the best approach for good quality latents is the simple sum rule. For bad quality images, the boosted max method outperformed manually marked minutiae by an average of 10% at all ranks (Fig. 3.21c). In the ugly quality images, boosted max fusion shows consistent improvement at all ranks compared to manually marked minutiae, and performs better for the first 30 ranks compared to the sum rule (Fig. 3.21d).



(a) All quality latents

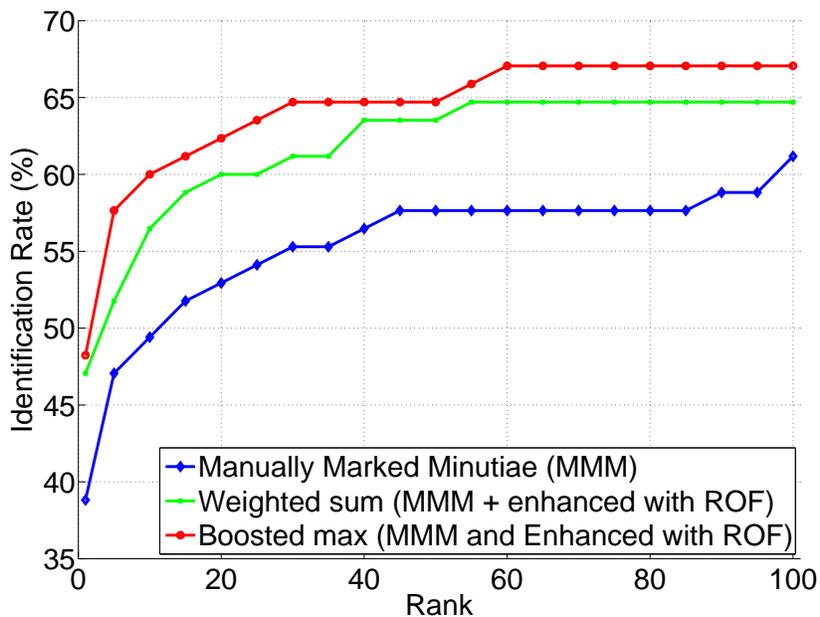


(b) Good quality latents

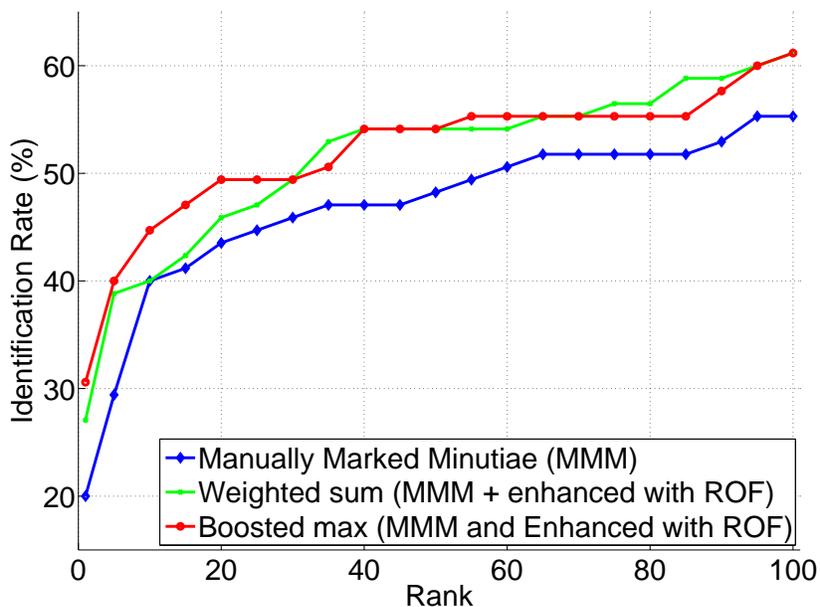
Figure 3.21: CMC curves for score level fusion for latents of all qualities.

As discussed in Section 3.4.1.2, we modified the boosted max approach to include a second weight. Therefore, two different parameters (weights w_1 and w_2) must be specified in order to

Fig. 3.21 (cont'd)



(c) Bad quality latents



(d) Ugly quality latents

apply the boosted max. These two weights were empirically determined and they do not need to sum to one because the purpose is to boost the scores of image pairs that are found to be consistent. The computation of the transformation matrices used to decide the consistency between a pair of

fingerprints in boosted max approach is based on the matched minutiae output by VeriFinger. Ideally, T_{ME} should be the Identity matrix because the two sets of minutiae are extracted from the same image. However, in practice manually markings and automatically extracted minutiae differ in their positions, which leads to a transformation matrix near to the Identity matrix, but not exactly the same.

We observed that in almost all the cases, if the mated rolled finger was ranked first in one of the matching experiments (manually marked minutiae or enhanced images), it was also ranked first in the boosted max approach. Fig. 3.22a shows an example of this case where the latent was ranked 1 by manually marked minutiae, ranked 2,403 by enhanced image, and ranked 1 by boosted max.

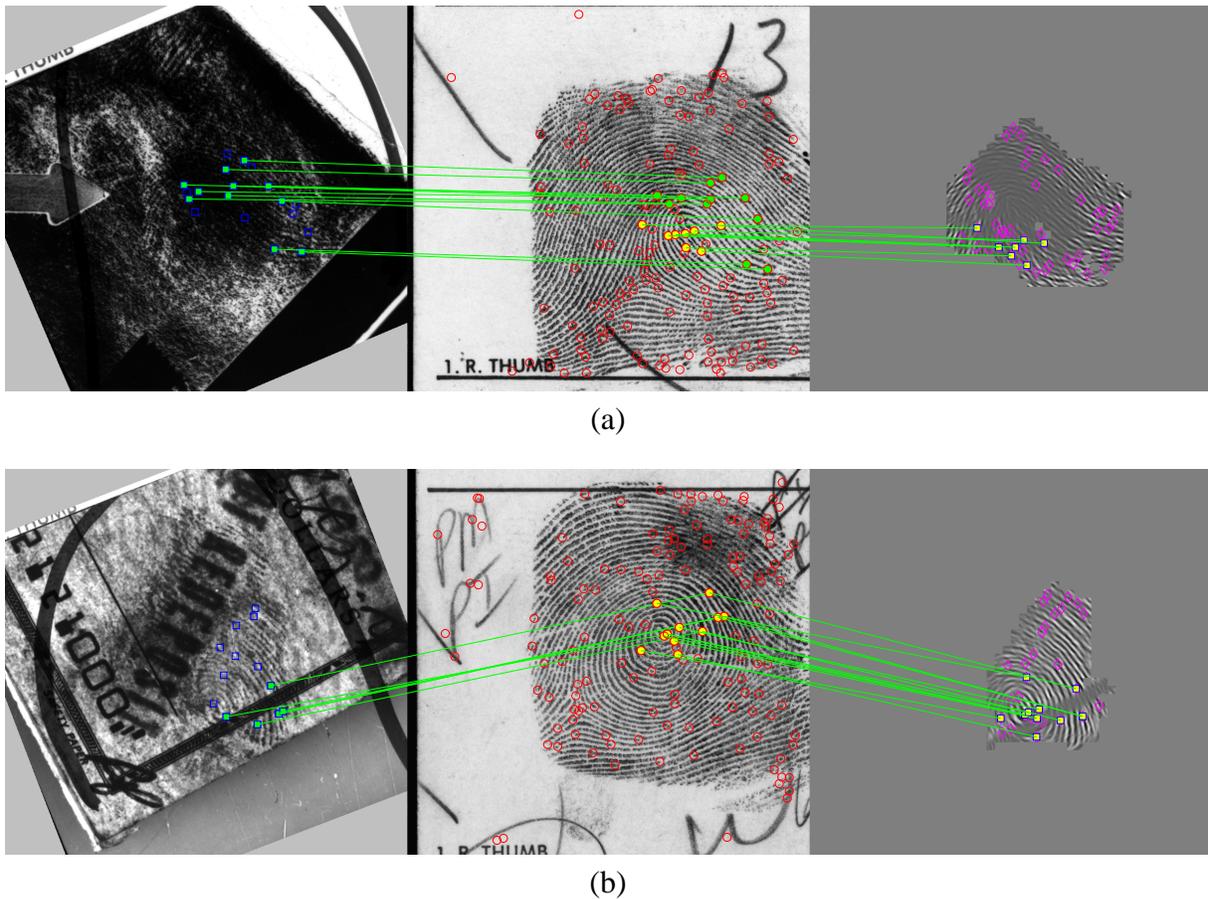


Figure 3.22: Matched minutiae shown for manually marked, rolled and enhanced latent images. (a) Boosted max retrieved the true mate for this latent at rank one even though one of the retrieved ranks (enhanced image) was 2,403 and (b) the retrieved rank for the true mate for this latent is rank one after boosted max was applied even though neither the manually marked minutiae nor enhanced image retrieved the true mate at rank one.

Boosted max ranked true mated rolled fingerprints at a lower rank in only two cases that were ranked first by manually marked minutiae (two good quality images). It also corrected the retrieved rank in eleven cases (4 good quality, 3 bad and 4 ugly quality latents). This means the mated rolled was neither ranked first in manually marked minutiae nor in enhanced images, but it was ranked first after boosted max fusion. Fig. 3.22b shows an example of a latent that was ranked as 268 by manually marked minutiae, ranked 2 by enhanced image, and ranked 1 by boosted max.

3.6.2 Latent Matching with Descriptor-Based Hough Transform

We first provide a description of the baseline algorithms to be compared with the proposed algorithm. Then we report the performances of alignment and matching. This is followed by the fusion of matchers and the effect of fingerprint quality. Then, we discuss the issue of computational cost.

3.6.2.1 Commercial Matchers

In order to compare the performance of the proposed latent fingerprint matcher, we used two commercial fingerprint matchers, Verifinger 4.2 [79] and Morpho as baseline. In addition, we also used the algorithm presented in [91, 118] as a baseline, for which the SDK was provided by the authors (MCC SDK). It should be pointed out that none of these three baseline matchers were designed specifically for the latent matching case. But, despite our efforts, we could not find any latent fingerprint matching SDK or a forensic AFIS that is available for evaluation purposes by a research lab. Still, the matchers we are using in our comparative study are well known: one of the COTS (VeriFinger) [79] has been widely used as a benchmark in fingerprint publications, and MCC is one of the best performing algorithms in FVC-onGoing [119].

3.6.2.2 Alignment Performance

In order to estimate the alignment error, we use ground truth mated minutiae pairs from NIST SD27, which are marked by fingerprint examiners, to compute the average distance between the

true mated pairs after alignment⁹. If the average Euclidean distance for a given latent is less than a pre-specified number of pixels in at least one of the ten best alignments (peaks in the Descriptor-Based Hough Transform), then we consider it a correct alignment. This alignment performance is shown in Fig. 3.23 for the NIST SD27 latent database. The x-axis shows the misalignment threshold¹⁰, and the y-axis shows the percentage of correctly aligned latent fingerprints in at least one of the ten top alignments. For comparison, we show the accuracy of aligning the minutiae sets based on the peaks of the Generalized Hough Transform (GHT) and based on the most similar minutiae pair (according to the MCC similarity)¹¹. Two latent alignment examples are given in Figs. 3.24 and 3.25 to show the alignment results by DBHT and GHT. As we can see from these figures, the proposed algorithm is superior to GHT for difficult latents where the number of minutiae is small.

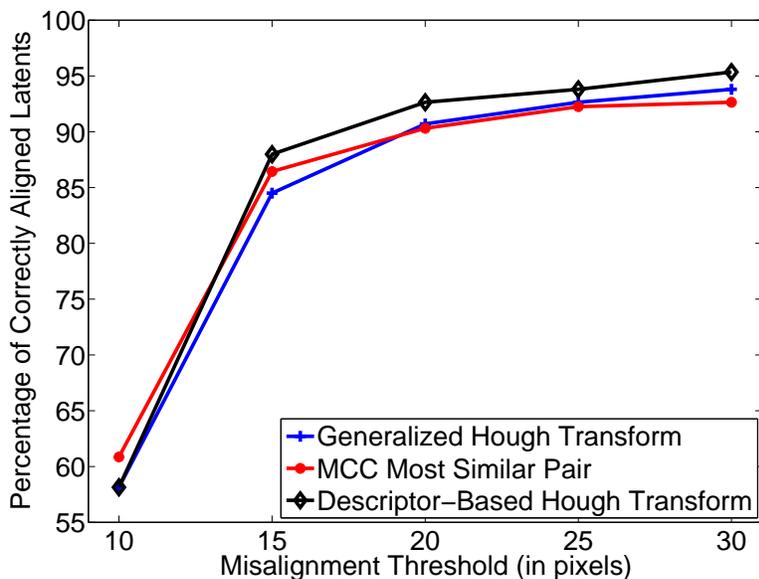


Figure 3.23: Alignment Accuracy: percentage of correctly aligned latents vs. misalignment threshold.

⁹Here we use the term ground truth minutiae to refer to minutiae which are marked by latent examiners by looking at the latent and the corresponding rolled print at the same time, and we use the term manually marked minutiae to refer to minutiae which are also marked in the latent by latent examiners, but without looking at the true mate (rolled print).

¹⁰The alignment is deemed as incorrect if the average distance between mated minutiae pairs after alignment is larger than this threshold.

¹¹In this case, each alignment is based on one of the ten most similar minutiae pairs.

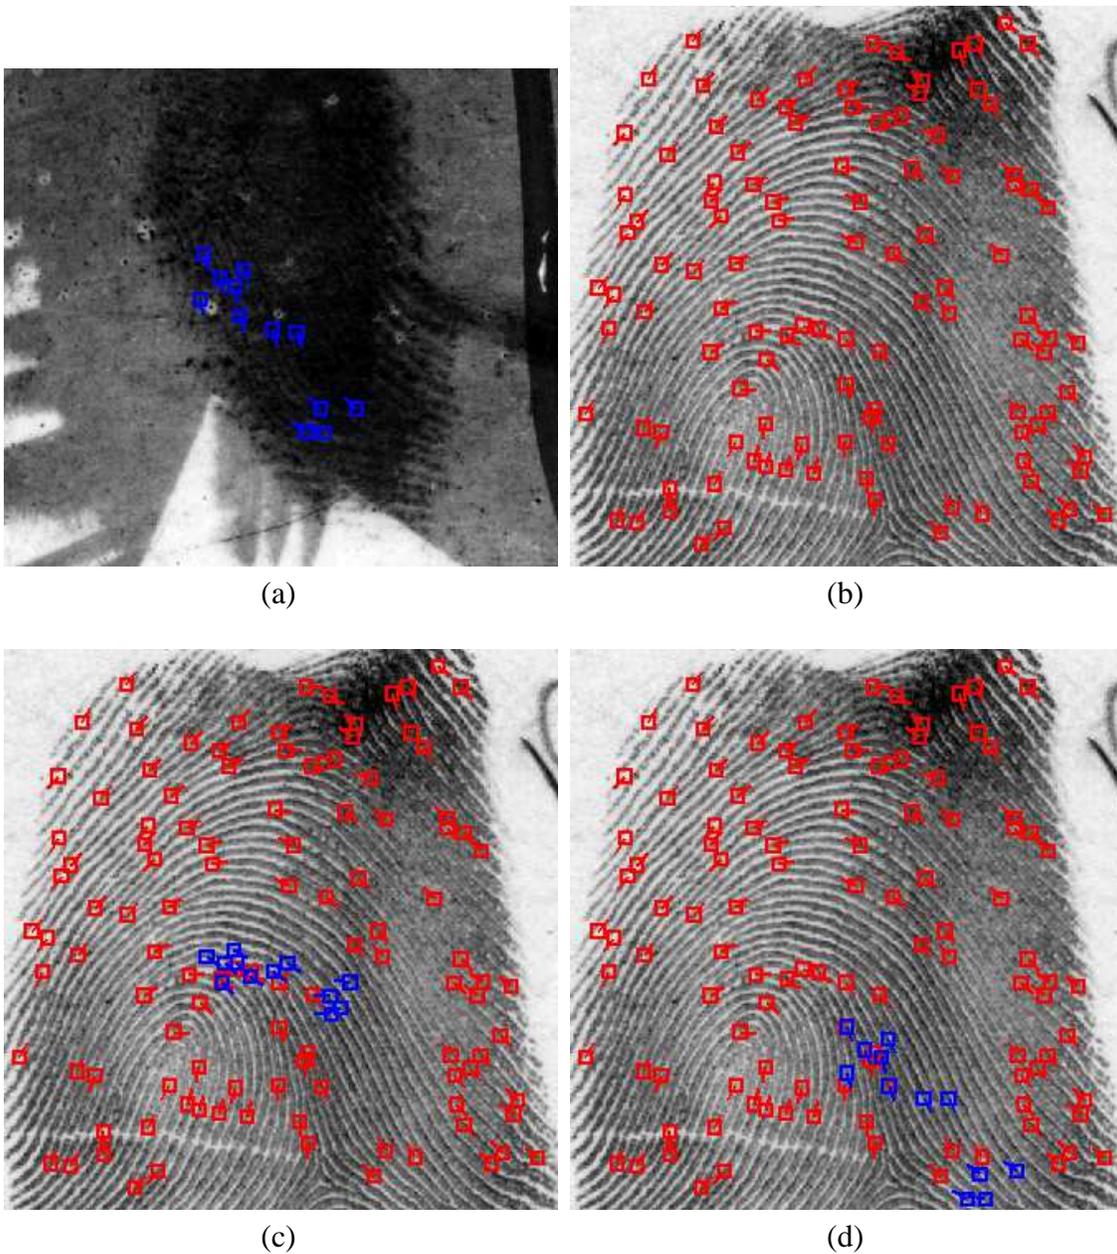


Figure 3.24: Example in which DBHT (Descriptor-Based Hough Transform) alignment is better than GHT (Generalized Hough Transform) alignment. (a) Latent with manually marked minutiae, (b) corresponding rolled print with automatically extracted minutiae, (c) rolled print with latent minutiae aligned by GHT, and (d) aligned by DBHT.

There are very few errors in alignment if we set the threshold value of misalignment as 20 pixels. One of the reasons for these failure cases is there are a very small number of true mated minutia pairs in the overlapping area between the latent and mated rolled print. As a result, not

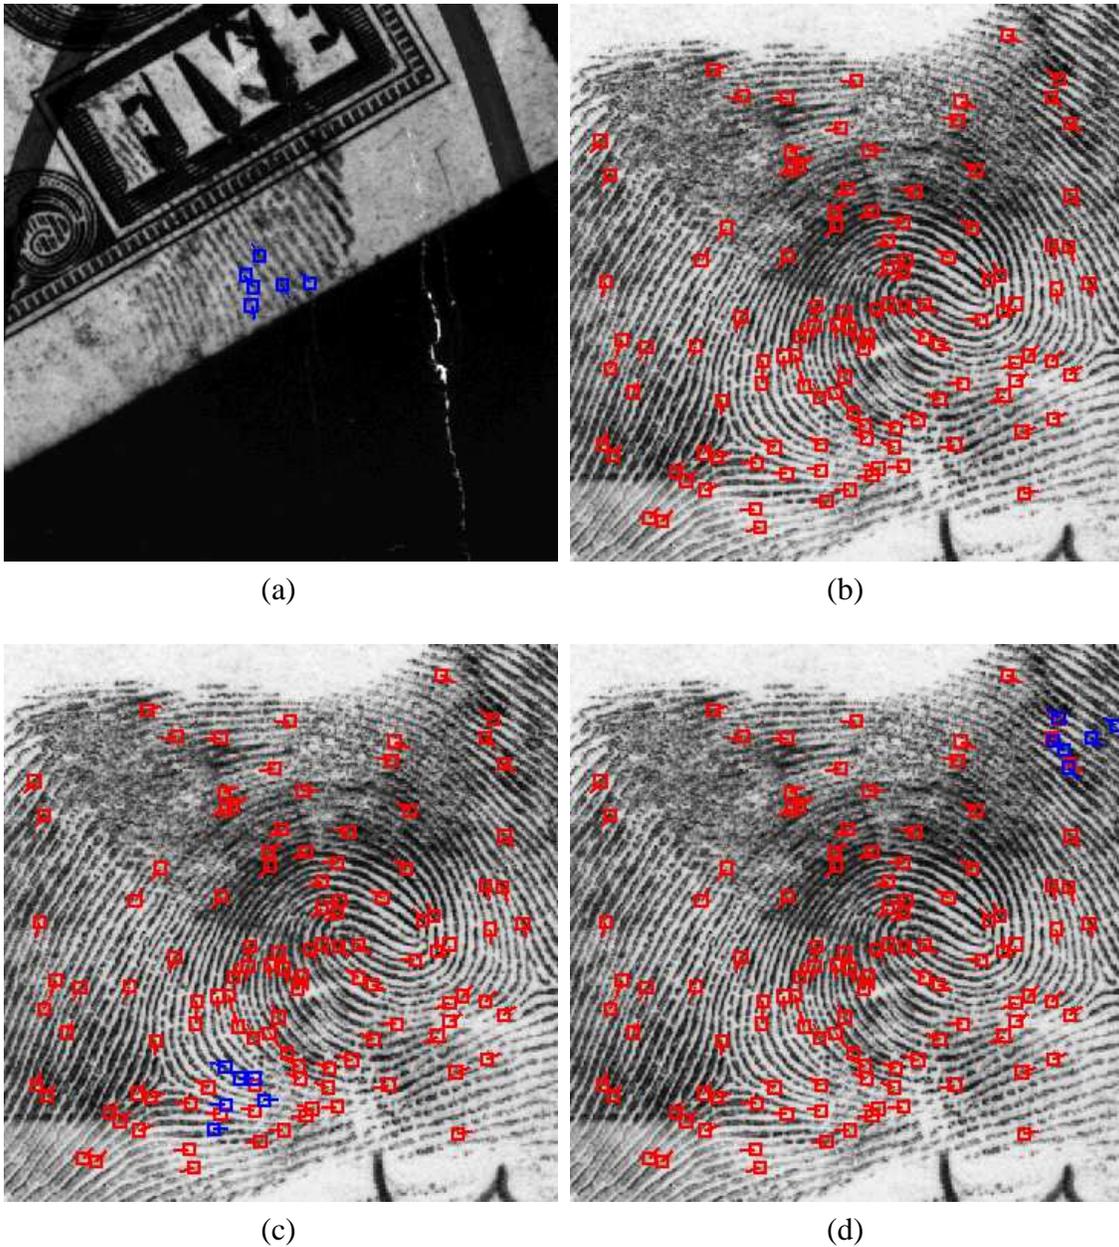


Figure 3.25: Example in which DBHT (Descriptor-Based Hough Transform) alignment is better than GHT (Generalized Hough Transform) alignment. (a) Latent with manually marked minutiae, (b) corresponding rolled print with automatically extracted minutiae, (c) rolled print with latent minutiae aligned by GHT, and (d) aligned by DBHT.

many true mated pairs vote for the correct alignment parameters. The absence of true mated pairs is due to a limited number of minutiae in latents and the error in minutiae detection in the rolled print. One such example is shown in Fig. 3.26. Blue squares represent manually marked minutiae

in the latent print (left), red squares represent automatically extracted minutiae in the rolled print (right), and the green line indicates the only true mated minutiae pair available for this (latent, rolled) image pair.

