

Advances in Computer Vision and Pattern Recognition



Mark J. Burge
Kevin W. Bowyer *Editors*

Handbook of Iris Recognition

 Springer

Handbook of Iris Recognition

Advances in Computer Vision and Pattern Recognition

For further volumes:

<http://www.springer.com/series/4205>

Mark J. Burge • Kevin W. Bowyer
Editors

Handbook of Iris Recognition

 Springer

Editors

Dr. Mark J. Burge
MITRE Corporation
McLean
Virginia
USA

Prof. Kevin W. Bowyer
Department of Computer Science
and Engineering
University of Notre Dame
Notre Dame
Indiana
USA

Series Editors

Prof. Sameer Singh
Research School of Informatics
Loughborough University
Loughborough
UK

Dr. Sing Bing Kang
Microsoft Research
Microsoft Corporation
Redmond, WA
USA

ISSN 2191-6586

ISBN 978-1-4471-4401-4

DOI 10.1007/978-1-4471-4402-1

Springer London Heidelberg New York Dordrecht

ISSN 2191-6594 (electronic)

ISBN 978-1-4471-4402-1 (eBook)

Library of Congress Control Number: 2012949519

© Springer-Verlag London 2013

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

*Divine Iris, what god sent you to me with a
message?*

– Iliad 18. 167 ff

Foreword

The arrival of this handbook in 2012 suitably marks a number of milestones and anniversaries for iris recognition. The most breathtaking of these is the fact that now on a daily basis more than 100 trillion, or 10-to-the-14th-power, iris comparisons are performed. This juggernaut (a Hindi word, appropriately) was unleashed by the Indian government to check for duplicate identities as the Universal Identification Authority of India, or UIDAI, enrolls the iris patterns of all its 1.2 billion citizens within 3 years. This vastly ambitious programme requires enrolling about 1 million persons every day, across 36,000 stations operated by 83 agencies. Its purpose is to issue each citizen a biometrically provable unique entitlement number (Aadhaar) by which benefits may be claimed, and social inclusion enhanced; thus, the slogan of UIDAI is “to give the poor an identity”. With about 200 million persons enrolled so far, against whom the daily intake of another million must be compared for deduplication, the daily number of iris cross-comparisons is about 10-to-the-14th-power and growing. Similar national projects are also under way in Indonesia and in several smaller countries.

Also breathtaking (but perhaps mainly just for me personally) is the fact that this year is only the 20-year anniversary of the first academic paper proposing an actual method for iris recognition. In August 1992, having recently arrived at Cambridge University as a Research Fellow, I submitted a paper about the method to *IEEE Transactions on Pattern Analysis and Machine Intelligence* (PAMI) entitled “High confidence visual recognition of persons by a test of statistical independence”. The core theoretical idea was that the *failure* of a test of independence could be a very strong basis for pattern recognition, if there is sufficiently high entropy (enough degrees of freedom of random variation) among samples from different classes, as I was able to demonstrate with a set of 592 iris images. The PAMI paper was published in 1993, shortly before my corresponding US Patent 5,291,560 was also issued. That original algorithm was widely licensed through a series of companies (IriScan, Iridian, Sarnoff, Sensor, LG-Iris, Panasonic, Oki, BI2, IrisGuard, Unisys, Sagem, Enschede, Securimetrics and L1 now owned by Safran/Morpho). With various improvements over the years, this algorithm remains today the basis of all significant public deployments of iris recognition. But academic research on many

aspects of this technology has exploded in recent years. To quote from the excellent survey chapter by Bowyer, Hollingsworth and Flynn in this book: During just the 3-year period 2008–2010, there were more papers published about iris recognition than during the entire 15-year period 1992–2007.