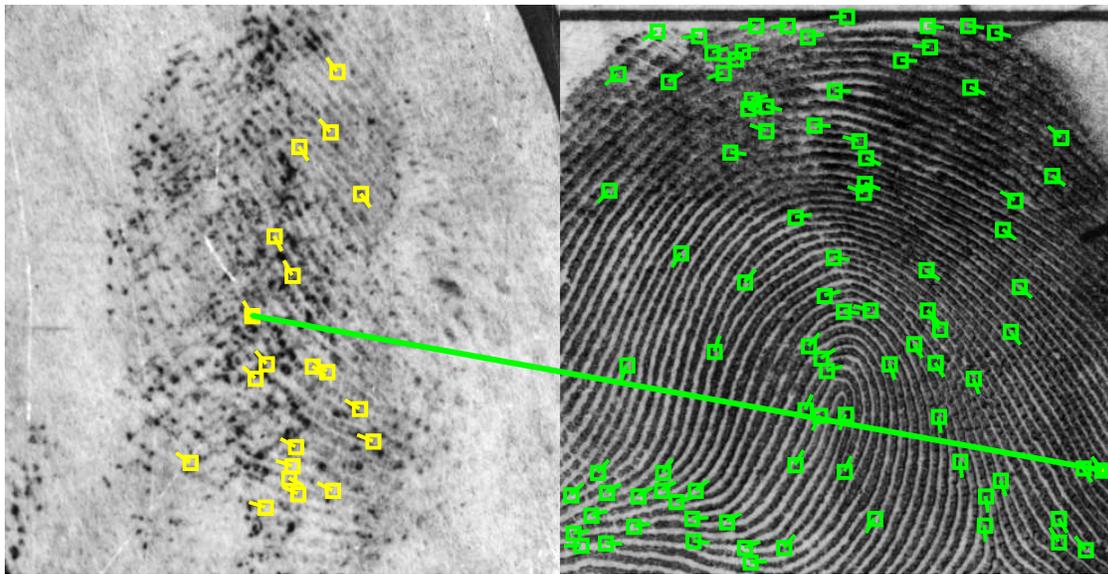


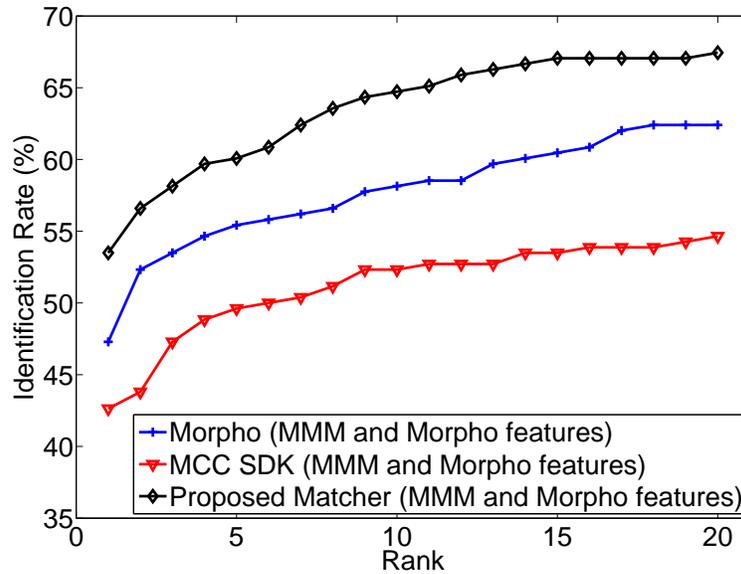
Figure 3.26: Example of alignment error due to the small number of true mated minutia pairs in the overlapping area between a latent and its mated rolled print. Note that there is only one aligned minutiae pair here.

3.6.2.3 Matching Performance

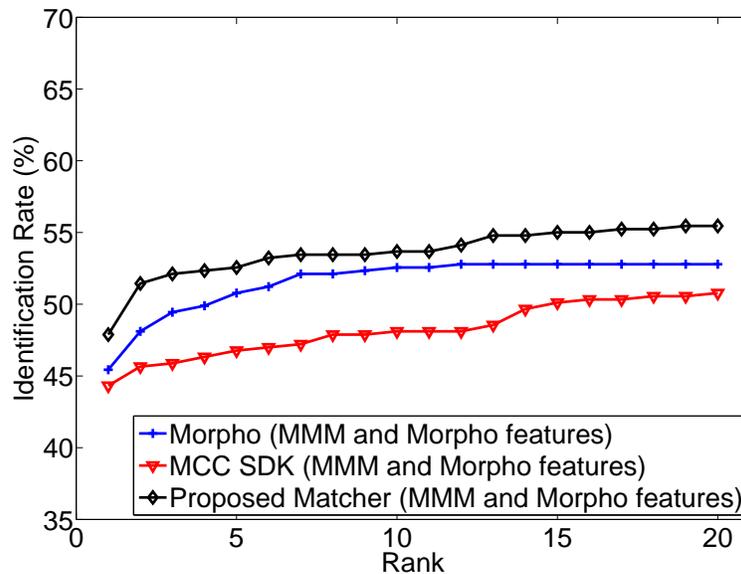
In the identification scenario, the size of the background database (or gallery) significantly affects the identification accuracy. Therefore, to make the problem more challenging and realistic, we built a large background database of rolled prints containing 258 mated rolled prints from NIST SD27, 4,740 rolled prints from WVU database, and 27,000 rolled prints from the NIST Special Database 14 [90]. Therefore, the total number of rolled prints in the background database is 31,998 from a combination of the rolled prints in the three databases.

Minutia Cylinder Code (MCC) is used as the local descriptor for minutiae in our experiments. The local descriptors are built using MCC SDK, which uses the bit-based implementation (binary descriptors) [118]. The parameters used for MCC are set as suggested in [118], with the number

of cells along the cylinder diameter as 16 (N_s). In our method, the local descriptor similarities are used in both the alignment and scoring modules.



(a) NIST SD27



(b) WVU LFD

Figure 3.27: Performance of Morpho, MCC SDK, and Proposed Matcher when the union of manually marked minutiae (MMM) extracted from latents and automatically extracted minutiae by Morpho from rolled prints is input to the matchers.

Our matcher and MCC SDK both take minutiae as input. In the latent cases, we use man-

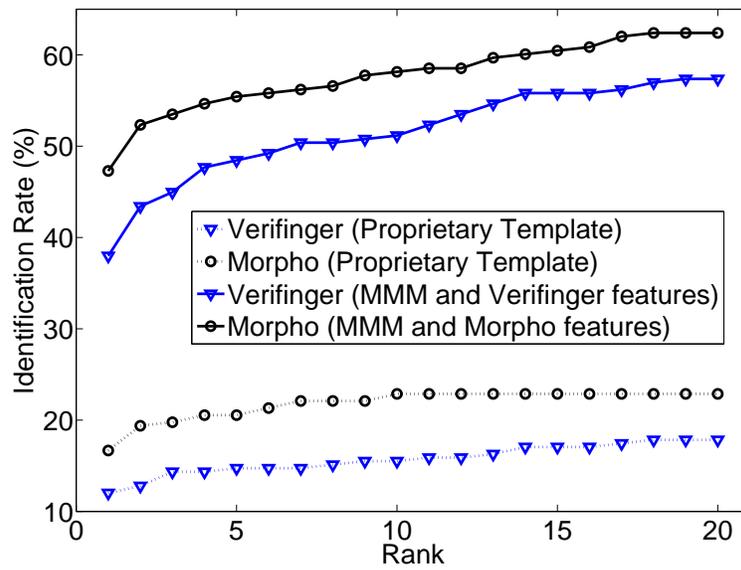
ually marked minutiae. For the rolled prints, we used both the COTS to extract minutiae. The performance of the proposed matcher using minutiae extracted from rolled prints using Morpho is slightly worse on the NIST SD27 database compared to the performance using minutiae extracted using Verifinger; however, for WVU LFD, using Morpho minutiae yielded a significantly better performance compared to the performance using minutiae extracted using Verifinger. This demonstrates that the performance of COTS can be significantly affected by the image quality. Overall, since minutiae extracted from Morpho resulted in a better performance, we only report the results in which minutiae are extracted using Morpho. Fig. 3.27 shows the performance of Morpho, MCC SDK, and the proposed matcher using manually marked minutiae in latents and automatically extracted minutiae by Morpho in rolled prints. The proposed approach outperforms the other fingerprint matchers used in our study.

In our earlier study [120], we did not use the bit-based implementation of MCC representation and we only evaluated the algorithm on NIST SD27. We used our own implementation of MCC with the parameters suggested in [91]. However, the bit-based implementation and the more optimized parameters suggested in [118] yielded a much better performance on WVU LFD, and a decrease in performance on the NIST SD27 database. We decided to use the bit-based implementation with the new parameters because of its better overall performance, and because we wanted to use a consistent framework for both the databases.

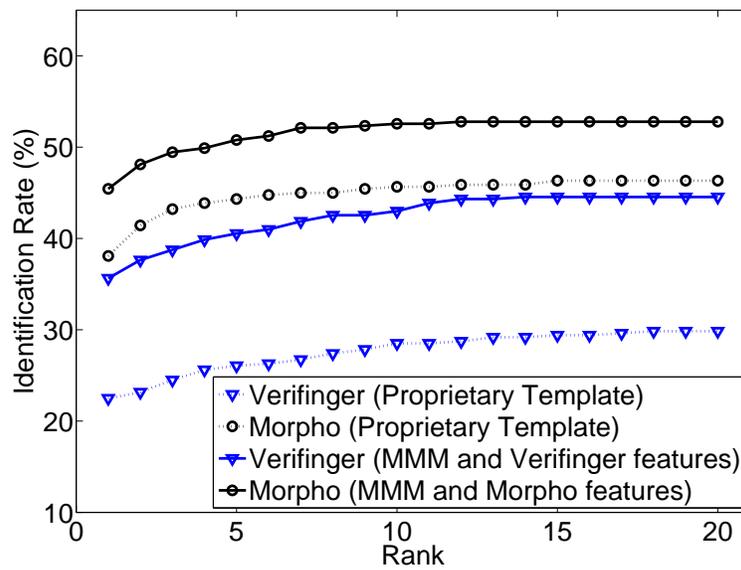
It is worth noting that the matching performance on WVU LFD when manually marked minutiae are used is generally worse than the performance on NIST SD27. We believe this is due to a number of factors: (i) there are 14 latents with less than 3 manually marked minutiae in WVU LFD, while the minimum number of manually marked minutiae in NIST SD27 latents is 5; (ii) while the genuine (latent, rolled) pairs were provided with the database, after we examined the images in the WVU database we identified some that appeared to be wrongly paired; (iii) the quality of the mates (rolled prints) is slightly worse in WVU LFD than in NIST SD27. We did not exclude any of the latents or (latent, rolled) mated pairs from the WVU database (from cases (i) and (ii)) to allow for a fair comparison by other researchers with our results.

The performance of the COTS matchers, each using its own proprietary templates for latents (including automatically extracted minutiae and possibly other features), is worse than using manually marked minutiae for both the databases. However, the gap between the performances of manually marked minutiae and of proprietary template is much larger in the case of NIST SD27 than in the case of WVU latent database. This is probably due to the characteristics of these two databases. Note that WVU is a laboratory collected database and so most of the latents in it do not contain background noise. On the other hand, in NIST SD27 the images are of operational casework quality and most of the latents contain a large amount of background noise, which poses a challenge in automatic feature extraction. Fig. 3.28 shows the performance of the two COTS matchers using both manually marked minutiae and proprietary templates (automatically extracted minutiae) for NIST SD27 and WVU databases.

There have been several studies on latent matching reported in the literature. Almost all of them are based on NIST SD27. Table 3.1 shows the reported results on the matching performance for NIST SD27 database. There is no reported performance on the WVU latent database. It should be noticed that most of the reported results cannot be directly compared mainly because of two factors: (i) the amount of input information related to the latent fingerprint, which could be automatically extracted features, or manually marked features such as minutiae, singular points, quality map, etc, or a combination of both; and (ii) some differences in the composition of the background databases and their size. In Table 3.1 we show the reported rank-1 accuracy, the manually marked latent features used in each method, and the size of the background database used. One of the results that could be almost directly compared to our results is the reported rank-1 accuracy (34.9%) in [28] where only manually marked minutiae was used as input, which is the same scenario as in our proposed matcher. The proposed matcher achieves a significantly higher rank-1 accuracy of 53.5% with similar background database size and images as in [28]. The other approach presented here (Paulino *et al.* [87]) can also be compared to this proposed matcher, although the former requires additional manually marked input features beyond minutiae.



(a) NIST SD27



(b) WVU LFD

Figure 3.28: Performance comparison using manually marked minutiae (MMM) and automatically extracted minutiae from latents.

Fig. 3.29 shows two examples of latent prints in WVU LFD correctly identified at rank-1 by

¹²SP: singular points, ROI: region of interest, RQM: ridge quality map, RFM: ridge flow map, RWM: ridge wavelength map.

Table 3.1: Comparison of rank-1 accuracies reported in the literature for the NIST SD27 database.

Work	Manually Marked Latent Features	Background Database Size	Rank-1 Accuracy (%)
Jain and Feng [28]	Minutiae, Skeleton, SP, ROI, RQM, RFM, RWM ¹²	29,257	74.0
Proposed Matcher [88] + Morpho	Minutiae	31,998	57.4
Proposed Matcher [88]	Minutiae	31,998	53.5
Paulino <i>et al.</i> [87]	Minutiae, SP, ROI ¹²	27,258	48.0
Jain and Feng [28]	Minutiae	29,257	34.9
Yoon <i>et al.</i> [101]	SP, ROI ¹²	27,258	26.0
Feng <i>et al.</i> [102]	None	27,258	25.0

the proposed matcher. Even though the number of minutiae in these latents is small, they could still be correctly identified. The ranks of the true mates using Morpho matcher are 1871 and 181, respectively, and the matched minutiae shown are from our proposed matcher.

Fig. 3.30 shows examples of latent prints in NIST SD27 and in WVU LFD whose mated reference prints are not included in the top 20 candidates by the proposed matcher, but were correctly identified at rank-1 by Morpho matcher. The ranks of these latents using the proposed matcher are 3626 and 64, respectively, and the matched minutiae shown are from our proposed matcher. In the first latent, a large number of minutiae do not have mated minutiae due to missing minutiae in the rolled print, and therefore the score is not as high as for impostor pairs in which many more minutiae could be matched. In the second case, we can see that the minutiae marked in the latent are relatively sparse, while the minutiae automatically extracted in the rolled print are denser. These facts make local neighborhoods (and descriptors) very different between the latent and its true mate, leading to a low match score.

3.6.2.4 Fusion of Matchers

We noticed that the two most accurate matchers (the proposed matcher and Morpho) perform differently on different latents, meaning they are complementary to each other. This suggests that a

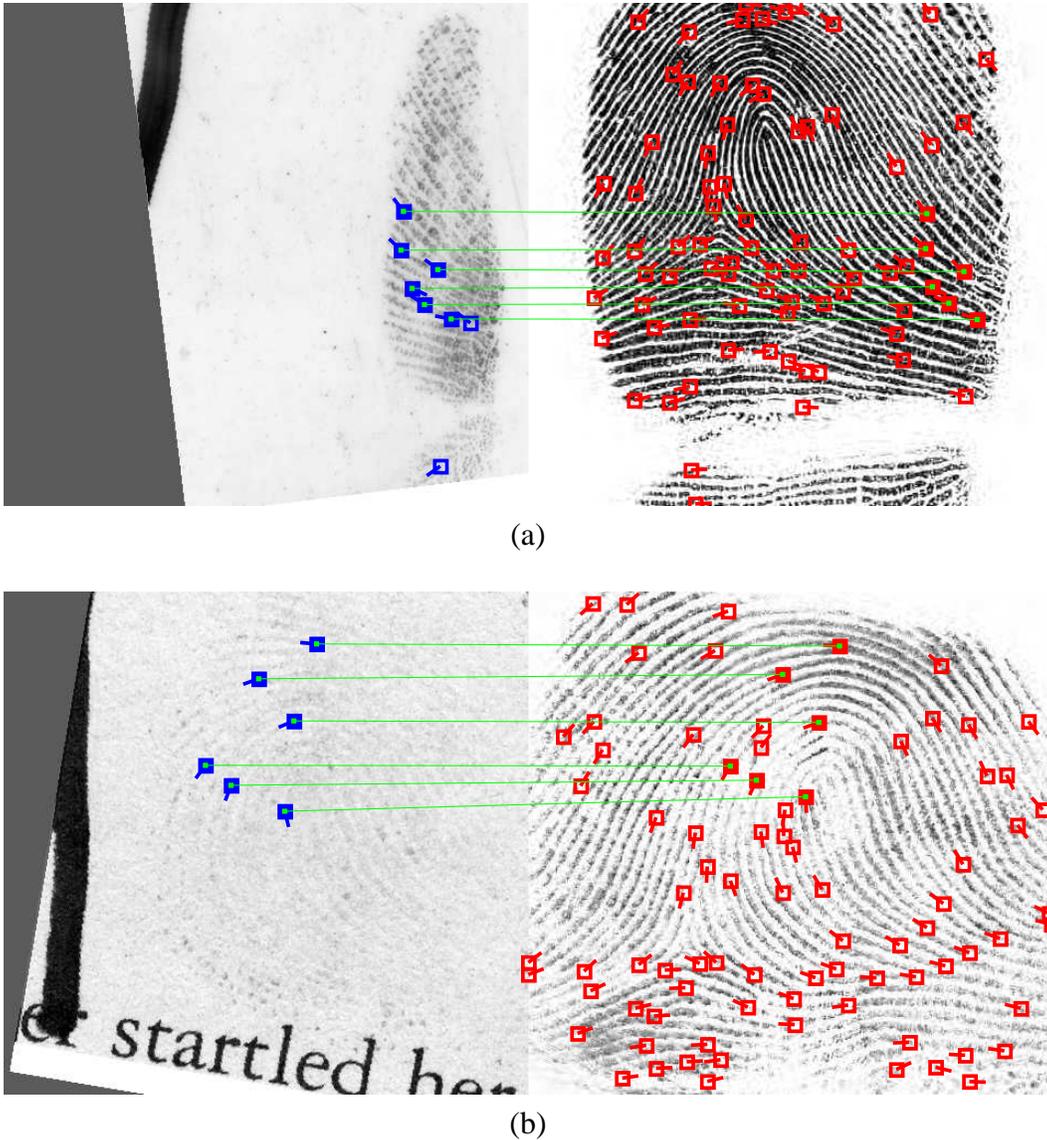
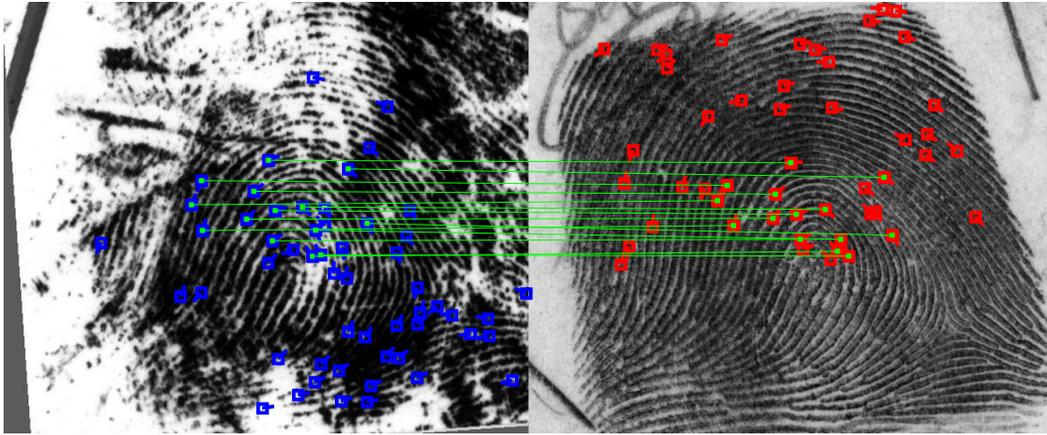
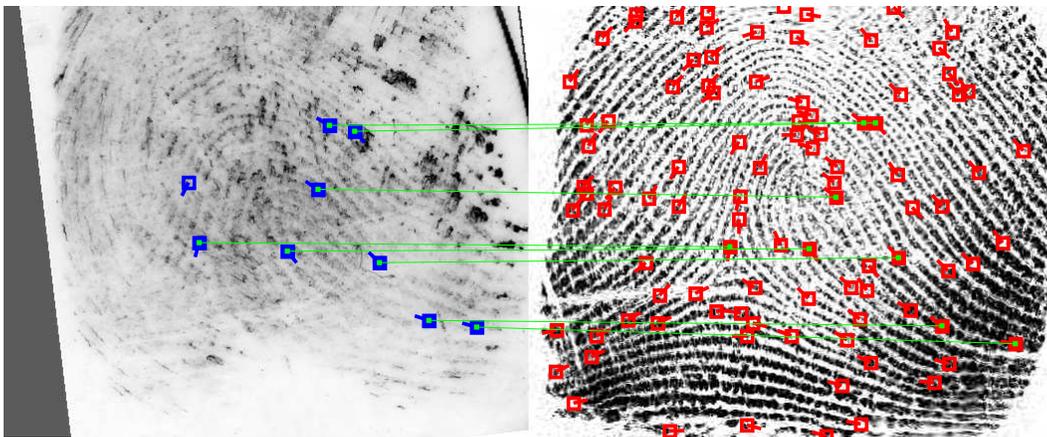


Figure 3.29: Two latent prints from the WVU database correctly identified at rank-1 by the proposed matcher but ranked below 20 by Morpho.

fusion of these two matchers would result in a better performance. We performed a score-level fusion of these two matchers. The scores from Morpho matcher were normalized to the range $[0, 1]$ for each latent (local min-max normalization) because local normalization was shown to perform better than global normalization in the identification scenario [121]. Although the proposed matcher and Morpho matcher have similar matching accuracy, the fusion weights selected (0.8 and 0.2) were not equal because of the large range of the scores for the Morpho matcher. The performance improvement obtained by the score-level fusion of Morpho matcher and the proposed



(a)

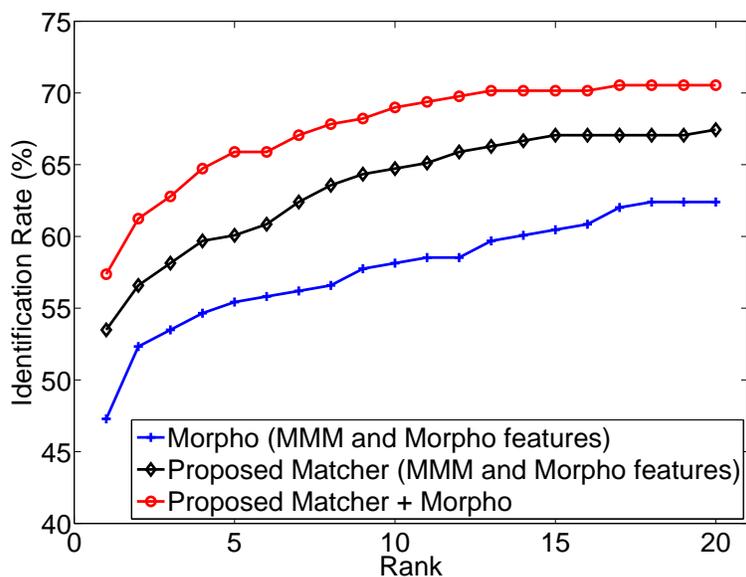


(b)

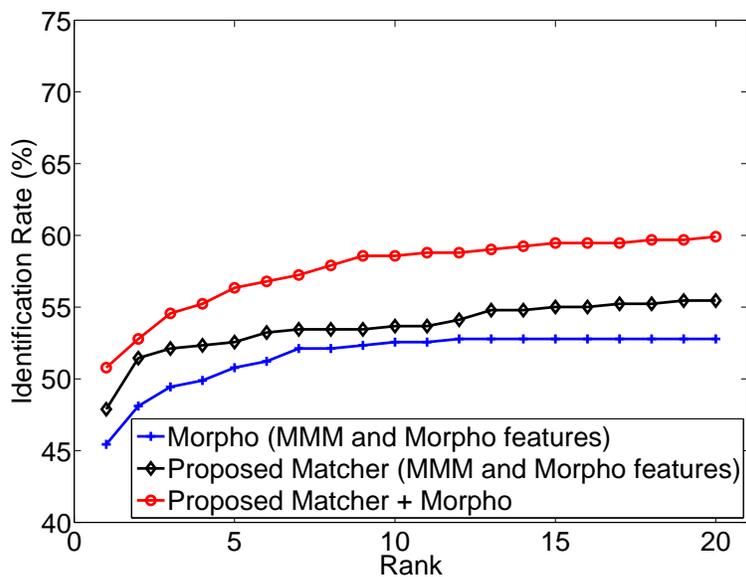
Figure 3.30: Latent prints ((a) from NIST SD27 and (b) from WVULFD) whose mates were not retrieved in the top 20 candidates by the proposed matcher but correctly identified at rank-1 by Morpho matcher.

matcher is shown in Fig. 3.31 for both the databases. Some latents for which the fusion of the two matchers (Morpho and proposed matcher) improved the ranks of the true mates compared to the retrieval ranks by the individual matchers separately are shown in Figs. 3.32 and 3.33. Note that like those mated pairs (shown in Fig. 3.29 and Fig. 3.30) identified at rank-1 by either one the two matchers, mated pairs (shown in Fig. 3.32 and Fig. 3.33) which both matchers failed to identify at rank-1 also benefit from the fusion. The reason is the scores of non-mated pairs given by the two matchers are not consistent.

Improvements in matching accuracy were also obtained by combining the proposed matcher



(a) NIST SD27



(b) WVU LFD

Figure 3.31: Score-level fusion of the proposed matcher and Morpho for NIST SD27 and WVU databases.

and other matchers in our study (Verifinger and MCC SDK), but they are not reported here because the fusion performance with Morpho was consistently better than the performance of Verifinger and of MCC SDK. We also performed rank-level fusion using the highest rank and Borda Count

Table 3.2: Rank-1 accuracies for various subjective quality levels of latents in NIST SD27.

Quality	Verifinger (%)	Morpho (%)	MCC (%)	Proposed (%)
All	38.0	47.3	42.6	53.5
Good	55.7	70.5	69.3	75.0
Bad	36.5	36.5	31.8	47.1
Ugly	21.2	34.1	25.9	37.6

Table 3.3: Rank-1 accuracies for various objective quality values of latents in NIST SD27 (large, medium and small refer to the number of minutiae in the latent).

Quality	Verifinger (%)	Morpho (%)	MCC (%)	Proposed (%)
All	38.0	47.3	42.6	53.5
Large	59.3	73.3	70.9	75.6
Medium	43.5	45.9	43.5	56.5
Small	11.5	23.0	13.8	28.7

methods [122]. However, since score-level fusion showed a better performance, we only report here results for score-level fusion.

3.6.2.5 Effect of Fingerprint Quality

In Section 3.5, we discuss how the quality of the latent fingerprints can be measured subjectively (assigned by latent experts as in NIST SD27) and objectively (based on the number of minutiae available). Rank-1 accuracies are shown for each quality separately in Tables 3.2, 3.3, and 3.4 for both the latent databases. We can see that the matching performance is highly correlated with the number of minutiae available in the latent prints. The performance of the proposed matcher is consistently better over all qualities and for both the databases.

The quality of reference prints also has a large impact on the matching accuracy. In Fig. 3.18, the histograms of NFIQ quality values for the corresponding rolled prints in each latent database are shown. According to the NFIQ quality measure, the quality of the rolled prints in WVU database is slightly worse than the quality of the rolled prints in NIST SD27. The NFIQ quality measure is an integer value in the range 1 to 5, where 1 is the highest quality and 5 is the worst

Table 3.4: Rank-1 accuracies for various objective quality values of latents in WVU LFD (large, medium and small refer to the number of minutiae in the latent).

Quality	Verifinger (%)	Morpho (%)	MCC (%)	Proposed (%)
All	35.6	45.4	44.3	47.9
Large	63.5	73.1	74.0	74.5
Medium	28.8	45.0	37.5	45.0
Small	3.1	9.9	9.3	14.9

Table 3.5: Rank-1 accuracies for latents grouped according to NFIQ quality values of corresponding rolled prints in NIST SD27.

Quality	Verifinger (%)	Morpho (%)	MCC (%)	Proposed (%)
All	38.0	47.3	42.6	53.5
NFIQ \leq 3	42.1	54.9	49.4	60.4
NFIQ $>$ 3	30.9	34.0	30.9	41.5

Table 3.6: Rank-1 accuracies for latents grouped according to NFIQ quality values of corresponding rolled prints in WVU LFD.

Quality	Verifinger (%)	Morpho (%)	MCC (%)	Proposed (%)
All	35.6	45.4	44.3	47.9
NFIQ \leq 3	36.9	50.0	48.0	52.4
NFIQ $>$ 3	34.4	41.1	40.6	43.3

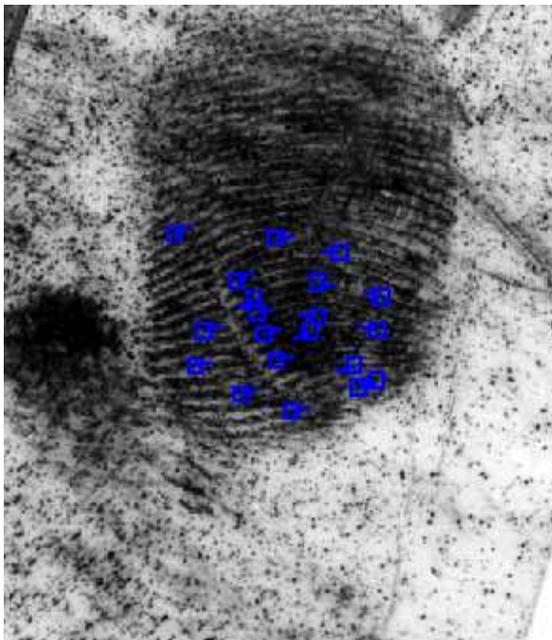
quality. We observed a significant difference in the matching performance when the latents were divided into the following two quality groups: (i) rolled prints are of good quality (NFIQ value of 1, 2 and 3), and (ii) rolled prints are of poor quality (NFIQ values of 4 and 5). The difference in matching performance between good NFIQ and poor NFIQ qualities for all matchers ranges from 11 – 21% for NIST SD27, while it ranges from 2 – 9% for WVU database (see Tables 3.5 and 3.6). As an example, the rank-1 accuracy of Morpho matcher on NIST SD27 is 54.9% and 34.0% for good and poor NFIQ quality, respectively.

3.6.2.6 Computational Cost

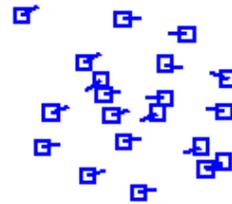
The implementation of our matching algorithm is in Matlab. The average speed of our matcher running on a PC with Intel Core 2 Quad CPU and Windows XP operating system is around 10 matches per second and it depends on the number of minutiae (the smaller the number of minutiae, the faster the matching speed). Multi-thread capability was not utilized. The majority of the running time (70%) is spent matching the local minutiae descriptors. In a C/C++ implementation, this matching would be much faster than in Matlab because of the nature of the MCC descriptors (binary). We have not optimized the code for speed.

3.6.2.7 Comparison of the two proposed methods

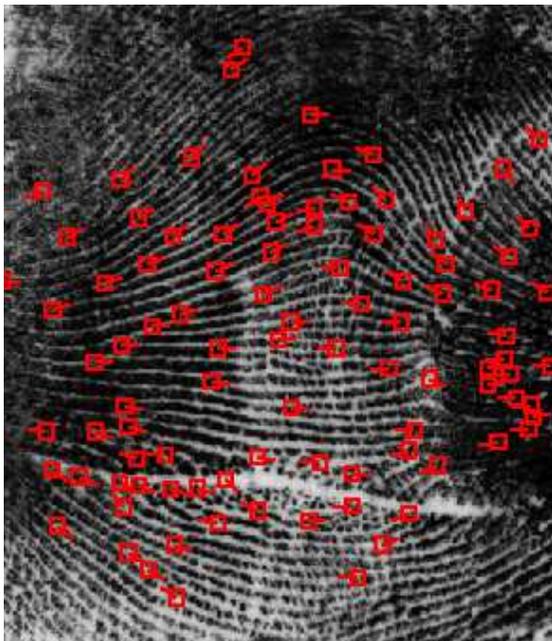
The manually marked input features are essentially the same for both the proposed methods: minutiae, singular points and region of interest in the enhancement, and minutiae in the descriptor-based Hough transform approach. The performance of the latent matching approach using descriptor-based Hough Transform is better than the performance of latent matching approach by enhancement. However, the improvement by enhancement has the advantage in terms of computational efficiency, because of the use of the commercial matcher in both steps (matching using manually marked minutiae and matching using minutiae automatically extracted from enhanced images). Furthermore, by simply combining the two approaches at the score level, the rank-100 accuracy increases from 72% to 81% compared to the descriptor-based Hough Transform matching alone (see Fig. 3.34).



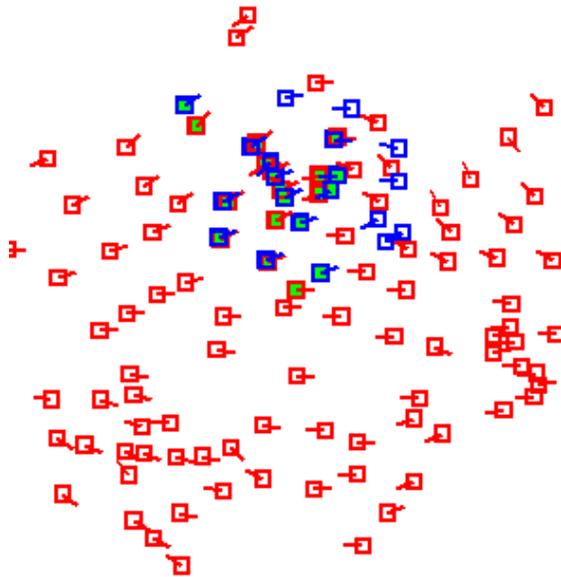
(a)



(b)



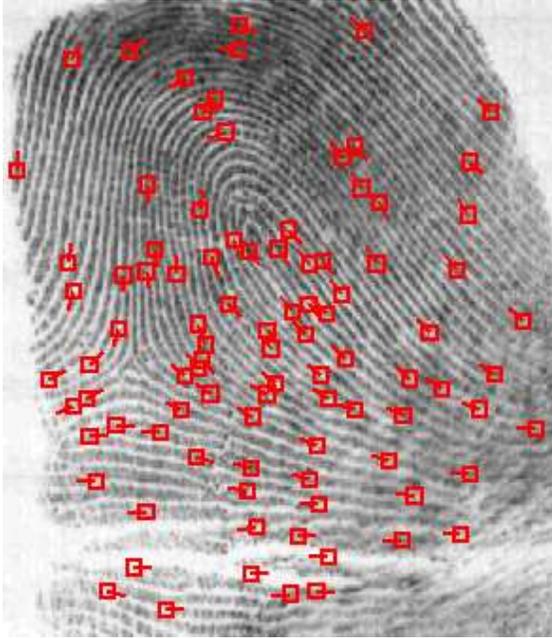
(c) (5, 10)



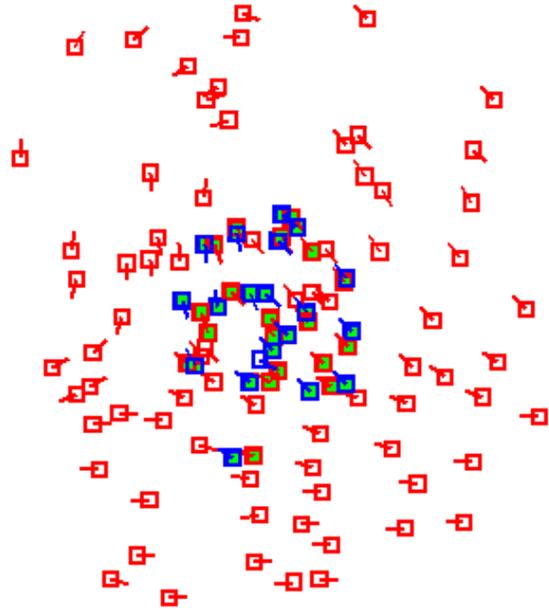
(d)

Figure 3.32: Latent print mate from NIST SD 27 identified at rank 1 after score-level fusion of Morpho and proposed matcher. The first column shows (a) a latent, (c) its true mate, (e) rank-1 non-mate by the proposed matcher, and (g) rank-1 non-mate by Morpho matcher. The second column shows (b) latent minutiae, (d), (f), and (h) latent minutiae (in blue) aligned by the proposed matcher to the rolled print minutiae shown in (c),(e) and (g). In (c), (e), and (g), the numbers in parentheses indicate the ranks at which true mated rolled print was retrieved by the proposed matcher and Morpho matcher, respectively.

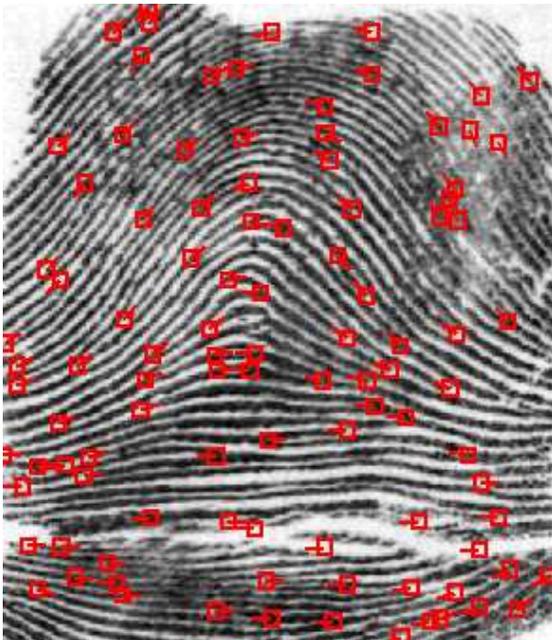
Fig. 3.32 (cont'd)



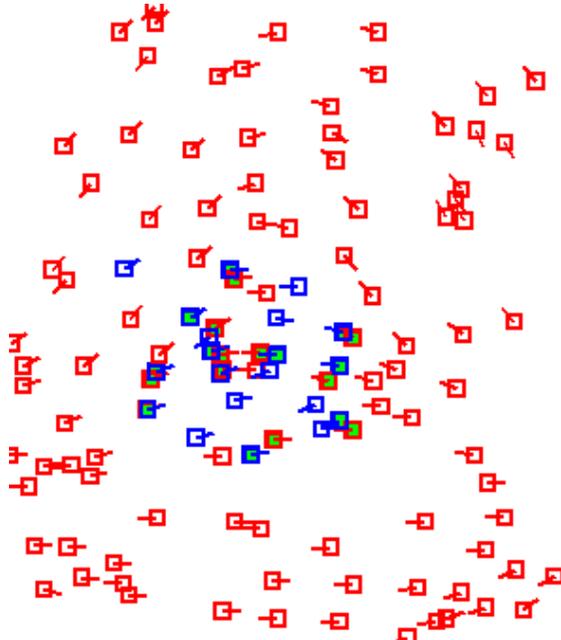
(e) (1, 30848)



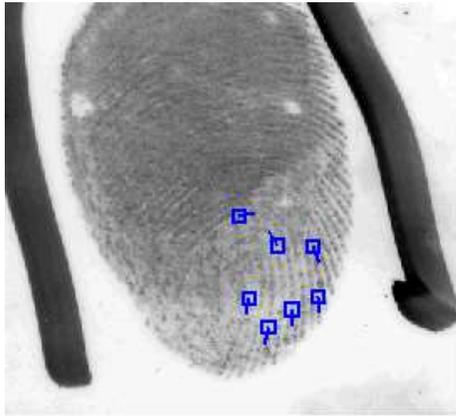
(f)



(g) (3617, 1)



(h)



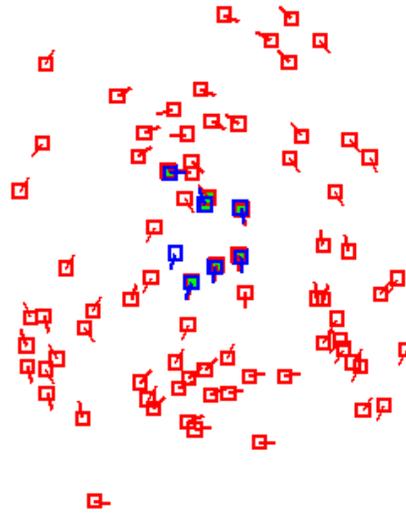
(a)



(b)



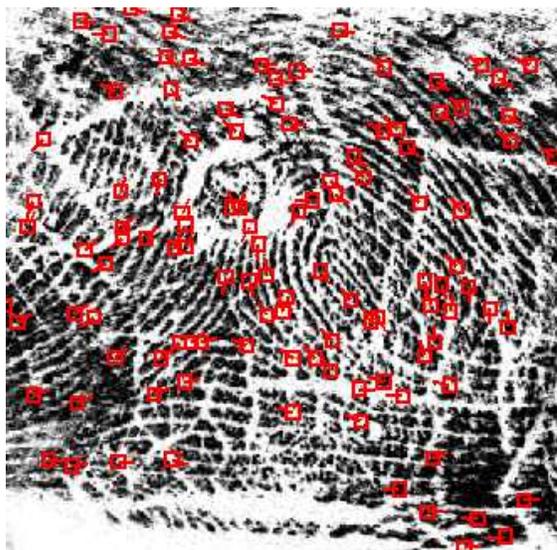
(c) (2, 42)



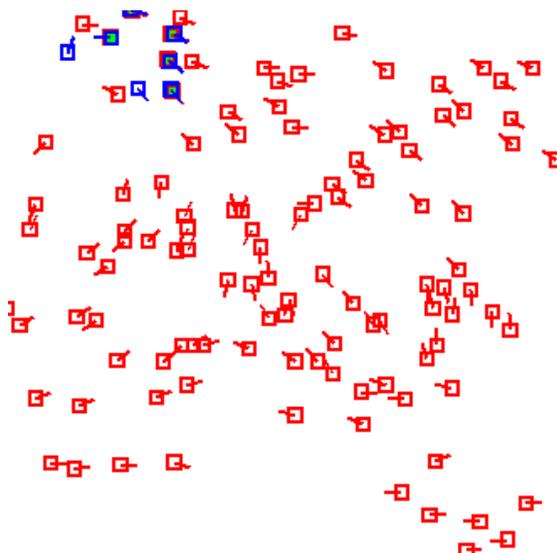
(d)

Figure 3.33: Latent print mate from WVU LFD identified at rank 1 after score-level fusion of Morpho and proposed matcher. The first column shows (a) a latent, (c) its true mate, (e) rank-1 non-mate by the proposed matcher, and (g) rank-1 non-mate by Morpho matcher. The second column shows (b) latent minutiae, (d), (f), and (h) latent minutiae (in blue) aligned by the proposed matcher to the rolled print minutiae shown in (c),(e) and (g). In (c), (e), and (g), the numbers in parentheses indicate the ranks at which true mated rolled print was retrieved by the proposed matcher and Morpho matcher, respectively.

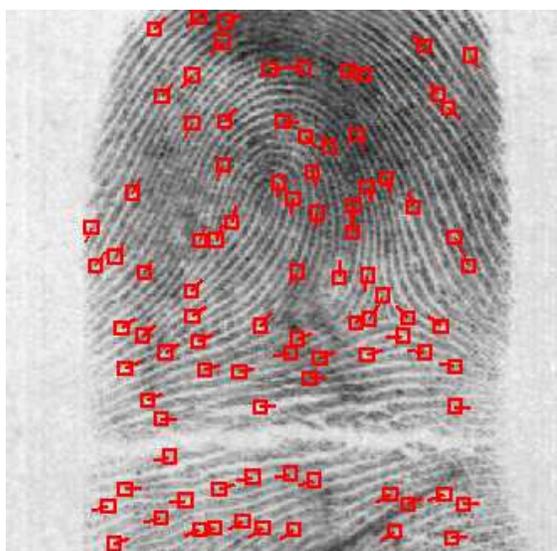
Fig. 3.33 (cont'd)



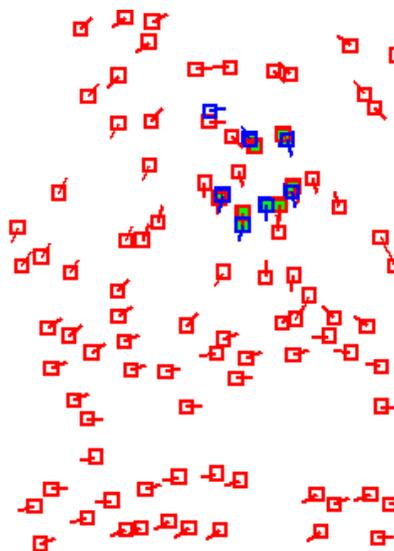
(e) (1, 29408)



(f)



(g) (48, 1)



(h)

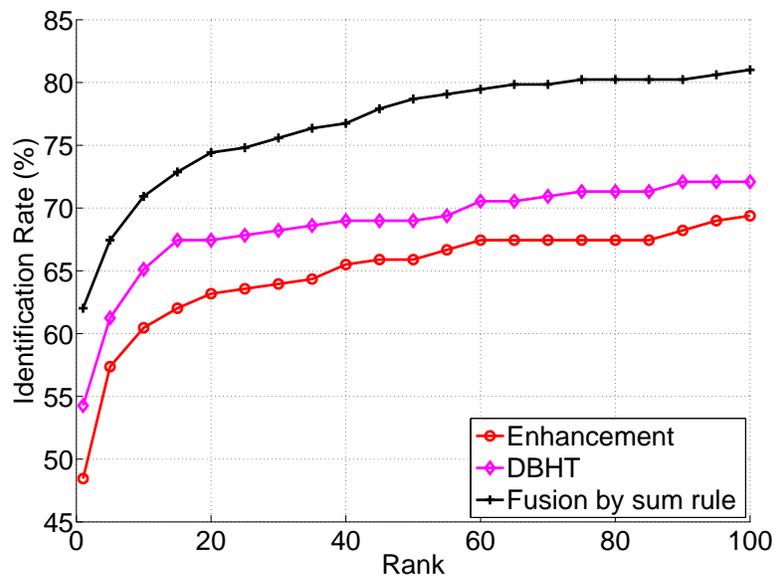


Figure 3.34: Score-level fusion of the two proposed approaches: by enhancement and using descriptor-based Hough Transform (DBHT) matching approach. By combining the two approaches, the rank-100 accuracy is 81%.

3.7 Summary and Conclusions

Latent fingerprint matching is a very challenging problem due to the characteristics of latent fingerprints (poor quality, small area, small number of minutiae, noisy background, etc). While a fully automatic latent matching system is desired, but given the difficulty of the problem and the relatively low matching performance of available AFIS, manual input is still needed. Thus, we have presented two different methods of improving latent fingerprint matching performance in the case where manually marked features are used.

We have shown that the performance of manually marked minutiae in latents can be improved by utilizing automatically extracted minutiae from enhanced latent images. This framework consists of the following steps: (i) reconstruct the orientation field based on manually marked minutiae and singular points; (ii) enhance the latent using median ridge frequency computed in small image blocks and the reconstructed orientation field; (iii) match enhanced latents and rolled fingerprints; (iv) combine the scores from two matchers using boosted max fusion. This framework improved the latent matching performance with a large background database irrespective of latent quality.