The conjecture that perhaps the iris could serve as a fingerprint has a much longer history, and this year marks the 60-year anniversary of the following statement in Adler's classic clinical textbook *Physiology of the Eye* (Chapter VI, page 143): "In fact, the markings of the iris are so distinctive that it has been proposed to use photographs as a means of identification, instead of fingerprints". Apparently, Adler referred to a proposal by the British ophthalmologist Doggart. In the 1980s, two American ophthalmologists, Flom and Safir, managed to patent Adler's and Doggart's conjecture, but they had no actual algorithm or implementation to perform it and so the patent was conjecture. The roots of the conjecture stretch back even further: In 1892, Alphonse Bertillon documented nuances in "Tableau de l'iris humain"; and divination of all sorts of things based on iris patterns goes back to ancient Egypt, Babylonia, and Greece. Iris divination persists today, as "iridology".

Optical systems for iris image acquisition have enjoyed impressive engineering advances, enabling generally a more flexible user interface and a more comfortable distance between camera and subject than the "in-your-face" experience and the "stop-and-stare" interface of the first cameras. Pioneering work by Jim Matey and his team at Sarnoff Labs led to the current generation of systems capturing "iris-at-a-distance" and "iris-on-the-move", in which capture volume is nearly a cubic metre and on-the-move means walking at 1 m/s, enabling throughput rates of a person per second. There has been a "long-distance race" to demonstrate the longest stand-off distance, with some claims extending to the tens of metres. The camera is then essentially a telescope, but the need to project enough radiant light safely onto the target to overcome its inverse square law dilution is a limitation. These developments bring two wry thoughts to my mind: First, I recall that when I originally began giving live demonstrations of iris recognition, the capture volume was perhaps a cubic inch; the hardware was a wooden box containing a video camera, a video display, a near-infrared light source, and a voice interface that replayed the name of a person when visually identified. Second, I read that the Hubble Space Telescope is to be decommissioned, and I wonder whether we might convert it into the Hubble iris camera for the ultimate "iris-at-a-distance" demonstration.

In the first dozen years after the 1993 PAMI paper, it was always very difficult to persuade leaders of the established biometrics community to take an interest in the claim that the iris algorithm had extraordinary resistance against false matches, as well as enormous matching speed. The encoding of an iris pattern into a sign bit sequence enables not only extremely fast XOR matching (e.g. on a 32-bit machine, 32 parallel bits from each of two IrisCodes can be simultaneously compared in a single machine instruction, in almost a single clock cycle at say 3 GHz). But even more important, the Bernoulli nature of random bit pair comparisons generates binomial distributions for the (dis)similarity scores between different eyes. The binomial distribution (for "imposter" comparisons) is dominated by combinatorial

terms with geometric tails that attenuate extremely rapidly. For example, if you accept as a match any IrisCode pair for which no more than 32% of the bits disagree, then the false match likelihood is about one in a million, but if your criterion is just slightly stricter, say that no more than 28% of the bits may disagree, then the false match likelihood is about one in a billion (*i.e.* reduced by a further thousandfold as result of a mere 4-percentile point [0.04] reduction in threshold). These claims became contentious in the year 2000 when the Director of the US “National Biometric Test Center” (NBTC) in San Jose wrote that in their testing of an iris recognition prototype at NBTC, many false matches had been observed. I received copies of all the images, ran all-against-all cross-comparisons, and sure enough, there were many apparent false matches. But when I inspected these putative false match images visually, it became clear that they were all in fact true matches but with changed identities. The Director of the NBTC later confirmed this and generously acknowledged: “Clearly we were getting scammed by some of our student volunteers (at \$25 a head, they were changing names and coming through multiple times)”.

Another obstacle to confirmation of the extreme resistance of this biometric to false matches was the decision in the first large-scale test (ICE 2006: *Iris Challenge Evaluation*) to evaluate at a false match rate of 1 in a thousand ($FMR = 0.001$). In this very nondemanding region of an ROC plot, most biometrics will appear equally powerful. Indeed, since ROC curves converge into the corners at either extreme, if one tested at say $FMR = 0.01$ then probably the length of one’s big toe would seem as discriminating as the iris. The long tradition of face recognition tests had typically used the $FMR = 0.001$ benchmark for obvious reasons: Face recognition cannot perform at more demanding FMR levels. Thus, the ICE 2006 Report drew the extraordinary conclusion that face and iris were equally powerful biometrics. Imagine how well face recognition would hold up in the 100 trillion daily cross-comparisons done by UIDAI. And if iris were operating at the $FMR = 0.001$ level, then every day in UIDAI there would be 100 billion false matches – a number equal to the number of stars in our galaxy or of neurons in the human brain.