We have also presented a fingerprint matching algorithm designed for matching latents to rolled/plain fingerprints which is based on a descriptor-based Hough Transform alignment. A comparison between the alignment performance of the proposed algorithm and the well-known Generalized Hough Transform shows the superior performance of the proposed method. We also reported matching results for two different latent fingerprint databases with a large background database of about 32K rolled prints. We compared the performance of the proposed matcher with three different state-of-the-art fingerprint matchers. Experimental results show that the proposed algorithm performs better than the three baseline fingerprint matchers used in the study across all image qualities. Furthermore, a score-level fusion of the proposed matcher and the best performing commercial matcher shows a further boost in the matching performance.

CHAPTER 4

LATENT FINGERPRINT INDEXING

4.1 Introduction

Law enforcement agencies routinely collect tenprint records of all apprehended criminals in two forms: rolled and plain. Latents – fingerprints lifted from the surfaces of objects at crime scenes – are regarded as an extremely important source of evidence to identify suspects since they can be compared to rolled or plain fingerprints in the background database of known identities. Compared to rolled and plain fingerprints that are obtained in an attended mode (see Figs. 4.1 (a) and (b)), latents typically have poor quality in terms of ridge clarity and complex background noise, and contain only a small part of a finger (friction ridge pattern) and large non-linear skin distortion (see Fig. 4.1 (c)).

Due to these characteristics of latents, the search for the source of a latent is a challenging problem in terms of both efficiency and identification accuracy [106, 50], especially when the background fingerprint database is extremely large. The size of the fingerprint database maintained by just the police department of a typical large city can be of the order of a few millions. The Integrated Automated Fingerprint Identification System (IAFIS) maintained by the Federal Bureau of Investigation (FBI) contains fingerprint records of over 74 million subjects [123] (approximately 740 million images) and this figure keeps growing on a daily basis as more individuals are added to the database. The identification accuracy generally deteriorates as the size of the fingerprint

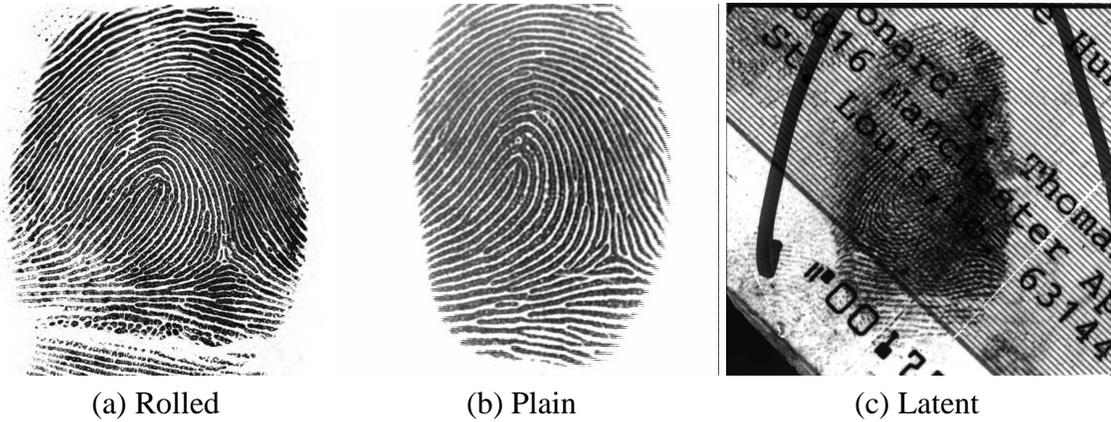


Figure 4.1: Examples of fingerprints. (a) Rolled print (Michigan State Police database), (b) plain or slap print (FVC 2002) and (c) latent print (NIST SD27).

database grows. Under simplifying assumptions, the performance in the identification mode can be estimated based on the FAR and FRR [5]. Assuming the same threshold is used in both the identification and verification scenarios, and that all identities for which the match score is above the aforementioned threshold is returned by the identification system, we have that the False Negative Identification Rate (FNIR)¹ is the same as the FRR. Furthermore, the False Positive Identification Rate (FPIR)² relates to FAR by

$$FPIR = 1 - (1 - FAR)^N \quad (4.1)$$

, where N is the number of subjects in the background database. Thus, when the value of N increases, the FPIR increases and the True Positive Identification Rate (TPIR) decreases, which means the matching accuracy of the identification system is degraded.

This problem of large background database can be alleviated by quickly filtering out a large number of fingerprints from the database that have very low similarity to the search latent before performing detailed one to one matching. Therefore, an effective scheme for fingerprint indexing, also referred to as retrieval or database filtering, is highly desirable to reduce the search space while maintaining or possibly even improving the identification accuracy. Fingerprint indexing is also

¹FNIR is the rate of subjects enrolled in the database who are not identified by the system.

²FPIR is the rate of subjects not enrolled in the database who are mistakenly identified by the system as one of the subjects enrolled in the database.

receiving more attention in recent years due to the limitations of the traditional Henry classification system. The Henry system of fingerprint classification attempts to divide the background database into a fixed number of classes – usually based on the five fingerprint types: right loop, left loop, whorl, arch and tented arch). The major problem with this classification approach lies in the fact that 90% of the fingerprints belong to only three of these five classes [124], thus the background database to be searched is not significantly reduced for most queries. Furthermore, some fingerprints cannot be reliably assigned to only one of these five categories. Indexing is the term used in the fingerprint recognition literature to refer to a significant reduction in the number of fingerprints in the background database to be considered in the finer matching stage [125, 124], or, in other terms, to retrieve a subset of fingerprints in the background database that are most similar to the search fingerprint. Here we use the term indexing to refer to an approach that provides continuous classification³ rather than exclusive classes or categories, while the terms database filtering and fingerprint retrieval can be used in both cases.

As mentioned in Chapter 1, fingerprint friction ridge features are generally described at three different levels. While level 3 features, like pores and ridge contours, are frequently adopted to assist in identification by latent print examiners [126], it is the level 1 (i.e., orientation field, singular point and type) and level 2 (i.e., minutiae) features that are widely used in Automated Fingerprint Identification Systems (AFIS). It is believed that minutiae are the most discriminating and reliable features. However, the information provided by minutiae set alone is limited.

In the case of indexing involving rolled/slap fingerprints, the use of one type of feature alone is usually sufficient to achieve good indexing performance (e.g. [38]), and the use of additional feature types might not significantly increase the indexing performance to justify the added features. However, in the case of latent fingerprints, one feature type alone might not be sufficient in most cases, since the availability and reliability of features in latents vary widely depending on the specific latent characteristics.

³Given the features of a search fingerprint, a distance or similarity to the features of reference prints in the database is computed and most similar fingerprints are retrieved based on the distance or similarity values.

We propose to use both level 1 and level 2 features to improve the indexing performance for latents. Firstly, orientation field in the neighborhood of each minutia is encoded into a rotation- and translation-invariant fixed length bit vector. The bit vectors are then indexed by means of Locality-Sensitive Hashing (LSH) [127]. Then, a conventional minutiae triplet based indexing [31] is boosted by incorporating rotation constraints. Orientation field indexing and triplet indexing are fused with fingerprint indexing technique based on the Minutia Cylinder-Code (MCC) representation [38], and this fusion is further boosted by combining singular points and ridge period filtering. The proposed indexing algorithm is tested by searching 258 latent fingerprints in NIST SD27 against a background database that contains 267,258 rolled fingerprints (including 258 rolled fingerprints in NIST SD27, 27,000 rolled fingerprints in NIST SD14, and 240,000 rolled fingerprints from the Michigan State Police). The experimental results demonstrate that the proposed algorithm outperforms the state-of-the-art indexing methods published in the literature. After filtering out a large fraction of background fingerprints for the finer matching stage, the overall identification accuracy is maintained while the computational efficiency is significantly improved.

The rest of this chapter is organized as follows: Section 4.2 gives a brief review of the related work; Section 4.3 presents feature extraction for both latents and rolled prints; Section 4.4 describes in detail the proposed indexing approach; experimental results are provided in Section 4.5; conclusions are drawn in Section 4.6.

4.2 Related Work

Fingerprint indexing techniques can be roughly categorized into three classes based on the types of features used: minutiae-based, orientation field-based and techniques based on other features (see Table 4.1).

Indexing based on some salient characteristics of the fingerprint (as opposed to ancillary information such as gender, race, age, etc.) is more important in the latent fingerprint case than in the rolled print case because ancillary information is usually not available with the latents. Rolled-to-rolled or plain-to-plain matching is usually performed for background checks or de-duplication.

In this case, a large portion of the database can be filtered out by using the subject's demographic information (e.g. gender, ethnicity, age range, date of birth). In the latent case, this information is usually not available because the latent comes from an unknown subject. In addition, there is a huge gap between the matching performance of rolled-to-rolled matching compared to latent-to-rolled matching due to the partial and noisy nature of the latents; this makes the latent matching problem — and thus the latent indexing problem — much more challenging than the rolled print matching problem.

Among the published works on indexing, only two of them reported the indexing performance on latent fingerprints. Since the focus of our research is latent indexing, we will briefly summarize only these two studies. Feng and Jain [11] proposed a multi-stage filtering scheme whose first stage depends on the fingerprint pattern type, followed by the use of singular points and orientation field. The features are manually marked for the latents, and automatically extracted in the rolled prints. The reported performance on NIST Special Database 27 [4] against a relatively small background of 10,258 rolled prints is 97.3% hit rate at a 39% penetration rate. This means that the true mates of 2.7% of the latents (among the 258 latents in NIST SD 27) were not present in the filtered database (39% of the 10,258 rolled prints).

Yuan *et al.* [10] improved the performance of triplet indexing technique by using triplets to count the number of polygons of various sides. Like [11], they also used manually marked minutiae in the latents and automatically extracted minutiae in the rolled prints. Yuan *et al.* reported an accuracy of 92.7% hit rate at a 40% penetration rate (80.7% at 10% penetration rate) for the NIST SD27 database against a large background database of 240,258 rolled prints. This results [10] could not be verified by us since by applying the algorithm in [10] (code was provided to us by the authors) to our background database of similar size, the hit rate was only around 82% at 40% penetration rate.

Table 4.1: Summary of studies on fingerprint indexing for rolled, plain and latent prints.

	Author(s)	Fingerprint Features	Approach	Fingerprint Database	HR @ PR = 10%
Rolled	Germain <i>et al.</i> [31]	Minutiae	Triples		
	Lumini <i>et al.</i> [124]	OF		1,204 queries and 1,204 templates (NIST SD4)	84%
	Bhanu and Tan [32]	Minutiae	Triples	2,000 queries and 2,000 templates (NIST SD4)	85.5%
	Lee <i>et al.</i> [128]	OF + RF	Histogram	1,204 queries and 1,204 templates (NIST SD4)	88%
	Liu <i>et al.</i> [129]	OF	Complex Filter Responses	2,000 queries and 2,000 templates (NIST SD4)	90%
	Li <i>et al.</i> [130]	OF	Complex Filter Responses	1,204 queries and 1,204 templates (NIST SD4)	96%
	Jiang <i>et al.</i> [33]	OF + RF	OF Clustering	2,000 queries and 2,000 templates (NIST SD4)	89.5%
	Wang <i>et al.</i> [34]	OF	Fingerprint Orientation Model based on 2D Fourier Expansion	2,700 queries and 2,700 templates (last 2,700 pairs of NIST SD14)	98%
	Gyaourova and Ross [131]	Scores	Set of Reference Prints	2,000 queries and 2,000 templates (NIST SD4)	84% @ PR=25%
	Cappelli <i>et al.</i> [38]	Minutiae	MCC	2,700 queries and 24,000 templates (NIST SD14)	95%
	Cappelli [132]	OF + RF		1,000 queries and 1,000,000 templates (generated by SFinGe v4.1)	99.6%
	Cappelli [133]	Minutiae + OF	MCC	2,700 queries and 2,700 templates (last 2,700 pairs of NIST SD14)	98.7%
Liu and Yap [134]	OF	Polar Complex Moments	2,000 queries and 2,000 templates (NIST SD4)	88%	

OF: orientation field, SP: singular points, SIFT: scale invariant feature transform, MCC: minutia cylinder code, RF: ridge frequency, HR: hit rate, PR: penetration rate.

*The images used as templates were randomly selected from each finger.

**Feng and Jain [11] only reported the hit rate at a single penetration rate of 39%.

***The hit rate of algorithm in [10] evaluated on the database used in this paper is 58.1% @ PR=10%.

Table 4.1 (cont'd)

	Author(s)	Fingerprint Features	Approach	Fingerprint Database	HR @ PR = 10%
Plain	Bhanu and Tan [32]	Minutiae	Triples	400 queries and 600 templates (collected by FIU-500-F01 sensor)	100%
	Jiang <i>et al.</i> [33]	OF + RF	OF Clustering	600 queries and 200 templates (FVC2000 DB2a & DB3a)	92.5%
	Liang <i>et al.</i> [135]	Minutiae	Triples	550 queries and 330 templates (FVC2004 DB1*)	99%
	Wang <i>et al.</i> [34]	OF	Fingerprint Orientation Model based on 2D Fourier Expansion	Queries and templates not indicated (FVC2002 DB1a)	99.9%
	Shuai <i>et al.</i> [30]	SIFT		500 queries and 300 templates (FVC2000 DB2a*)	98%
	Gyaourova and Ross [131]	Scores	Set of Reference Prints	2,400 queries and 2,400 templates (WVU)	85%
	Cappelli <i>et al.</i> [38]	Minutiae	MCC	700 queries and 100 templates (FVC2002 DB1a)	99%
	Iloanusi <i>et al.</i> [136]	Minutiae	Quadruplets	400 queries 400 templates (FVC2004 DB1a*)	98%
	Cappelli [132]	OF + RF		500 queries and 300 templates (FVC2002 DB1a*)	99.9%
	Cappelli [133]	Minutiae + OF	MCC	700 queries and 100 templates (FVC2002 DB1a)	100%
	Liu and Yap [134]	OF	Polar Complex Moments	700 queries and 100 templates (FVC2002 DB1a)	85%

Table 4.1 (cont'd)

	Author(s)	Fingerprint Features	Approach	Fingerprint Database	HR @ PR = 10%
Latent	Feng and Jain [11]	Fingerprint Type + SP + OF	Multi-stage filtering	258 latent queries and 10,258 templates (NIST SD27 and NIST SD14)	97.3% @ PR=39% ^{**}
	Yuan <i>et al.</i> [10]	Minutiae	Triplets	258 latent queries and 240,258 templates (NIST SD27 and a private database)	80.7% ^{***}
	Proposed Approach [137]	Minutiae + OF + SP + RF	Triplets + MCC + OF Descriptor Indexing	258 latent queries and 267,258 templates (NIST SD27, NIST SD14 and MSP)	81.8% (95.7% @ PR=39%)

4.3 Feature Extraction

In this section we describe various features that are extracted from reference prints and latents for indexing purposes. More specifically, the following features are extracted: minutiae, orientation field, singular points and ridge period. Note these features are essentially the same that are used in latent matching. The difference between latent indexing and matching is the order in which the features are used and the multistage nature of indexing. In fact, some will argue that an efficient fingerprint matcher(both for rolled prints and latent) has the indexing scheme implicitly built in the matcher itself. However, in this dissertation, we view indexing and matching as two separate modules.

4.3.1 Reference Prints

For reference (rolled) prints in the background database, two commercial off-the-shelf (COTS) matchers are used to automatically extract all the features. Based on our experience with the characteristics and performance of these two matchers, one of the COTS matcher is used to extract the minutiae and the other one is used to extract the skeleton image. Fig. 4.2 shows a reference print from Michigan State Police database, along with the skeleton, orientation field and minutiae. The orientation field and ridge period of a reference print are extracted from the skeleton image as follows:

1. Morphological operations (dilation and erosion) are applied to the skeleton image to get a region of interest (ROI) of the reference print.
2. The skeleton image is filtered with an averaging filter of size 5×5 pixels to convert a binary skeleton image to a grayscale image (see Fig. 4.3).
3. The gradient based algorithm proposed in [8] is used to extract orientation field from the ROI region of the grayscale image (see Fig. 4.3).

4. The X-signature proposed in [115] is used to estimate the ridge period from the ROI region of the grayscale image.

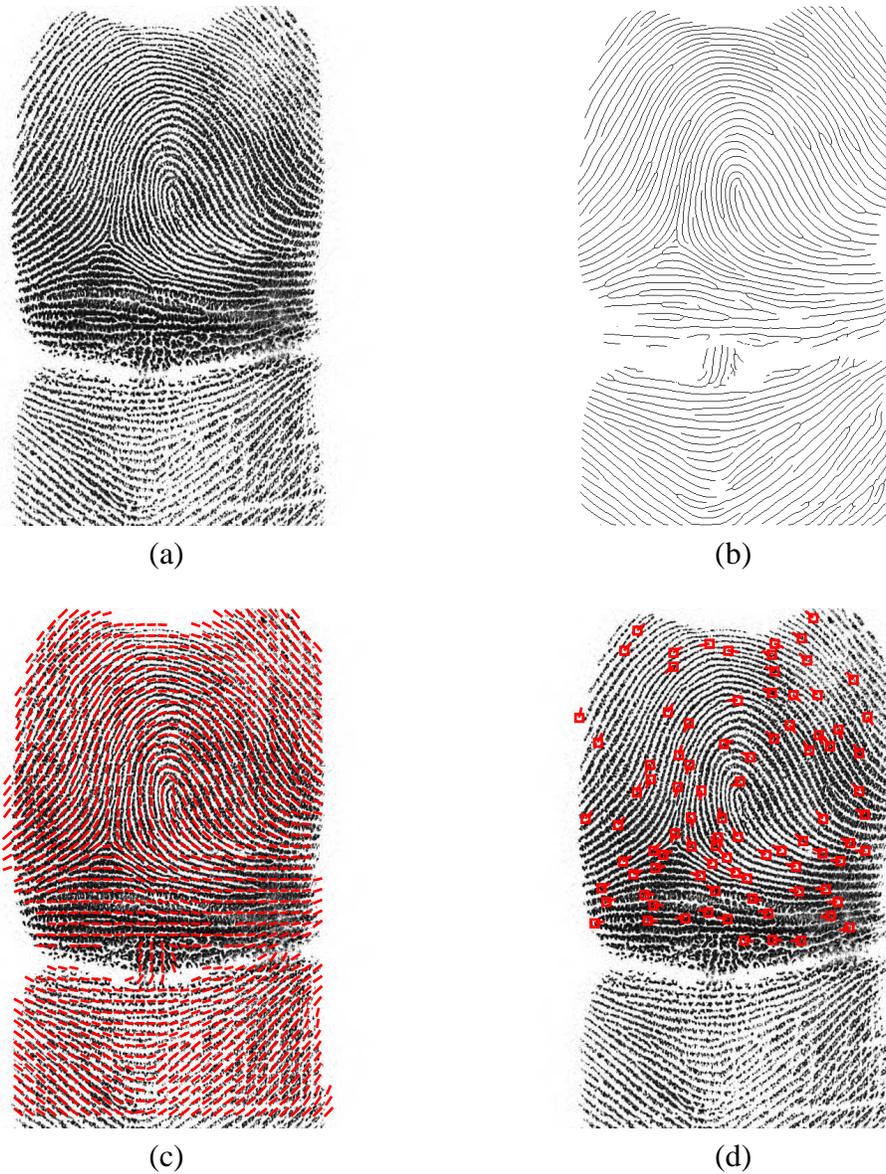


Figure 4.2: (a) Reference print from Michigan State Police database, (b) its skeleton extracted by a commercial matcher, (c) orientation field extracted from skeleton image, and (d) minutiae extracted by a commercial matcher.

From the orientation field, the singular points are extracted, if present in the reference print, by the complex filtering approach proposed in [138]. For both core and delta points, their directions can also be extracted.

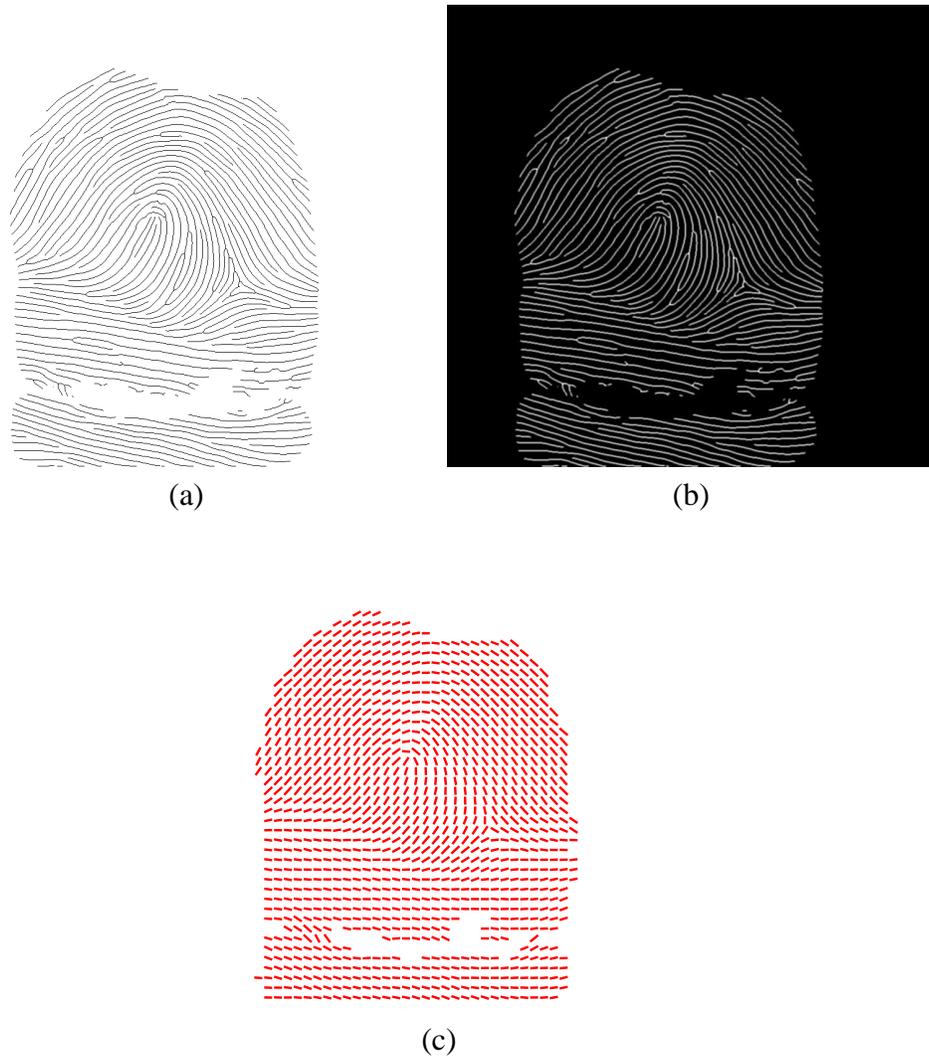


Figure 4.3: (a) Skeleton of a reference print from Michigan State Police database, (b) the grayscale image obtained after applying image filtering and (c) the extracted orientation field.

4.3.2 Latents

Feature extraction in latents is a challenging problem due to heavy background noise. Since the objective of our study is indexing, we assume, similar to the earlier studies in [11] and [10], that minutiae features have been manually marked in the latents by an expert. This is the case for the NIST SD27 database, where manually marked minutiae have been provided along with singular points. The orientation field and ridge period information in latents were marked by the authors in [28] and are available to us. Fig. 4.4 shows some of the manually marked features (orientation

field, singular points and minutiae) available on a latent print from NIST SD27 latent database.

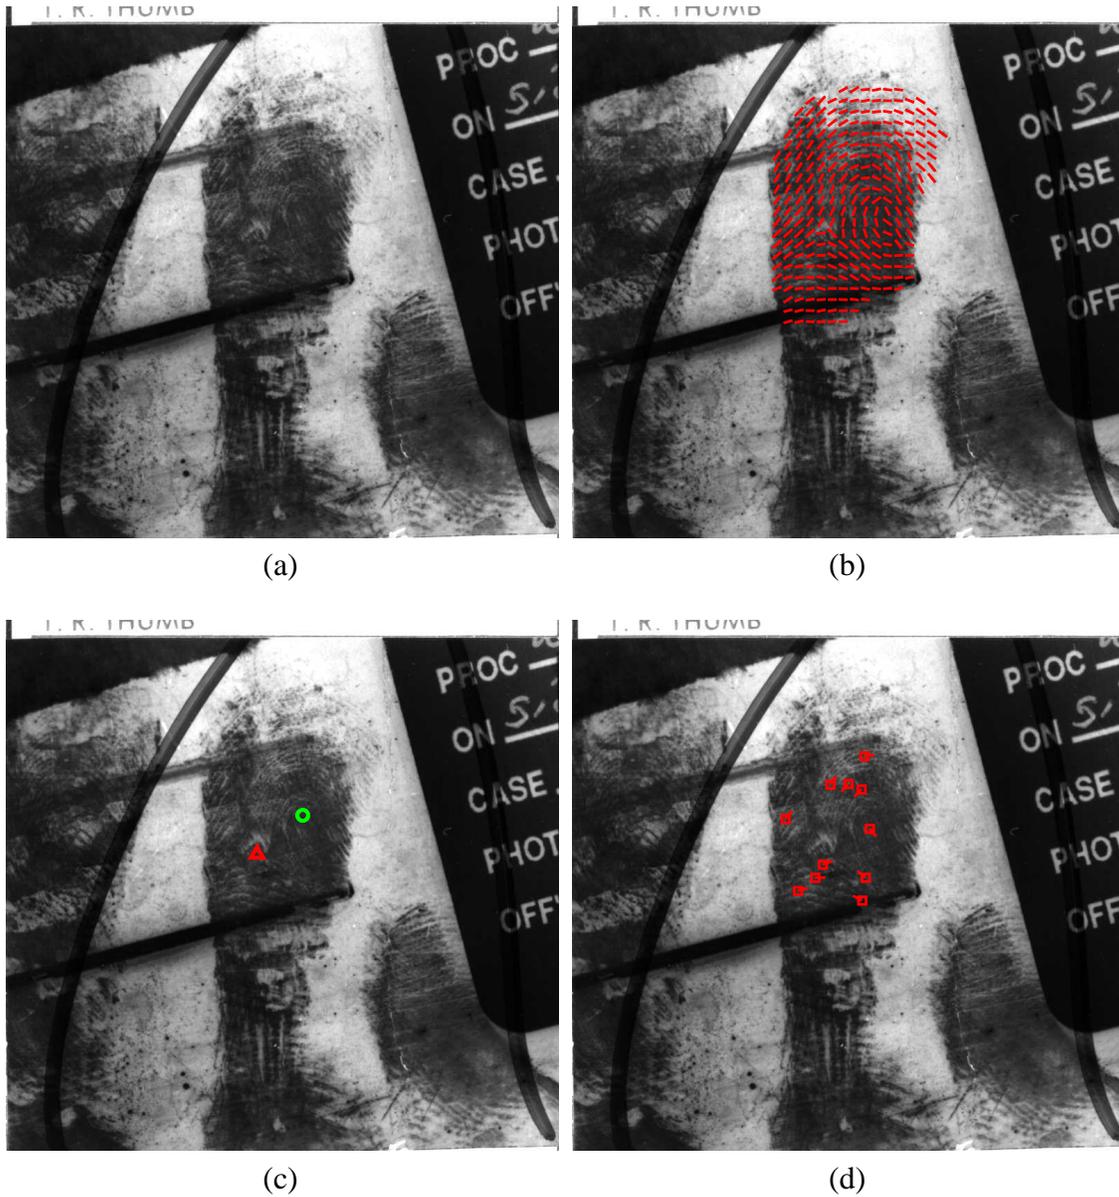


Figure 4.4: (a) Latent print from NIST SD27 latent database, and some manually marked features: (b) orientation field, (c) singular points and (d) minutiae.

4.4 Indexing Approach

We propose to combine different features and indexing techniques to improve latent fingerprint indexing performance. For full (rolled, slap) fingerprints, one feature type alone (e.g. minutia

cylinder code) has been shown to be successful for indexing. But, in the case of latents, availability of features is limited due to small area, distortion, etc. Therefore, the proposed indexing approach using a combination of features is more appropriate for latents.

Our approach consists of combining (i) a constrained version of triplet indexing, (ii) MCC indexing as proposed in [38], (iii) a new orientation field descriptor indexing technique that uses hash function, (iv) filtering based on singular points as proposed in [11], and (v) averaged ridge period comparison. Basic triplets indexing technique is improved by applying a rotation constraint to the matched triplets. Orientation field descriptor indexing is carried out first by converting the descriptor to a binary vector, and then by applying hash functions, similar to the approach proposed in [38]. Each of these specific features used outputs an indexing score, which are then combined to obtain the final indexing score. A description of the techniques used in our approach is presented below, and the overall scheme is shown in Fig. 4.5.

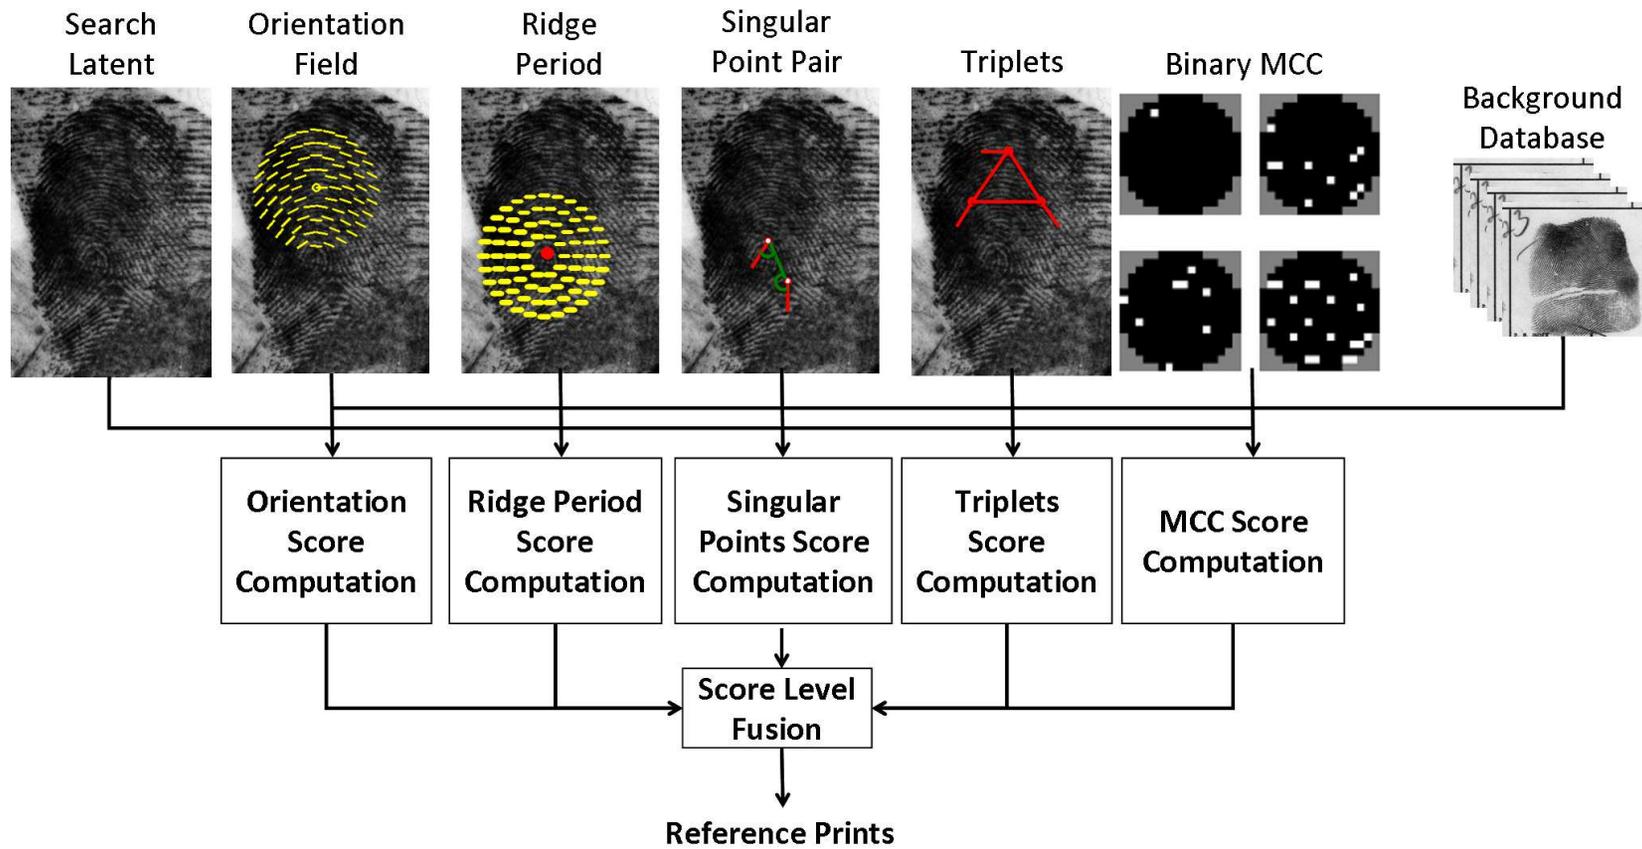


Figure 4.5: Overview of the proposed indexing scheme.

4.4.1 Triplets

Features extracted from the triangles (triplets) formed by any subset of three minutia points have been popular for fingerprint indexing [31]. Several modifications of the first triplet-based algorithm [31] are now available [32, 135, 10]. The basic features in [31] consisted of the length of the sides of the triangles, the ridge count between every pair of minutiae in the triplet, and the angle between minutiae direction and the side of the triangle. The ordering of the sides of the triangle was defined in the clockwise direction.

In our approach, we use the three side lengths of the triangle, and the difference between minutiae direction and one side of the triangle for indexing. Ridge count between minutiae in latents is not sufficiently reliable so we do not use this feature. As proposed in [32], we order the three sides as $L_{max}, L_{min}, L_{med}$ based on their lengths (see Fig. 4.6). Each minutia point $m_i = (x_i, y_i, \beta_j), i = 1, 2, 3$, is associated with a vertex of the triangle by the following rule: P_1 is the point associated with the largest angle, which is opposite to the largest length side, P_2 is associated with the minimum angle in the triangle, which is opposite to the minimum length side, and P_3 is associated with the median angle in the triangle, which is opposite to the median length side, as illustrated in Fig. 4.6. After the ordering of the sides and minutiae, the directional differences ($\theta_i, i = 1, 2, 3$) between each minutiae and one side of the triangle can be consistently obtained.

We also use the handedness of the triangle proposed in [32] as one of the triplet features. Let $P_i = (x_i, y_i)$ be the point corresponding to the i^{th} minutiae in the triplet, $Z_i = x_i + jy_i$ be the complex number associated with $m_i, i = 1, 2, 3, j = \sqrt{-1}$, and $Z_{21} = Z_2 - Z_1, Z_{32} = Z_3 - Z_2, Z_{13} = Z_1 - Z_3$. Then, the triangle handedness ϕ is defined as

$$\phi = \text{sign}(Z_{21} \times Z_{32}), \quad (4.2)$$

where $\text{sign}(\cdot) \in \{-1, 1\}$ is the sign function and \times is the cross product.

In summary, given a set of minutiae points in a latent, the feature vector extracted for every possible triplet t with side length in the range $[l_{min}, l_{max}]$ — in our implementation, 20 and 150

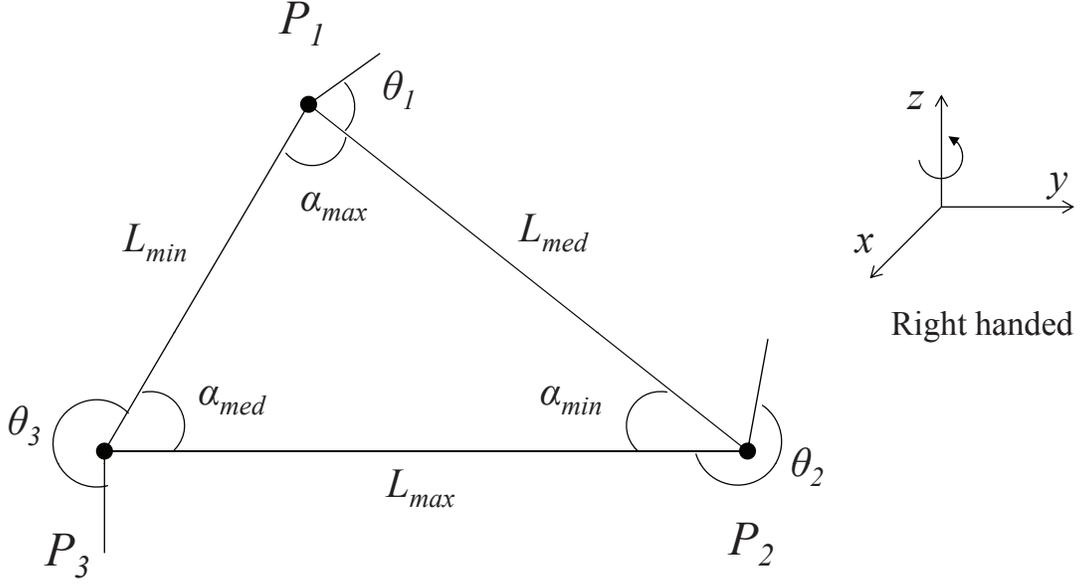


Figure 4.6: Illustration of triplets features.

respectively — is given by

$$F_t = (L_{max}, L_{min}, L_{med}, \theta_1, \theta_2, \theta_3, \phi). \quad (4.3)$$

These triplet features are quantized into bins (bin size is 20 for both triplet side lengths and directional differences; ϕ , which can be either -1 or 1 , is mapped to $\{0, 1\}$). A key is generated based on the possible values of each bin and associated with each triplet. We have that each bin $j, j = 1, \dots, 7$ can assume possible values $v_k, k = 1, \dots, M_j$, and these values are mapped to $1, \dots, M_j, j = 1, \dots, 7$. Assume $b_1 = 1, b_j = b_{j-1}M_{j-1}, j = 2, \dots, 7$ and v_j is the value corresponding to the j^{th} bin for triplet t , then the key to triplet t is computed as

$$K_t = \sum_{j=1}^7 v_j b_j. \quad (4.4)$$

During the indexing or search, triplets in reference prints with the same key as triplets in the latents are considered as matched triplets.

Latent examiners can usually position or align a latent very close to its true position in the finger based on their knowledge of fingerprint structure. Also, rolled prints in background databases usually present only a small rotation compared to the finger upright position. Thus, an additional

step is added to further improve the indexing performance of the triplets by constraining the possible rotation between the two matched triplets. We call this as the constrained triplets approach. The constraint is imposed by first computing the rotation among each pair of matched triplets. Given a matched triplets pair, the differences in the direction of the corresponding three minutiae are averaged and taken as the rotation between that matched triplets pair. If this difference is larger than a threshold (60 degrees in our implementation), then that triplet pair is excluded from the set of matched triplets. The triplet indexing score between the latent and a reference print in the background database is computed as the number of matched triplets between them.

4.4.2 Minutia Cylinder Code (MCC) Indexing

In [38], the authors proposed an indexing technique based on a local minutiae descriptor, called Minutia Cylinder Code (MCC). MCC descriptor [91] represents the neighborhood minutiae information in the form of a 3D function (cylinder). Each level of the cylinder corresponds to a direction difference range, and for each direction difference range, the descriptor contains spatial information about the neighboring minutiae relative to the center minutia (refer to Chapter 3 for details of MCC). MCC descriptor has been successfully used in both rolled-to-rolled fingerprint matching [91] and latent-to-rolled fingerprint matching [88].

Hash functions are utilized to calculate an indexing score for two given fingerprints (search and reference print). The indexing score is computed by averaging the estimate of the Hamming distance between the descriptors, without the need to explicitly perform the distance computations. The estimate of the Hamming distance is based on the collisions between two descriptors under the hash functions so that the indexing is more efficient than directly comparing the descriptors.

More specifically, given two fingerprint templates T_1 and T_2 , with V_1 and V_2 being the sets of cylinders for each minutia in each template, the indexing score proposed in [38] is given by

$$S_M(T_1, T_2) \cong \frac{\sum_{\mathbf{v} \in V_1} (\max_{\mathbf{v}_j \in V_2} \{C_F(\mathbf{v}, \mathbf{v}_j)\})^{\frac{p}{h}}}{|V_1| \times (H_M)^{\frac{p}{h}}}, \quad (4.5)$$

where $|V_1|$ is the number of cylinders in the search print, H_M is the number of hash tables, h is the number of bits selected in each hash table, p is a parameter of the distance function (set to 30), and C_F represents the number of collisions over all hash functions between cylinders v and v_j . We used MCC SDK provided by the authors [38] with the default parameters, with the exception of the number of hash tables that was increased to 64 instead of the default 32 because we found that the former worked better for the latent indexing.

4.4.3 Orientation Field Descriptor Indexing

Given a set of minutiae $M = \{(x_i, y_i, \beta_i)\}_{i=1}^N$ and orientation field $f(x, y)$, where N is the number of minutiae, x_i and y_i are the coordinates of the i^{th} minutia, respectively, and β_i is the direction of minutia, a set of descriptors $O = \{\mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_N\}$ is extracted, where $\mathbf{o}_i = \{o_{i,1}, o_{i,2}, \dots, o_{i,n}\}$ is the descriptor for the i^{th} minutia and n is the total number of sample points in the orientation descriptor [27].

Given a minutia, L concentric circles are centered at the minutia with the j^{th} circle of radius $\frac{R \times j}{L}$, and K_j sample points are equally spaced on each circle starting from the projection location of the reference minutia along its direction on the circle. In this paper, the number of circles L is 4 and the numbers of sample points from inner circle to outer circle are 10, 16, 22 and 28, respectively, the maximum radius R for the outer most circle is 80 pixels. So, the total number of sample points for each orientation descriptor is $n = \sum_{j=1}^L K_j = 76$. The feature value at each sample point is computed as the counter-clock wise rotation angle from local orientation field at the sample point to the local orientation at the central minutia.

In the indexing algorithm, given a seed, a set of l pairs of sample points $\{(h_j, b_j)\}_{j=1}^l$ from the orientation descriptor is randomly generated, where $h_j, b_j \in \{1, 2, \dots, n\}$ and $h_j \neq b_j$. Then, a bit vector $\mathbf{r} = \{r_1, r_2, \dots, r_l\}$, is generated as

$$r_j = \begin{cases} 1, & \text{if } o_{i,h_j} > o_{i,b_j}, \\ 0, & \text{otherwise,} \end{cases} \quad j = 1, 2, \dots, l. \quad (4.6)$$

The random generation of the bit vectors is repeated so that the total number of hash functions is H_O , and the number of bits in each hash function is fixed (l). Each bit vector corresponds to a decimal number, which is then used as one of the index keys. The generation of the orientation field descriptor indexing tables is illustrated in Fig. 4.7. Given two sets of orientation descriptors O_1 and O_2 , similar to the MCC index score computation, the index score for orientation descriptor is computed as

$$S_O(T_1, T_2) \cong \frac{\sum_{\mathbf{r} \in O_1} (\max_{\mathbf{r}_j \in O_2} \{C_F(\mathbf{r}, \mathbf{r}_j)\})^{\frac{p}{l}}}{|O_1| \times (H_O)^{\frac{p}{l}}}, \quad (4.7)$$

where C_F represents the number of collisions over all hash functions as in Equation 4.5.

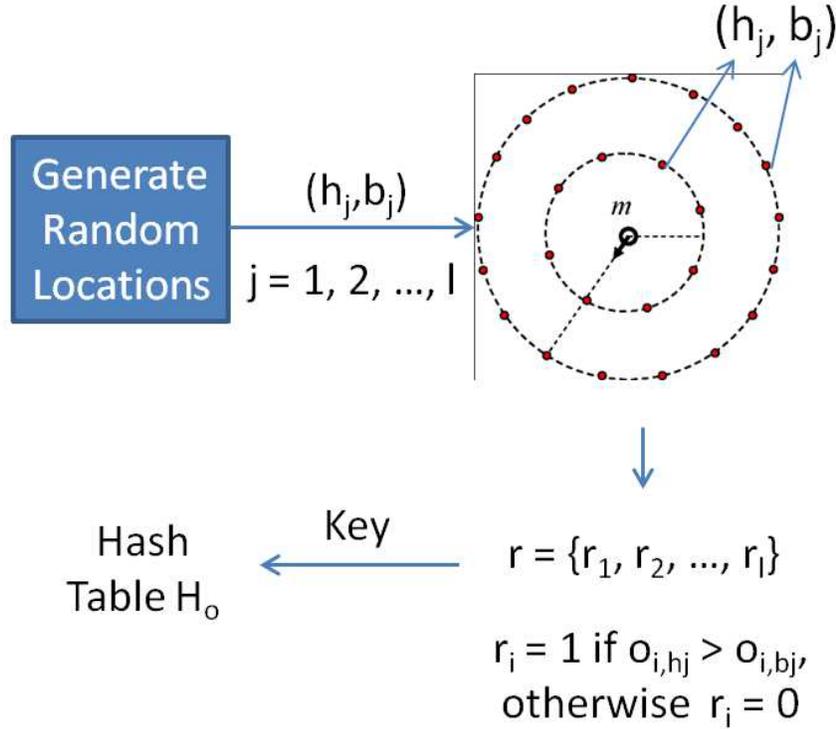


Figure 4.7: Illustration of the key generation in the orientation field descriptor indexing scheme.

4.4.4 Singular Points

Singular points provide useful characterization of a fingerprint. Given a list of core points and delta points with each singular point containing coordinate and direction information, three types

of singular point pair features can be extracted [11]: 1) core pair; 2) delta pair; and 3) core-delta pair.

To order the singular points in a pair of singular points, each core point pair generates two pairs with each core point being ranked first once. For delta pairs, the delta points are ordered based on the x -coordinate, and for core-delta pairs, the core points are ordered first. For each singular pair, three features are computed: 1) the distance (d) between the singular points; 2) the counter-clock wise rotation angle (α) from the direction of the first singular point to the line connecting singular points; 3) the counter-clock wise rotation angle (β) from the direction of the second singular point to the line connecting singular points.

The singular pair similarity score between latent and reference print is defined as

$$S_{SP} = \frac{1}{N_{SP}} \sum_{i=1}^{N_{SP}} \max_j s_{i,j}, \quad (4.8)$$

where N_{SP} is the number of singular pairs in latent and $s_{i,j}$ is the similarity between the i^{th} singular pair in latent and the j^{th} singular pair of the same type in reference print, which can be computed as

$$s_{i,j} = \frac{1}{3} \left(f_{20,60}(\Delta d_{i,j}) + f_{\frac{\pi}{6}, \frac{\pi}{3}}(\Delta \alpha_{i,j}) + f_{\frac{\pi}{6}, \frac{\pi}{3}}(\Delta \beta_{i,j}) \right), \quad (4.9)$$

where $\Delta d_{i,j}$ is the distance difference, $\Delta \alpha_{i,j}$ and $\Delta \beta_{i,j}$ are the angle differences between the first and second singular points, respectively, and $f_{a,b}(x)$ is a piece-wise linear function:

$$f_{a,b}(x) = \begin{cases} 1 & x < a, \\ 0 & x > b, \\ \frac{b-x}{b-a} & \text{otherwise.} \end{cases} \quad (4.10)$$

4.4.5 Ridge Period

Given a fingerprint (either latent or reference print), the average ridge period is computed within a circle of radius 10 blocks centered at the core point, with each block being 16×16 pixels; only foreground blocks are used. If there is no core point in the print, the center point of the foreground

region is used. The ridge period similarity is then computed as

$$S_R = \exp\left(-\frac{|R_l - R_r|}{\sigma_R}\right), \quad (4.11)$$

where R_l and R_r are the average ridge periods in the latent and the reference print, respectively, and σ_R is a normalization term which is usually set to the average of $|R_l - R_r|$. In this paper, σ_R is set to 2.