A critical feature of iris recognition is that it produces very flat ROC or DET curves. By threshold, adjustment the FMR can be shifted over four or five orders of magnitude while the FnMR hardly changes. Thus, at $FMR = 0.001$, iris may appear unremarkable, as in ICE 2006, and so Newton and Phillips (2007) disputed “the conventional wisdom” that iris was a very powerful biometric. But hardly any price is paid in iris FnMR when its FMR is shifted by several log units, to 0.0000001 or smaller, as required for national-scale deployments. Fortunately, tests by NIST subsequent to ICE have understood this point about the likelihood ratio (the slope of the ROC curve) and have pushed iris testing into the billions of cross-comparisons (IREX-I) and indeed now 1,200 billion cross-comparisons (IREX-III). IREX-I confirmed (7.3.2) that “there is little variation in FnMR across the five decades of FMR” and also confirmed exactly the exponential decline in FMR with minuscule (percentile point) reductions in threshold as I had tabulated in earlier papers. IREX-III results (presented by Patrick Grother in London, October 2011) included a comparison of iris and face performance using the best face algorithms

from 2010 on a database of 1.6 million mugshot face images (compliant with a police mugshot standard) and also 1.6 million DoD detainee iris images. These NIST tests showed that for any plausible FnMR target, iris recognition makes 100,000 times fewer false matches than face.

I am delighted to see the range of topics included in this handbook, which reflects in part the richness of our subject and all the connections it draws among biology, photonics, optical engineering, security engineering, mathematics, algorithms and standardisation. Especially hot current topics include iris image quality metrics, with the recent NIST report (IREX-II or IQCE) on quality-performance covariates and their predictive powers across matchers, and current development of an ISO/IEC Standard (29794-6) for quality. One area that remains rather unexplored is the role of information theory, which lies at the heart of our subject since it measures both the complexity of random variation (the key to biometric collision avoidance) and discriminating power.

Twenty years is a remarkably short time to get from 0 to 100 trillion iris comparisons per day. But also, 20 years is perhaps a generation. It feels as though the real potential of this technology is just beginning to be understood (as can probably also be said about its limitations). This handbook – the first book to be devoted entirely to iris recognition – is full of excellent contributions from a new generation of researchers. If I have been a torchbearer, I am all too happy to “pass the torch” to them while remaining, I hope, still on the field amidst increasing numbers of colleagues captivated by the entropy of the eye.

Cambridge, UK

John Daugman

Preface

Overview and Goals

Iris recognition became a practical area of technology and study with John Daugman's pioneering work about two decades ago. The development of the field was at first slow but has expanded dramatically in recent years. There are now various national identity schemes in progress that make use of iris recognition technology. There is also a large and vibrant research community focused on iris recognition, studying ways to make it even more accurate in even larger-scale applications. The primary goal of this book is to give an authoritative introduction to the current state of the art in iris recognition technology. The field has already in large part moved past the study of alternative segmentation algorithms and texture filters applied to pristine iris images. One major current emphasis is how to deal with varying quality iris images acquired with less explicit user cooperation. Another major current emphasis is on methods for improving accuracy in the context of varying quality images. Still another major current emphasis is on a better understanding of the basic science underlying iris recognition. Each of these emphases is represented by multiple chapters in this book.