4.4.6 Fusion of Indexing Scores

The indexing scores based on different features (triplets, MCC, orientation field descriptor, singular points and ridge period) take values in different ranges. Thus, normalization of individual indexing scores is done before the fusion. Let $\mathbf{s} = \{s_1, s_2, \dots, s_{NR}\}$ be the list of scores corresponding to a latent, and N^R be the number of reference prints in the background database. The normalized score $\mathbf{s}' = \{s'_1, s'_2, \dots, s'_{NR}\}$ is obtained as

$$s'_i = \frac{s_i - \min(\mathbf{s})}{\max(\mathbf{s}) - \min(\mathbf{s})}. \quad (4.12)$$

The fusion is based on the saliency of individual features. Triplets, MCC and orientation field descriptor provide similar individual performances; thus, with proper normalization, the simple sum rule is successful [139]. Singular points are not always present in fingerprints, especially in latents that usually contain a partial area. Therefore, the indexing score from the other features is increased based on whether singular points are available in both the prints being compared.

Given the indexing scores from the constrained triplets (S_T), MCC indexing (S_M), and orientation field descriptor (S_O), we obtain locally normalized scores S'_T, S'_M and S'_O . Given the indexing scores from singular points (S_{SP}) and ridge period (S_R), the final indexing score S is computed as

$$S = (S'_T + S'_M + S'_O) \times (1 + S_{SP}) \times S_R. \quad (4.13)$$

The final indexing scores are then used to order the background database such that, for each latent, the top- k reference prints that are most similar to the latent print are retrieved.

4.5 Experimental Results

4.5.1 Databases

Indexing experiments were performed on the only publicly available latent fingerprint database, NIST Special Database 27 [4]. This database consists of 258 latent fingerprint images from operational scenarios (i.e., latents collected from crime scenes) and 258 rolled prints corresponding to the mates of the latent fingerprints. We used 500 ppi resolution for both latents and rolled prints. Features (i.e., minutiae) marked by latent examiners are provided with this database.

The purpose of latent fingerprint indexing is to filter out a large portion of the background database so that the speed of matching stage is enhanced. It is also important to ensure that this gain in matching speed is not at the cost of loss in matching performance. To show the effectiveness of the proposed indexing scheme, we build a large background database of 267,258 rolled fingerprints. This background database includes the 258 mated rolled prints from NIST SD27, 27,000 rolled prints from NIST SD 14 [90], and 240K rolled prints from a database provided by the Michigan State Police (see Fig. 4.8).

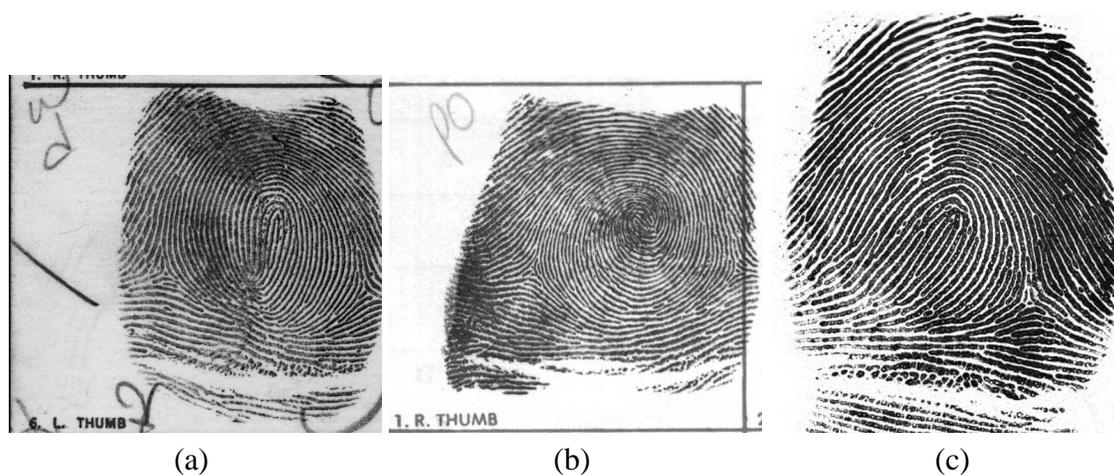


Figure 4.8: Rolled prints in (a) NIST SD27, (b) NIST SD14, and (c) Michigan State Police databases.

4.5.2 Indexing Results

Indexing performance is usually measured by computing the hit rate at various penetration rates. A hit rate at a given penetration rate p refers to the portion of the search prints for which the true mates are retrieved within $p\%$ of the background database (penetration rate). The desirable outcome is high hit rates at low penetration rates and high computational efficiency. Further, it is also desired that the matching accuracy with indexing be no lower than matching accuracy without indexing.

As a baseline, the hit rates at the fixed penetration rate of 20% are 62.0%, 73.6%, and 74.0% for indexing schemes based on basic triplets, MCC indexing and the constrained triplets, respectively. The hit rates at varying penetration rates are shown in Fig. 4.9. Note that for the three methods shown, only minutiae features are used.

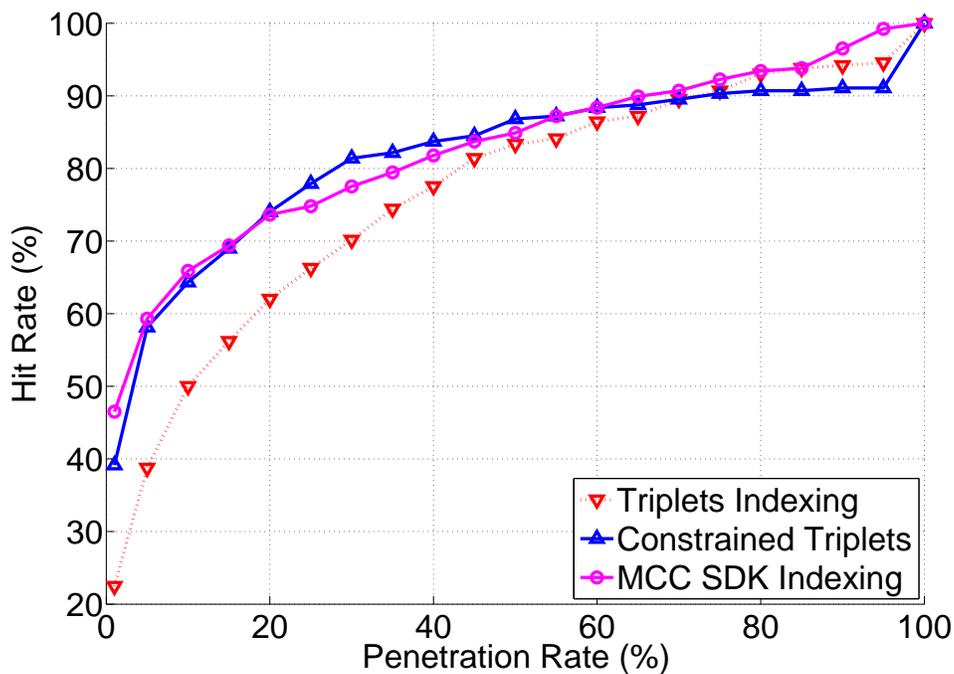


Figure 4.9: Comparison of the indexing performance of basic triplets, constrained triplets and MCC indexing.

Fig. 4.10 shows the performance of the proposed indexing approach along with the performance of two state-of-the-art approaches proposed in [11] and [10]. The results of the algorithm in

[10] shown here are based on an executable file provided by the authors [10], applied to the same database and features as used in our approach. The hit rate at a penetration rate of 39% reported in [11] is shown here for Feng and Jain 2008, which is based on a background database of only 10,258 rolled prints. Note that our approach cannot be directly compared with Feng and Jain’s approach [11] because they used a much smaller background database size compared to our approach and Yuan *et al.*’s approach [10]. However, in our experiments we observed that the differences in background database size do not significantly affect the hit rates (usually about 2-3% difference). Thus, if we fix the penetration rate at 39%, both performances, ours and that reported in [11], are comparable. Furthermore, we will show that the overall latent matching performance is improved by using our approach that filters out 80% of the reference prints, compared to filtering out only 61% in [11].

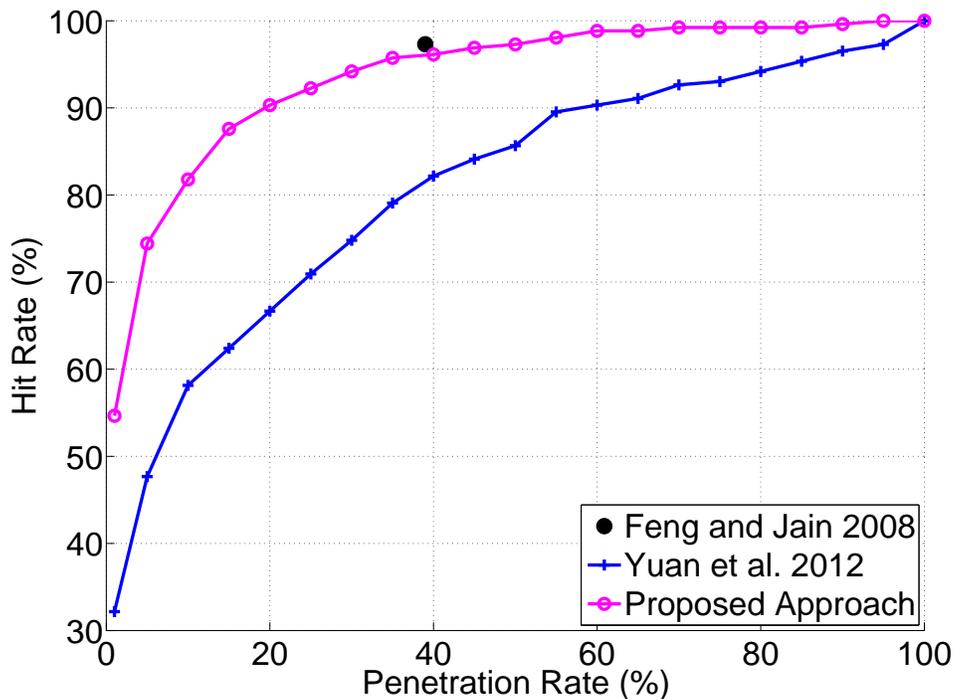


Figure 4.10: Comparison of indexing performance based on the proposed approach against Yuan *et al.* [10] and Feng and Jain [11] on the NIST SD27 latent database. Note that Feng and Jain [11] used only a small background database of 10,258 versus the proposed approach and [10] that use 267,258 reference prints in the background database. Further, Feng and Jain [11] did not provide the complete hit rate vs. Penetration rate curve, but only a single point at a penetration rate of 39% (●).

Given that the proposed approach combines different features and techniques for indexing, we also determine the incremental contributions to hit rates with the addition of different features at a fixed penetration rate (20%). This is shown in Fig. 4.11, where the x -axis denotes the successive utilization of different fingerprint features, with the leftmost feature CTrip being the constrained triplets. It can be observed that the incremental addition of features steadily improves the indexing performance.

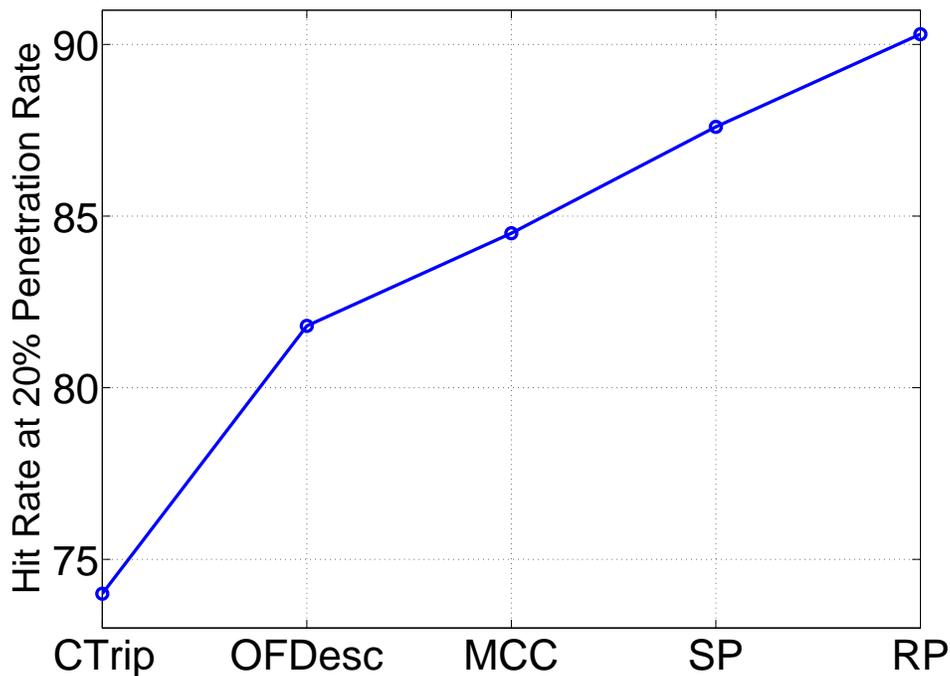


Figure 4.11: Increase in the cumulative hit rate at a fixed penetration rate (20%) as we incrementally add features for indexing (from left to right), where CTrip denotes constrained triplets, OFDesc denotes orientation field descriptor, SP denotes singular points, RP denotes ridge period. The plot shows that we can achieve about 91% hit rate at 20% penetration rate when we use all five features (CTrip, OFDesc, MCC, SP and RP) for indexing.

In order to further evaluate the performance of the proposed indexing algorithm, we also analyzed the influence of the filtering on the latent matching performance. For this purpose, three fingerprint matchers were used: (i) an in-house matcher (DBHT) [88], (ii) a commercial fingerprint matcher (Morpho) and (iii) a commercial latent fingerprint matcher (here referred to as LCOTS).

The input to the two first matchers consisted of manually marked minutiae in the latents and

automatically extracted minutiae in the rolled print, which is the same set of minutiae that was used in the indexing experiments. For the LCOTS, the raw image was provided as input.

For this experiment, for computational efficiency we used a smaller background database of 27,258 rolled prints (258 from NIST SD27 and 27K from NIST SD14). Our indexing approach consisted of not only filtering out a portion of the database, but also of combining both the match scores and indexing scores to improve the matching performance. After retaining a small portion of the background database, we combine the indexing score with the match score by the sum rule. Then, the overall latent matching performance is improved.

For the DBHT and Morpho matchers, the overall latent matching performance on NIST SD27 before and after filtering out 80% of the background database (penetration rate of 20%) is improved, as shown in Fig. 4.12. This demonstrates that our indexing scheme has the desirable property, namely a significant reduction in matching time while improving the matching accuracy. While our indexing is not perfect, most of the errors in the indexing are also errors in the latent matching stage as well. By filtering out 80% of the database, the total computational time for the search is reduced by a factor of 5 for both the in-house matcher and Morpho, while the latent matching performance is improved.

In the case where the indexing is applied, three scenarios are possible regarding the retrieval rank before and after the filtering: (i) the retrieval rank of the true mate stays the same, (ii) the retrieval rank of the true mate improves, or (iii) the true mate is excluded from the background database. The first case occurs when none of the non-mated reference prints with match scores higher than the true mate match score are excluded from the background database. In the second case, the retrieval rank improves because of the exclusion of reference prints with higher match scores to the latent compared to the true mate match score. The third scenario occurs when the true mate is incorrectly filtered out by the indexing module. The first and second cases are a great advantage due to the fact that the search time is greatly reduced while the true mate retrieval rank is maintained or improved. The third scenario might reduce the matching performance, unless the true mate was not going to be retrieved at higher ranks by the fingerprint matcher.

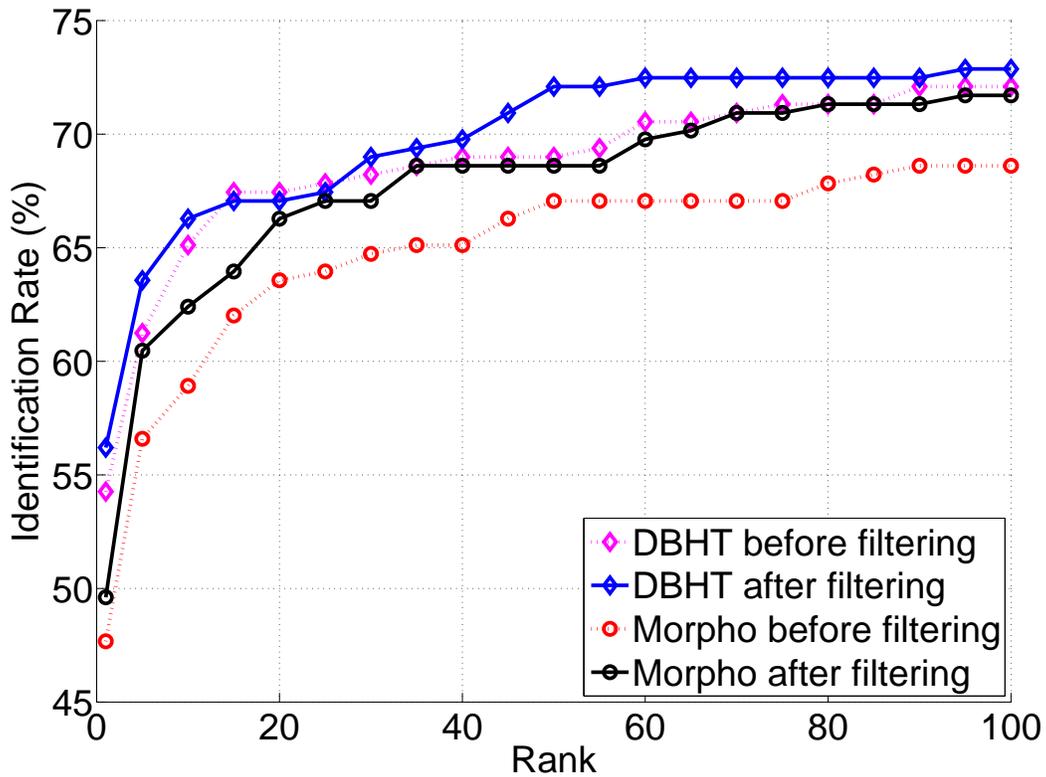
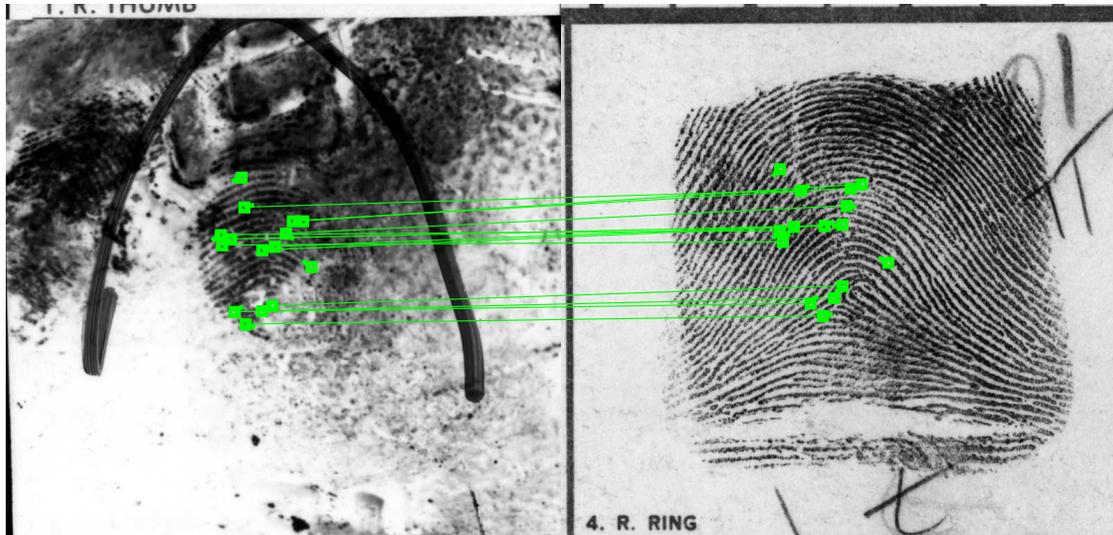


Figure 4.12: Latent matching performance of two different matchers (DBHT in-house latent matcher and Morpho matcher) before and after applying the proposed indexing to filter out 80% of the background database. Note that after the filtering and the fusion of indexing and match scores, the overall matching performance is improved while the matching time is reduced by a factor of five (20% penetration rate).

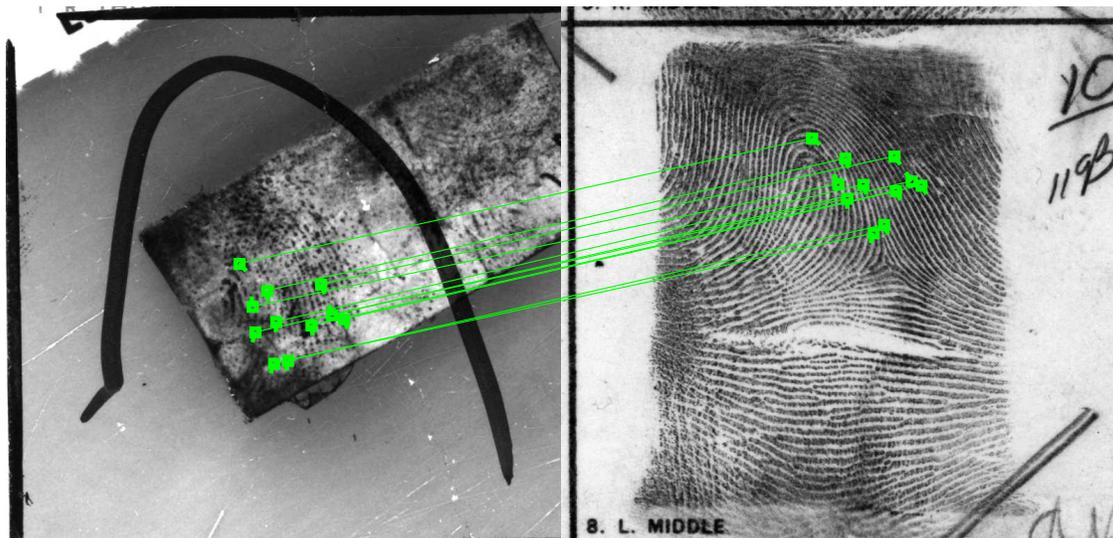
Two latents and their true mates are shown in Fig. 4.13. In both cases, the retrieval rank of the true mate by the in-house matcher was improved after the filtering was applied (3 to 1 and 142 to 49, respectively). In the second case, the true mate was not in the top-100 candidates before the filtering, while it became the top 49 candidate after filtering. Ground truth matched minutiae are shown for reference.

For the commercial latent fingerprint matcher, the overall latent matching performance on NIST SD27 before and after filtering out two thirds of the background database (penetration rate of 33.33%) is improved, as shown in Fig. 4.14. Thus, our proposed indexing scheme improves the matching accuracy of one of the top commercial latent matchers.

In the cases for which the commercial latent matcher provided zero score between the latent



(a)



(b)

Figure 4.13: Examples of latents for which the true mate was correctly retrieved when 20% of the background database was considered are shown. Note that by applying the indexing, the retrieval rank went from (a) 3 to 1 and (b) 142 to 49 due to the exclusion of non-mated reference prints with high match scores to the latent. In (b), the true mate only appeared in the top-100 candidate list after the filtering.

and its true mate, we consider that the true mate is retrieved in the last rank. Fig. 4.15 shows two latents and their true mates for which the true mate could be correctly retrieved at rank 1 after the indexing scheme proposed here was applied and the indexing scores and match scores were combined. The true mates were both ranked last by the commercial latent matcher before the

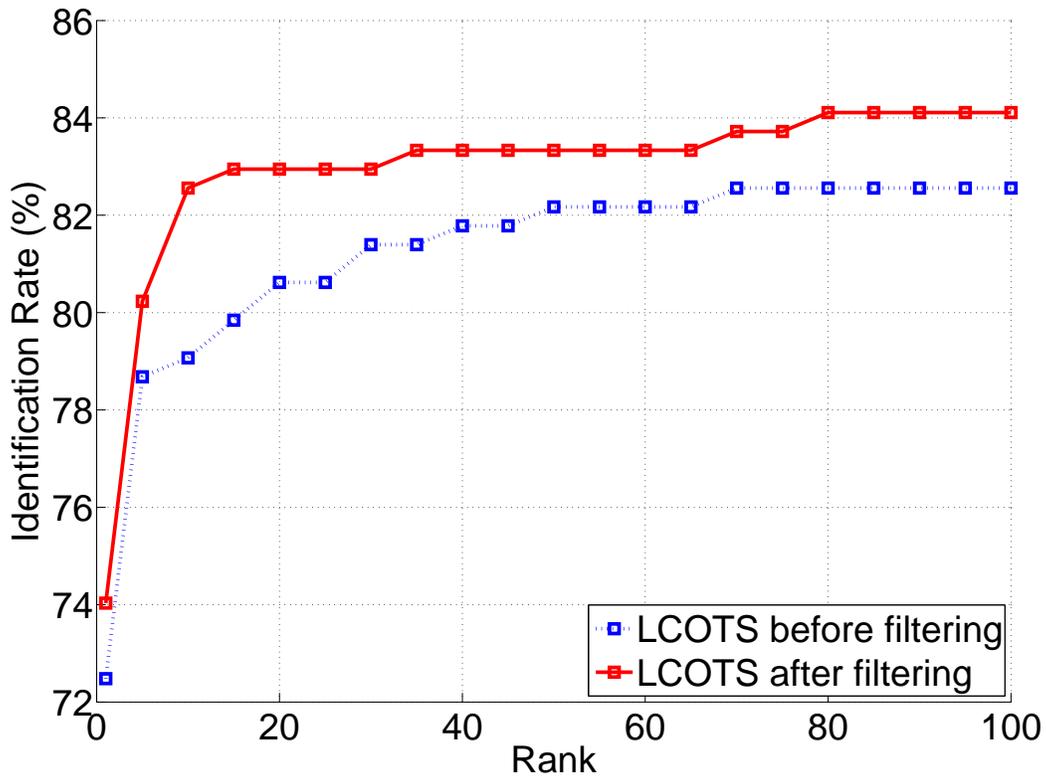
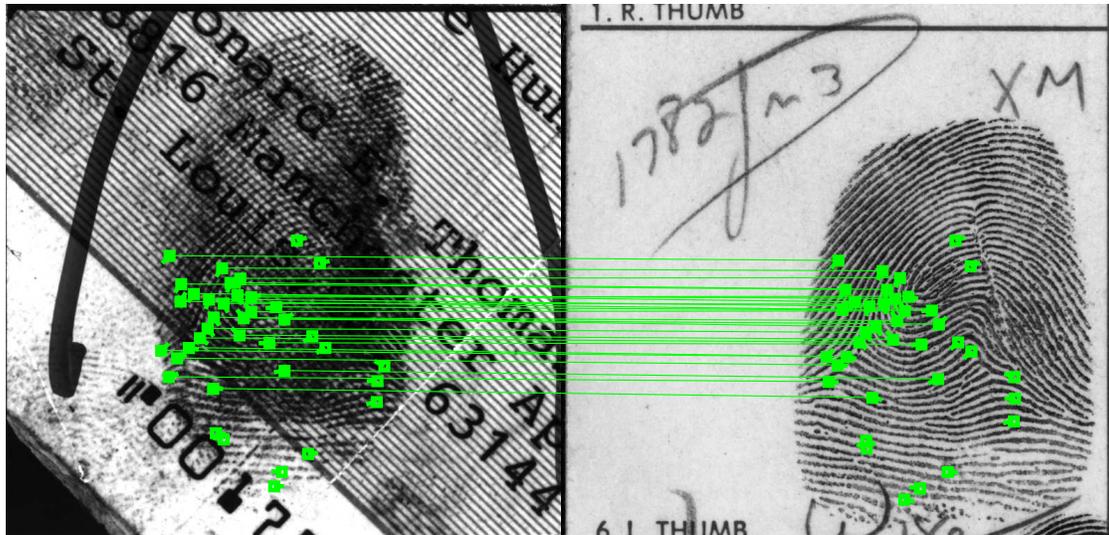


Figure 4.14: Latent matching performance of a commercial latent fingerprint matcher before and after applying the proposed indexing and retaining one third of the background database. The filtering improves the overall matching performance of the commercial latent fingerprint matching while reducing the matching time.

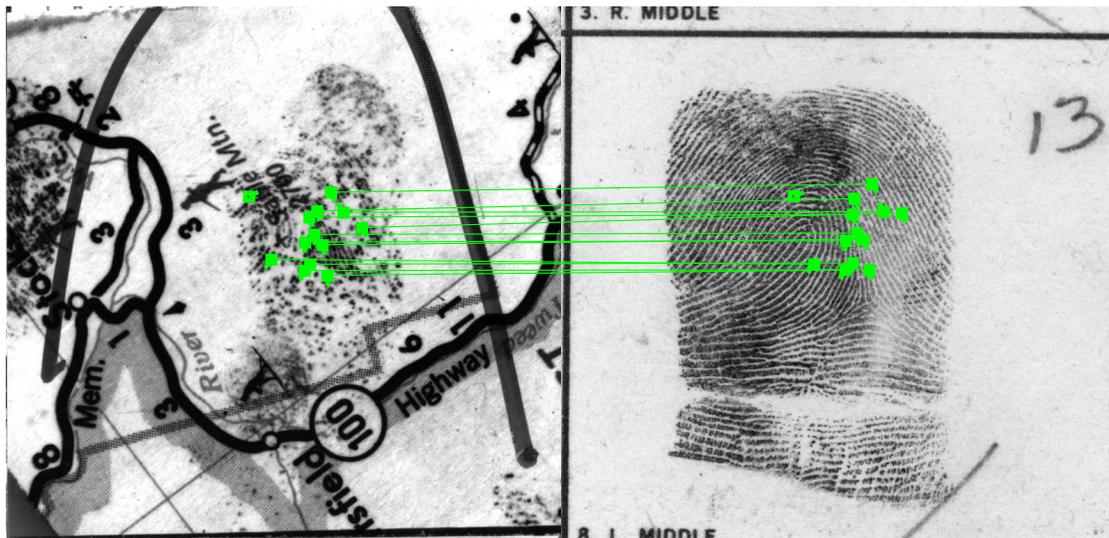
filtering.

In Fig. 4.16, two latents and their true mates are shown. In both cases, the true mates were excluded by the indexing, while they were both ranked first by the commercial latent matcher. The manually marked minutiae and the minutiae extracted by Morpho in the rolled print provide only a small number of corresponding minutiae, therefore the indexing approach is not successful and the true mates are excluded during the filtering.

The proposed approach is more suitable for latent indexing. For indexing of reference prints, the state of the art indexing methods [38] already report very high hit rate at very low penetration rate. For example, the indexing performance of MCC applied on NIST SD4 (the 2,000 second impressions are searched against the 2,000 first impressions and 10,000 reference prints from the



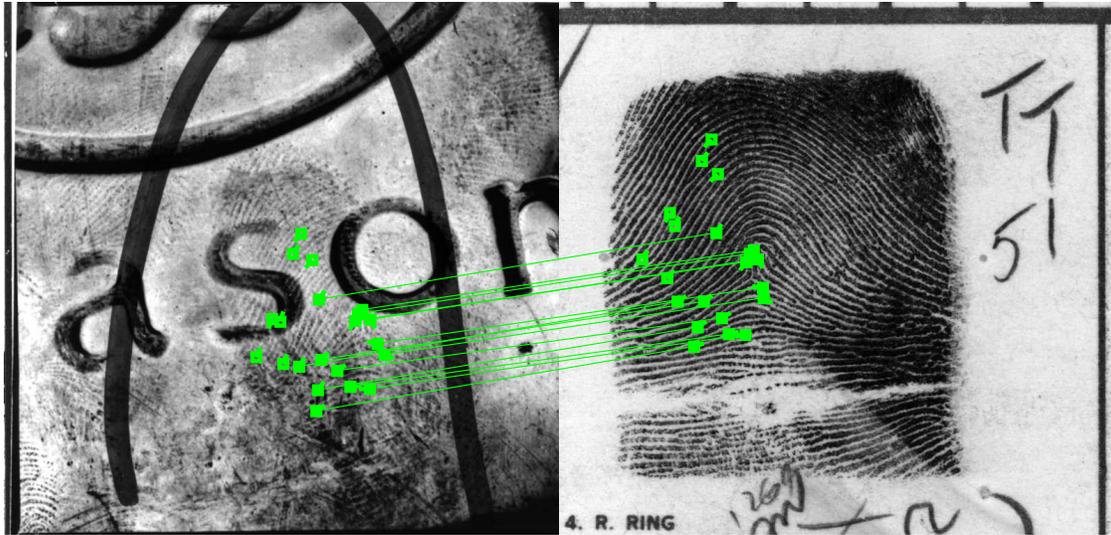
(a)



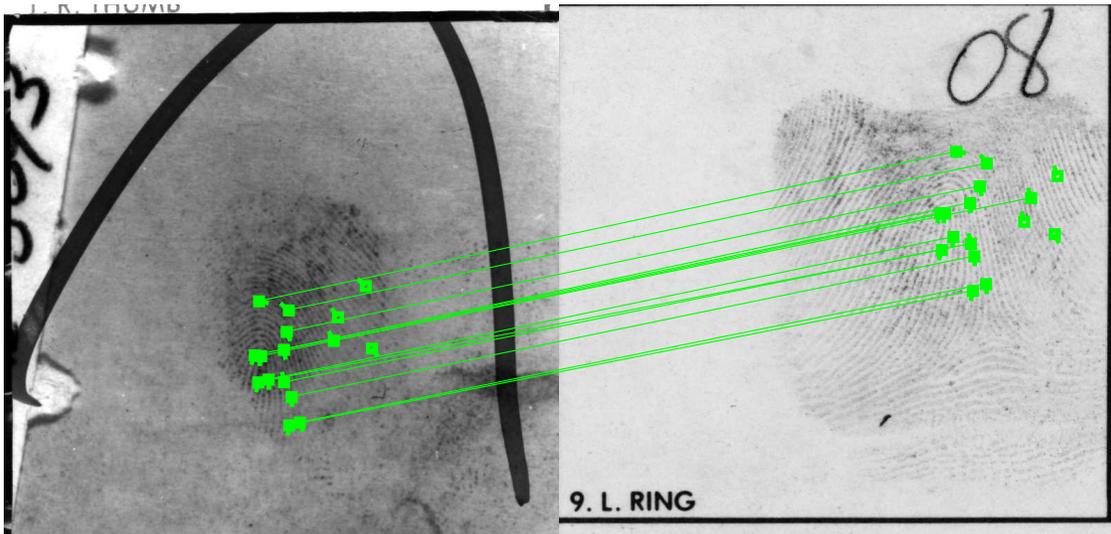
(b)

Figure 4.15: Examples of latents for which the true mate was correctly retrieved when one third of the background database was considered are shown. Note that by applying the indexing, the retrieval rank went from last to 1 due to the fusion of indexing and match scores.

Michigan State Police database) shows a hit rate of 97.8% at 5% penetration rate. Including additional features made only a minor improvement to the indexing performance in this case. Most of the failure cases in the rolled print indexing are due to poor quality of one of the two impressions.



(a)



(b)

Figure 4.16: Examples of latents for which the true mate was not retrieved when one third of the background database was retained. The true mates retrieval ranks before the filtering were 1, so the indexing degrades the matching performance for these two latents.

4.5.3 Implementation Details

Here we report the implementation details of each module used in our indexing scheme. The triplets indexing algorithm, singular point and ridge period estimation modules are implemented in Matlab; MCC SDK (implemented in C#) is provided by the authors of [38] and can be downloaded

from their webpage⁴; the orientation field descriptor-based indexing algorithm is implemented in C++. MCC indexing timing was obtained on a PC (Intel(R) Core(TM)2 Quad CPU Q9650 3.00GHz, 3GB memory), and the timing for all the other modules were obtained on a server with 48 cores (AMD Opteron(tm) Processor 6176, 2.3GHz, 193GB memory); no parallel computing resource was used, thus only one core was used when measuring the computation time.

Table 4.2 reports the computation time for each search module. The modules are all independent and can be run in parallel; in this case, the total time for the indexing would be 28,667 reference prints searched per second, per latent, plus the time spent extracting features from the latent image (less than 3 seconds for all the 258 latents, given the manually marked features minutiae, orientation field, singular points and ridge frequency).

Table 4.2: Computation time for individual indexing modules.

Module	# of reference prints searched per second	Programming language
Basic triplets	39,509	Matlab
Constrained triplets	30,777	Matlab
MCC SDK	58,604	C#
Orientation descriptor	28,667	C++
Singular points	38,444	Matlab
Ridge period	382,960	Matlab

4.6 Summary

We have presented an indexing approach for latent fingerprints that outperforms state-of-the-art indexing techniques on the public domain latent database NIST SD27 against a background database of $\sim 267K$ rolled prints. The proposed approach combines various level 1 and level 2 features, including minutia cylinder code, minutia triplets, singular points, orientation field and ridge period. Improvement in the overall indexing accuracy due to the addition of each individual feature is

⁴<http://biolab.csr.unibo.it/>

shown. The experimental results show that filtering not only improves the computational efficiency of latent search, but the latent matching accuracy is maintained. For the NIST SD 27, the matching accuracy of two different matchers, with and without filtering, remains the same at 20% penetration rate. While the proposed indexing scheme is not perfect, the indexing errors reflected in the hit rate are also reflected in the matching errors.

CHAPTER 5

SUMMARY AND FUTURE WORK

5.1 Summary

This dissertation studied two very challenging problems in biometric recognition: biometric traits of identical twins and latent fingerprint matching. The primary contributions of this dissertation are (i) to provide a better understanding of the similarities and dissimilarities of the biometric traits of identical twins and (ii) to advance latent fingerprint indexing and matching performance.

In Chapter 2, we have presented an analysis of the biometric characteristics of identical twins (fingerprint, face and iris). The discriminability of these three biometric traits is supported by anatomy and the formation process of the biometric characteristics. We have assessed the capacity of state-of-the-art commercial biometric matchers in distinguishing identical twins, as well as their capacity to determine whether two subjects are identical twins or not.

Face biometric system can distinguish two different persons who are not identical twins much better than it can distinguish identical twins. More recent studies support this conclusion, and the face recognition becomes even more challenging when occlusion, expression changes or different lighting conditions are involved.

Fingerprint biometric system also can discriminate two different persons who are not identical twins only slightly better than it can discriminate identical twins, this difference is not as large as for the face biometric system. In the fingerprint matching experiments, the identical twin impostor

distribution is shifted to the right (indicating higher similarity) compared to the general impostor distribution; a statistical hypothesis test shows that the difference between these two distributions is significant.

The iris matching experiment results show that while there is a significant difference in the performance of the biometric system for the identical twin data and for the general population data, the iris biometric system can successfully distinguish identical twins.

Multimodal biometric systems that combine different units of the same biometric modality (e.g., fusion of four fingers or fusion of two irises) lead to an almost perfect separation between genuine and impostor distributions for identical twins.

Biometric traits can also be used to determine whether two subjects are identical twins. By using face and iris modalities together, for example, we can correctly determine 80% of the identical twin pairs as such, while only 2% of subject pairs who are not identical twins will be incorrectly considered identical twins.

In Chapter 3, we have presented two different methods of improving latent fingerprint matching performance. It is assumed, as is the current practice, that manually marked features (minutiae, singular points and ROI) are available for latents.

We have shown that the matching performance based on manually marked minutiae in latents can be improved by utilizing automatically extracted minutiae from enhanced latent images. This framework consists of the following steps: (i) reconstruct the orientation field based on manually marked minutiae and singular points; (ii) enhance the latent using median ridge frequency computed in small image blocks and the reconstructed orientation field; (iii) match minutiae extracted from enhanced latents and rolled fingerprints; (iv) combine the scores from the two matchers using boosted max fusion technique. This framework improved the latent matching performance irrespective of latent quality.

We have also presented a latent fingerprint matching algorithm which is based on a descriptor-based Hough Transform alignment. Experimental results show that the proposed algorithm performs better than the three fingerprint matchers (that includes two COTS matchers) used as baseline

across all latent image qualities. Furthermore, a score-level fusion of the proposed matcher and the best performing COTS matcher shows a further boost in the latent matching performance.

In Chapter 4, we have presented an indexing approach for latent fingerprints that outperforms state-of-the-art indexing techniques on the public domain latent database NIST SD27 against a large background database of $\sim 267\text{K}$ rolled prints. The proposed approach combines a number of level 1 and level 2 features to boost the indexing performance. Contribution to the overall indexing accuracy from each individual feature is shown. Indexing experimental results show that filtering not only improves the computational efficiency of latent search, but also improves the latent matching accuracy when indexing scores are combined with match scores.

5.2 Contributions

In Chapter 2, an analysis of the biometric traits of identical twins provides:

- Capability of state-of-the-art biometric systems to successfully distinguish between identical twins using the three most common biometric traits: fingerprint, face and iris.
- Capability of state-of-the-art biometric systems to successfully determine whether two subjects are identical twins based on their biometric traits (face, fingerprint, iris) alone.

In Chapter 3, we studied the latent matching problem and we presented two different approaches to improve the matching performance given manually marked features (minutiae, singular points, ROI) in the latents. The contributions of this chapter are as follows:

- An approach that combines manually marked features with automatically extracted features from enhanced latents to obtain a better matching performance using a COTS fingerprint matcher.
- Orientation field reconstructed from minutiae along with singular points are used to enhance the latent image to obtain reliable automatically extracted minutiae.

- A fusion scheme that combines match scores based on manually marked minutiae and automatically extracted minutiae from enhanced latents is able to increase the performance consistently over all latent fingerprint qualities in NIST SD27.
- A latent fingerprint matcher is developed using a descriptor-based Hough Transform for alignment (between latent and reference print) and manually marked minutiae from the latents as input.
- The proposed latent matcher advances the matching performance compared to three baseline matchers (that includes two COTS matchers).

In Chapter 4, we presented an indexing approach that combines different features to advance the performance of latent fingerprint indexing. In latents, since there is no guarantee that all the features will be present, a combination of different features for indexing makes more sense than rolled/plain fingerprints. The contributions of the proposed indexing approach are:

- An analysis of the contribution of different features for latent indexing.
- A reduction in the total search time while improving the latent matching accuracy on the NIST SD27 database with a background of 267,258 reference prints.

5.3 Future Work

In Chapter 2, we presented an analysis of the biometric traits of identical twins. Identical twins are the most similar persons in terms of genetics. The next group considering the degree of similarity would be siblings. To extend the study on the similarity of biometrics of identical twins, the use of siblings data would be the next step. For example, can biometrics be used to find siblings of unidentified victims? We used matching scores to derive a measure for identical twin determination. However, additional features need to be investigated for identical twin determination. Some possible features to consider include fingerprint pattern type and number of matched minutiae. It has been shown in the literature that fingerprint pattern type is more correlated between fingerprints

of identical twins and the number of matched minutiae is also slightly higher compared to matched minutiae from general impostor pairs.

In Chapter 3, our purpose was to improve the latent matching performance by developing a matcher based on the availability of a limited amount of reliable features, i.e., manually marked features. Research on extracting reliable features from the latents would be very beneficial to further improve state-of-the-art in latent fingerprint matching.

In Chapter 4, our indexing technique assumes the availability of manually marked features in latents, and it considers all the features in parallel for indexing. A better fusion scheme would take into account the differences in the latents so that the weights assigned to different features used in indexing can be adaptively determined. Further, different features could be used sequentially as opposed to in parallel as implemented here.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar, *Handbook of Fingerprint Recognition*. Springer-Verlag, 2nd ed., 2009.
- [2] NIST Special Database 4, “NIST 8-bit gray scale images of fingerprint image groups FIGS.” <http://www.nist.gov/srd/nist4.cfm>.
- [3] Integrated Pattern Recognition and Biometrics Lab, “West virginia university latent database.” <http://www.cse.msu.edu/~rossarun/i-probe/>.
- [4] NIST Special Database 27, “Fingerprint minutiae from latent and matching tenprint images.” <http://www.nist.gov/srd/nistsd27.cfm>.
- [5] A. K. Jain, A. A. Ross, and K. Nandakumar, *Introduction to Biometrics*. Springer, 2011.
- [6] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [7] P. J. Phillips, P. J. Flynn, K. W. Bowyer, R. W. V. Bruegge, P. J. Grother, G. W. Quinn, and M. Pruitt, “Distinguishing identical twins by face recognition,” in *IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, pp. 185–192, 3 2011.
- [8] A. M. Bazen and S. H. Gerez, “Systematic methods for the computation of the directional fields and singular points of fingerprints,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 24, pp. 905–919, July 2002.
- [9] E. Tabassi, C. L. Wilson, and C. I. Watson, “NIST fingerprint image quality,” tech. rep., NIST Interagency/Internal Report (NISTIR) - 7151, August 2004.
- [10] B. Yuan, F. Su, and A. Cai, “Fingerprint retrieval approach based on novel minutiae triplet features,” in *Biometrics: Theory, Applications and Systems (BTAS)*, pp. 170–175, September 2012.
- [11] J. Feng and A. K. Jain, “Filtering large fingerprint database for latent matching,” in *International Conference on Pattern Recognition (ICPR)*, pp. 1–4, December 2008.
- [12] NSTC Subcommittee on Biometrics, “Biometrics history.” <http://www.biometrics.gov/documents/biohistory.pdf>, August 2006. Accessed November 20, 2012.
- [13] H. T. F. Rhodes, *Alphonse Bertillon, Father of Scientific Detection*. New York: Abelard-Schuman, 1956.

- [14] National Law Enforcement Museum Insider, “Bertillon system of criminal identification.” <http://www.nleomf.org/museum/news/newsletters/online-insider/november-2011/bertillon-system-criminal-identification.html>, November 2011.
- [15] H. Faulds, “On the skin-furrows of the hand,” *Nature*, vol. 22, p. 605, October 1880.
- [16] W. Herschel, “Skin furrows of the hand,” *Nature*, vol. 23, p. 76, November 1880.
- [17] F. Galton, “Personal identification and description,” *Nature*, vol. 38, pp. 201–202, 1888.
- [18] H. Cummins and C. Midlo, *Finger Prints, Palms and Soles: An Introduction to Dermoglyphics*. Dover Publications, Inc., 1961.
- [19] Biography, “Juan Vucetich (18581925).” <https://www.nlm.nih.gov/visibleproofs/galleries/biographies/vucetich.html>, February 2006.
- [20] C. Wilson *et al.*, “Fingerprint vendor technology evaluation 2003: Summary of results and analysis report.” http://biometrics.nist.gov/cs_links/fpvte/report/ir_7123_summary.pdf, June 2004. NISTIR 7123.
- [21] P. Grother, G. W. Quinn, J. R. Matey, M. Ngan, W. Salamon, G. Fiumara, and C. Watson, “IREX III: Performance of iris identification algorithms.” <http://www.nist.gov/itl/iad/ig/irexiii.cfm>, April 2012. NISTIR 7836.
- [22] S. Z. Li and A. K. Jain, eds., *Encyclopedia of Biometrics*, vol. 1. Springer Science + Business Media, 2009.
- [23] E. R. Henry, *Classification and Uses of Finger Prints*. London: Routledge, 1900.
- [24] SWGFAST, “ANSI/NIST ITL 1-2000 American National Standard for Information Systems - data format for the interchange of fingerprint, facial, & scar mark & tattoo (SMT) information.” http://fingerprint.nist.gov/standard/cdeffs/Docs/SWGFAST_Memo.pdf, November 2005.
- [25] A. M. Bazen, G. T. B. Verwaaijen, S. H. Gerez, L. P. J. Veelenturf, and B. J. van der Zwaag, “A correlation-based fingerprint verification system,” in *11th Annual Workshop on Circuits, Systems and Signal Processing ProRISC*, pp. 205–213, November-December 2000.
- [26] T. Hatano, T. Adachi, S. Shigematsu, H. Morimura, S. Onishi, Y. Okazaki, and H. Kyuragi, “A fingerprint verification algorithm using the differential matching rate,” in *International Conference on Pattern Recognition ICPR*, pp. 799–802, 2002.
- [27] M. Tico and P. Kuosmanen, “Fingerprint matching using and orientation-based minutia descriptor,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 25, pp. 1009–1014, August 2003.
- [28] A. K. Jain and J. Feng, “Latent fingerprint matching,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 33, pp. 88–100, January 2011.

- [29] N. K. Ratha, K. Karu, S. Chen, and A. K. Jain, "A real-time matching system for large fingerprint databases," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 18, pp. 799–813, August 1996.
- [30] X. Shuai, C. Zhang, and P. Hao, "Fingerprint indexing based on composite set of reduced SIFT features," in *International Conference on Pattern Recognition (ICPR)*, pp. 1–4, December 2008.
- [31] R. S. Germain, A. Califano, and S. Colville, "Fingerprint matching using transformation parameter clustering," *IEEE Computational Science and Engineering*, vol. 4, no. 4, pp. 42–49, 1997.
- [32] B. Bhanu and X. Tan, "Fingerprint indexing based on novel features of minutiae triplets," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 25, no. 5, pp. 616–622, 2003.
- [33] X. Jiang, M. Liu, and A. C. Kot, "Fingerprint retrieval for identification," *IEEE Trans. on Information Forensics and Security*, vol. 1, pp. 532–542, December 2006.
- [34] Y. Wang, J. Hu, and D. Phillips, "A fingerprint orientation model based on 2D fourier expansion (FOMFE) and its application to singular-point detection and fingerprint indexing," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 29, pp. 573–585, April 2007.
- [35] C. Champod, C. Lennard, P. Margot, and M. Stoilovic, *Fingerprints and Other Ridge Skin Impressions*. CRC Press LLC, 2004.
- [36] R. Cappelli, M. Ferrara, and D. Maltoni, "FVC-onGoing FIDX benchmark." <https://biolab.csr.unibo.it/fvcongoing/UI/Form/ICB2013FIDX.aspx>, 2013. Accessed May 9, 2013.
- [37] "FVC 2002 fingerprint verification competition." <http://bias.csr.unibo.it/fvc2002/>, 2002.
- [38] R. Cappelli, M. Ferrara, and D. Maltoni, "Fingerprint indexing based on minutia cylinder-code," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 33, pp. 1051–1057, May 2011.
- [39] B. Klare, "On a taxonomy of facial features," in *Biometrics: Theory, Applications and Systems (BTAS)*, pp. 1–8, September 2010.
- [40] M. Turk and A. Pentland, "Eigenfaces for recognition," *Journal of Cognitive Neuroscience*, vol. 3, no. 1, pp. 71–86, 1991.
- [41] P. N. Belhumeur, J. P. Hespanha, and D. Kriegman, "Eigenfaces vs. fisherfaces: Recognition using class specific linear projection," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 711–720, 1997.
- [42] NIST, "Face recognition grand challenge." <http://www.nist.gov/itl/iad/ig/frgc.cfm>.

- [43] P. J. Grother, G. W. Quinn, and P. J. Phillips, "Multiple-biometric evaluation (MBE): Report on the evaluation of 2D still-image face recognition algorithms." http://www.nist.gov/customcf/get_pdf.cfm?pub_id=905968, 2010. NISTIR 7709.
- [44] J. G. Daugman, "High confidence visual recognition of persons by a test of statistical independence," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 15, pp. 1148–1161, November 1993.
- [45] "Unique identification authority of india." <https://portal.uidai.gov.in/uidwebportal/dashboard.do#>.
- [46] "Biometrics provides economy and security." <http://www.scribd.com/doc/313998/Jornal-Impressao-Digital,8> 2007.
- [47] S. de Sainte Croix, "Brazil's most secure voting ever." <http://riotimesonline.com/brazil-news/front-page/brazils-most-secure-voting-ever/#>, March 2010.
- [48] P. Kallender, "Japanese banks choose vein-recognition security system." http://www.computerworld.com/s/article/95545/Japanese_banks_choose_vein_recognition_security_system, August 2004.
- [49] "Monozygotic twins." Dorland's Medical Dictionary for Health Consumers, 2007. <http://medical-dictionary.thefreedictionary.com/monozygotic+twins>.
- [50] M. Indovina, V. Dvornychenko, E. Tabassi, G. Quinn, P. Grother, S. Meagher, and M. Garris, "An evaluation of automated latent fingerprint identification technology (Phase II)," tech. rep., NIST Interagency/Internal Report (NISTIR) - 7577, April 2009.
- [51] K. D. Kochanek, S. E. Kirmeyer, J. A. Martin, D. M. Strobino, and B. Guyer, "Annual summary of vital statistics: 2009," *Pediatrics*, vol. 129, pp. 338–348, February 2012.
- [52] J. A. Martin, H.-C. Kung, T. J. Mathews, D. L. Hoyert, D. M. Strobino, B. Guyer, and S. R. Sutton, "Annual summary of vital statistics: 2006," *Pediatrics*, vol. 121, no. 4, pp. 788–801, 2008.
- [53] J. G. Hall, "Twins and twinning," *American Journal of Medical Genetics*, vol. 61, pp. 202–204, 1996.
- [54] J. J. Nora, F. C. Fraser, J. Bear, C. R. Greenberg, D. Patterson, and D. Warburton, *Medical Genetics: Principles and Practice*, ch. Twins and Their Use in Genetics, pp. 395–402. Lea & Febiger, 4 ed., 1993.
- [55] "Identical twins guilty in bank fraud scheme," *United Press International*, September 24, 1986. Section: Domestic News.
- [56] "Twin convicted in third rape trial," March 22nd 2006. <http://origin.foxnews.com/story/0,2933,188801,00.html>.

- [57] Digital Persona, “Unimed case study.” http://www.digitalpersona.com/uploadedFiles/Collateral/Case_Studies/CS_Unimed20080905.pdf, 2008.
- [58] A. K. Jain, S. Prabhakar, and S. Pankanti, “On the similarity of identical twin fingerprints,” *Pattern Recognition*, vol. 35, pp. 2653–2663, 2002.
- [59] A. W.-K. Kong, D. Zhang, and G. Lu, “A study of identical twins’ palmprints for personal verification,” *Pattern Recognition*, vol. 39, pp. 2149–2156, 2006.
- [60] K. Kodate, R. Inaba, E. Watanabe, and T. Kamiya, “Facial recognition by a compact parallel optical correlator,” *Measurement Science and Technology*, vol. 13, pp. 1756–1766, 2002.
- [61] A. Ariyaeinia, C. Morrison, A. Malegaonkar, and S. Black, “A test of the effectiveness of speaker verification for differentiating between identical twins,” *Science and Justice*, vol. 48, pp. 182–186, 2008.
- [62] J. Daugman and C. Downing, “Epigenetic randomness, complexity and singularity of human iris patterns,” in *Proceedings of the Royal Society of London*, vol. 268, pp. 1737–1740, 2001.
- [63] Y. Han, C. Ryu, J. Moon, H. Kim, and H. Choi, “A study on evaluating the uniqueness of fingerprints using statistical analysis,” in *International Conference on Information Security and Cryptology*, pp. 467–477, 2005.
- [64] H. A. Patil and T. K. Basu, “The teager energy based features for identification of identical twins in multi-lingual environment,” in *International Conference on Neural Information Processing*, pp. 333–337, 2004.
- [65] A. M. Bronstein, M. M. Bronstein, and R. Kimmel, “Three-dimensional face recognition,” *International Journal of Computer Vision*, vol. 64, no. 1, pp. 5–30, 2005.
- [66] S. N. Srihari, H. Srinivasan, and G. Fang, “Discriminability of fingerprints of twins,” *Journal of Forensic Identification*, vol. 58, no. 1, pp. 109–127, 2008.
- [67] Z. Sun, A. A. Paulino, J. Feng, Z. Chai, T. Tan, and A. K. Jain, “A study of multibiometric traits of identical twins,” in *Biometric Technology for Human Identification VII*, vol. SPIE 7667, 2010.
- [68] K. Hollingsworth, K. W. Bowyer, S. Lagree, S. P. Fenker, and P. J. Flynn, “Genetically identical irises have texture similarity that is not detected by iris biometrics,” *Computer Vision and Image Understanding*, vol. 115, pp. 1493–1502, 11 2011.
- [69] M. T. Pruitt, J. M. Grant, J. R. Paone, and P. J. Flynn, “Facial recognition of identical twins,” in *International Joint Conference on Biometrics (IJCB)*, pp. 1–8, 10 2011.
- [70] S. Biswas, K. W. Bowyer, and P. J. Flynn, “A study of face recognition of identical twins by humans,” in *International Workshop on Information Forensics and Security (WIFS)*, pp. 1–6, December 2011.

- [71] B. Klare, A. A. Paulino, and A. K. Jain, "Analysis of facial features in identical twins," in *International Joint Conference on Biometrics (IJCB)*, pp. 1–8, October 2011.
- [72] V. Vijayan, K. W. Bowyer, P. J. Flynn, D. Huang, L. Chen, M. Hansen, O. Ocegueda, S. K. Shah, and I. A. Kakadiaris, "Twins 3D face recognition challenge," in *International Joint Conference on Biometrics (IJCB)*, pp. 1–7, 10 2011.
- [73] N. Srinivas, G. Aggarwal, P. Flynn, and R. W. V. Bruegge, "Analysis of facial marks to distinguish between identical twins," *IEEE Trans. on Information Forensics and Security*, vol. 7, pp. 1536–1550, October 2012.
- [74] X. Tao, X. Chen, X. Yang, and J. Tian, "Fingerprint recognition with identical twin fingerprints," *PLoS ONE*, vol. 7, p. e35704, 4 2012.
- [75] J. M. Lytle, "Biometric face scanner tells identical twins apart," May 9th 2008. <http://www.techradar.com/news/world-of-tech/future-tech/face-scanner-tells-identical-\\twins-apart-363150>.
- [76] U. Park and A. K. Jain, "Face matching and retrieval using soft biometrics," *IEEE Trans. on Information Forensics and Security*, vol. 5, pp. 406–415, 9 2010.
- [77] Z. Sun and T. Tan, "Ordinal measures for iris recognition," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 31, no. 12, pp. 2211–2226, 2009.
- [78] "Symwave: BioPrint USB Portable and Peripheral Solutions." http://www.symwave.com/prod/_biometric.shtml.
- [79] Neurotechnology Inc, "Neurotechnology: VeriFinger SDK." <http://www.neurotechnology.com/verifinger.html>.
- [80] A. A. Afifi and S. P. Azen, *Statistical Analysis: A computer oriented approach*. New York: Academic Press, 1972.
- [81] "Cognitec: FaceVACS - SDK." <http://www.cognitec-systems.de/FaceVACS-SDK.19.0.html>.
- [82] "SDK-Beijing IrisKing Ltd. Co." www.irisking.com/en/sdk.html.
- [83] Neurotechnology Inc, "Neurotechnology: VeriEye SDK." <http://www.neurotechnology.com/verieye.html>.
- [84] A. K. Jain, P. Flynn, and A. A. Ross, *Handbook of Biometrics*, ch. 23. Springer, 2008.
- [85] D. R. Ashbaugh, *Quantitative-Qualitative Friction Ridge Analysis*. CRC Press, 1999.
- [86] FBI, "Next generation identification." http://www.fbi.gov/about-us/cjis/fingerprints/_biometrics/ngi.
- [87] A. A. Paulino, A. K. Jain, and J. Feng, "Latent fingerprint matching: Fusion of manually marked and derived minutiae," in *23rd SIBGRAPI - Conference on Graphics, Patterns and Images*, pp. 63–70, August 2010.

- [88] A. A. Paulino, J. Feng, and A. K. Jain, "Latent fingerprint matching using descriptor-based Hough transform," *IEEE Trans. on Information Forensics and Security*, vol. 8, pp. 31–45, January 2013.
- [89] J. Feng and A. K. Jain, "Fingerprint reconstruction: from minutiae to phase," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 33, pp. 209–223, February 2011.
- [90] NIST Special Database 14, "Mated fingerprint card pairs 2." <http://www.nist.gov/srd/nistsd14.cfm>.
- [91] R. Cappelli, M. Ferrara, and D. Maltoni, "Minutia cylinder-code: a new representation and matching technique for fingerprint recognition," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 32, pp. 2128–2141, December 2010.
- [92] J. Qi, S. Yang, and Y. Wang, "Fingerprint matching combining the global orientation field with minutia," *Pattern Recognition Letters*, vol. 26, pp. 2424–2430, 2005.
- [93] J. Gu, J. Zhou, and C. Yang, "Fingerprint recognition by combining global structure and local cues," *IEEE Trans. on Image Processing*, vol. 15, pp. 1952–1964, July 2006.
- [94] X. Jiang and W.-Y. Yau, "Fingerprint minutiae matching based on the local and global structures," in *International Conference on Pattern Recognition (ICPR)*, pp. 1038–1041, 2000.
- [95] X. Chen, J. Tian, and X. Yang, "A new algorithm for distorted fingerprints matching based on normalized fuzzy similarity measure," *IEEE Trans. on Image Processing*, vol. 15, pp. 767–776, March 2006.
- [96] W. Xu, X. Chen, and J. Feng, "A robust fingerprint matching approach: Growing and fusing of local structures," in *International Conference on Biometrics (ICB)*, (Seoul, Korea), pp. 134–143, August 2007.
- [97] J. Feng, "Combining minutiae descriptors for fingerprint matching," *Pattern Recognition*, vol. 41, no. 1, pp. 342–352, 2008.
- [98] J. Feng, S. Yoon, and A. K. Jain, "Latent fingerprint matching: Fusion of rolled and plain fingerprints," in *International Conference on Biometrics (ICB)*, June 2009.
- [99] A. K. Jain, J. Feng, A. Nagar, and K. Nandakumar, "On matching latent fingerprints," in *CVPR Workshops on Biometrics*, pp. 1–8, June 2008.
- [100] S. Yoon, J. Feng, and A. K. Jain, "On latent fingerprint enhancement," in *SPIE*, vol. 7667-766707, April 2010.
- [101] S. Yoon, J. Feng, and A. K. Jain, "Latent fingerprint enhancement via robust orientation field estimation," in *International Joint Conference on Biometrics (IJCB)*, October 2011.
- [102] J. Feng, J. Zhou, and A. K. Jain, "Orientation field estimation for latent fingerprint enhancement," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 35, pp. 925–940, April 2013.

- [103] A. Sankaran, T. I. Dhamecha, M. Vatsa, and R. Singh, “On matching latent to latent fingerprints,” in *International Joint Conference on Biometrics (IJCB)*, pp. 1–6, October 2011.
- [104] M. Vatsa, R. Singh, A. Noore, and K. Morris, “Simultaneous latent fingerprint recognition,” *Applied Soft Computing*, vol. 11, pp. 4260–4266, October 2011.
- [105] NIST, “Evaluation of latent fingerprint technologies.” <http://www.nist.gov/itl/iad/ig/latent.cfm>.
- [106] NIST, “Summary of the results of Phase I ELFT testing.” http://biometrics.nist.gov/cs/_links/latent/elft07/phase1/_aggregate.pdf, September 2007.
- [107] M. Indovina, R. A. Hicklin, and G. I. Kiebusinski, “NIST evaluation of latent fingerprint technologies: Extended feature sets [Evaluation #1],” tech. rep., NIST Interagency/Internal Report (NISTIR) - 7775, April 2011.
- [108] B. T. Ulery, R. A. Hicklin, J. Buscaglia, and M. A. Roberts, “Accuracy and reliability of forensic latent fingerprint decisions,” *Proc. of the National Academy of Sciences of the USA*, vol. 108, no. 19, pp. 7733–7738, 2011.
- [109] B. T. Ulery, R. A. Hicklin, J. Buscaglia, and M. A. Roberts, “Repeatability and reproducibility of decisions by latent fingerprint examiners,” *PLoS ONE*, vol. 7, p. e32800, March 2012.
- [110] I. E. Dror, D. Charlton, and A. E. Peron, “Contextual information renders experts vulnerable to making erroneous identifications,” *Forensic Science International*, vol. 156, no. 1, pp. 74–78, 2006.
- [111] B. G. Sherlock and D. M. Monro, “A model for interpreting fingerprint topology,” *Pattern Recognition*, vol. 26, no. 7, pp. 1047–1055, 1993.
- [112] J. Zhou and J. Gu, “Modeling orientation fields of fingerprints with rational complex functions,” *Pattern Recognition*, vol. 37, pp. 389–391, 2004.
- [113] S. Huckemann, T. Hotz, and A. Munk, “Global models for the orientation field of fingerprints: an approach based on quadratic differentials,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 30, pp. 1507–1519, September 2008.
- [114] J. Feng and J. Zhou, “A performance evaluation of fingerprint minutia descriptors,” in *International Conference on Hand-Based Biometrics*, pp. 1–6, 2011.
- [115] L. Hong, Y. Wan, and A. K. Jain, “Fingerprint enhancement: Algorithm and performance evaluation,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 20, pp. 777–789, August 1998.
- [116] T. K. Ho, J. J. Hull, and S. N. Sihari, “Decision combination in multiple classifier systems,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 16, pp. 66–75, January 1994.
- [117] A. K. Jain, P. Flynn, and A. A. Ross, eds., *Handbook of Biometrics*. Springer, 2008.

- [118] R. Cappelli, M. Ferrara, D. Maltoni, and M. Tistarelli, “MCC: a baseline algorithm for fingerprint verification in FVC-onGoing,” in *11th Int. Conf. Control, Automation, Robotics and Vision*, pp. 19–23, December 2010.
- [119] Biometric System Laboratory team, “FVC-onGoing.” <https://biolab.csr.unibo.it/FVConGoing/UI/Form/Home.aspx>.
- [120] A. A. Paulino, J. Feng, and A. K. Jain, “Latent fingerprint matching using descriptor-based Hough transform,” in *International Joint Conference on Biometrics (IJCB)*, pp. 1–7, October 2011.
- [121] V. N. Dvornychenko, “Evaluation of fusion methods for latent fingerprint matchers,” in *International Conference on Biometrics (ICB)*, March-April 2012.
- [122] A. A. Ross, K. Nandakumar, and A. K. Jain, *Handbook of Multibiometrics*. Springer, 2006.
- [123] The Federal Bureau of Investigation (FBI), “Next generation identification.” http://www.fbi.gov/about-us/cjis/fingerprints_biometrics/ngi.
- [124] A. Lumini, D. Maio, and D. Maltoni, “Continuous versus exclusive classification for fingerprint retrieval,” *Pattern Recognition Letters*, vol. 18, pp. 1027–1034, 1997.
- [125] X. Tan, B. Bhanu, and Y. Lin, “Fingerprint identification: Classification vs. indexing,” in *IEEE Conference on Advanced Video and Signal Based Surveillance, AVSS '03*, (Washington, DC, USA), pp. 151–156, IEEE Computer Society, 2003.
- [126] A. K. Jain, Y. Chen, and M. Demirkus, “Pores and ridges: High-resolution fingerprint matching using level 3 features,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 29, no. 1, pp. 15–27, 2007.
- [127] A. Gionis, P. Indyk, and R. Motwani, “Similarity search in high dimensions via hashing,” in *Proceedings of the 25th International Conference on Very Large Data Bases, VLDB '99*, pp. 518–529, 1999.
- [128] S.-O. Lee, Y.-G. Kim, and G.-T. Park, “A feature map consisting of orientation and inter-ridge spacing for fingerprint retrieval,” in *Audio- and Video-Based Biometric Person Authentication* (T. Kanade, A. Jain, and N. K. Ratha, eds.), vol. 3546 of *Lecture Notes in Computer Science*, pp. 184–190, Springer Berlin Heidelberg, 2005.
- [129] M. Liu, X. Jiang, and A. C. Kot, “Fingerprint retrieval by complex filter responses,” in *International Conference on Pattern Recognition ICPR*, pp. 1042–1045, 2006.
- [130] J. Li, W.-Y. Yau, and H. Wang, “Fingerprint indexing based on symmetrical measurement,” in *International Conference on Pattern Recognition ICPR*, pp. 1038–1041, 2006.
- [131] A. Gyaourova and A. Ross, “A novel coding scheme for indexing fingerprint patterns,” in *Joint IAPR International Workshop on Structural, Syntactic, and Statistical Pattern Recognition*, (Berlin, Heidelberg), pp. 755–764, Springer-Verlag, 2008.

- [132] R. Cappelli, “Fast and accurate fingerprint indexing based on ridge orientation and frequency,” *IEEE Trans. on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 41, pp. 1511–1521, December 2011.
- [133] R. Cappelli and M. Ferrara, “A fingerprint retrieval system based on level 1 and level 2 features,” *Expert Systems with Applications*, vol. 39, pp. 10465–10478, 2012.
- [134] M. Liu and P.-T. Yap, “Invariant representation of orientation fields for fingerprint indexing,” *Pattern Recognition*, vol. 45, pp. 2532–2542, 2012.
- [135] X. Liang, A. Bishnu, and T. Asano, “A robust fingerprint indexing scheme using minugia neighborhood structure and low-order delaunay triangles,” *IEEE Trans. on Information Forensics and Security*, vol. 2, pp. 721–733, December 2007.
- [136] O. Iloanusi, A. Gyaourova, and A. Ross, “Indexing fingerprints using minutiae quadruplets,” in *IEEE CVPR Workshop on Biometrics*, pp. 127–133, June 2011.
- [137] A. A. Paulino, E. Liu, K. Cao, and A. K. Jain, “Latent fingerprint indexing: Fusion of level 1 and level 2 features,” in *Biometrics: Theory, Applications and Systems (BTAS)*, (Washington, D.C.), September 29 - October 2 2013.
- [138] K. Nilsson and J. Bigun, “Localization of corresponding points in fingerprints by complex filtering,” *Pattern Recognition Letters*, vol. 24, no. 13, pp. 2135–2144, 2003.
- [139] K. Nandakumar, Y. Chen, S. C. Dass, and A. K. Jain, “Likelihood ratio-based biometric score fusion,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 30, pp. 342–347, February 2008.