Organisation and Features

This book includes a foreword by Professor John Daugman, along with a collection of nineteen chapters contributed by researchers from around the world. It includes theoretical studies, such as the chapter by Clark, Culp, Herron and Ross on iris dynamics and the chapter by Kong, Zhang and Kamel on the IrisCode. It also includes very empirical studies, such as the chapter by Baker, Bowyer, Flynn and Phillips on iris template aging and the chapter by Phillips and Flynn analysing results from the Iris Challenge Evaluation 2006. The 44 authors contributing to the

book come from companies, government agencies and universities. They also come from many different countries, including Lithuania, Canada, Singapore, Denmark, Portugal, Hong Kong, the United Kingdom and the USA.

Target Audiences

The target audience for this book is anyone who wants a better understanding of the current state of the art in iris recognition. Practitioners in industry should find new insights and possibilities in the breadth of topics covered. Managers and executives in government should find a more sober appraisal of the field than exists in the marketing literature of the industry. Researchers in government, industry and academia should find new ideas for productive research efforts.

Acknowledgements

We want to thank the editors at Springer for their patience and advice during the development of this project. We also want to thank all of the contributors to this book for their prompt replies on various points. We want to thank all of our collaborators at our respective institutions for the vibrant research atmosphere that they have provided. Finally, we want to thank our families; for without their continual support and encouragement, this book would not have been possible.

Washington, DC, USA
Notre Dame, IN, USA

Mark J. Burge
Kevin W. Bowyer

Contents

1	Introduction to the Handbook of Iris Recognition	1
	Kevin W. Bowyer and Mark J. Burge	
2	A Survey of Iris Biometrics Research: 2008–2010	15
	Kevin W. Bowyer, Karen P. Hollingsworth, and Patrick J. Flynn	
3	Standard Iris Storage Formats	55
	George Quinn, Patrick Grother, and Elham Tabassi	
4	Iris Quality Metrics for Adaptive Authentication	67
	Natalia A. Schmid, Jinyu Zuo, Francesco Nicolo, and Harry Wechsler	
5	Quality and Demographic Investigation of ICE 2006	85
	P. Jonathon Phillips and Patrick J. Flynn	
6	Iris Recognition with Taylor Expansion Features	103
	Algirdas Bastys, Justas Kranauskas, and Volker Krüger	
7	A Theoretical Model for Describing Iris Dynamics	129
	Antwan D. Clark, Scott A. Kulp, Isom H. Herron, and A.A. Ross	
8	Iris Recognition in the Visible Wavelength	151
	Hugo Proença	
9	Multispectral Iris Fusion and Cross-Spectrum Matching	171
	Mark J. Burge and Matthew Monaco	
10	Robust and Secure Iris Recognition	183
	Jaishanker K. Pillai, Vishal M. Patel, Rama Chellappa, and Nalini K. Ratha	

11	Template Aging in Iris Biometrics	205
	Sarah E. Baker, Kevin W. Bowyer, Patrick J. Flynn, and P. Jonathon Phillips	
12	Fusion of Face and Iris Biometrics	219
	Ryan Connaughton, Kevin W. Bowyer, and Patrick J. Flynn	
13	Methods for Iris Segmentation	239
	Raghavender Jillela and Arun A. Ross	
14	Iris Segmentation for Challenging Periocular Images	281
	Raghavender Jillela, Arun A. Ross, Vishnu Naresh Boddeti, B.V.K. Vijaya Kumar, Xiaofei Hu, Robert Plemmons, and Paúl Pauca	
15	Periocular Recognition from Low-Quality Iris Images	309
	Josh Klontz and Mark J. Burge	
16	An Introduction to the IrisCode Theory	321
	Adams Wai Kin Kong, David Zhang, and Mohamed Kamel	
17	Application of Correlation Filters for Iris Recognition	337
	B.V.K. Kumar, Jason Thornton, Marios Savvides, Vishnu Naresh Boddeti, and Jonathon M. Smereka	
18	Iris Spoofing: Reverse Engineering the Daugman Feature Encoding Scheme	355
	Shreyas Venugopalan and Marios Savvides	
19	Optics of Iris Imaging Systems	367
	David A. Ackerman	
	Afterword	395
	Index	399

List of Contributors

- David A. Ackerman** SRI Sarnoff, Princeton, NJ, USA
- Sarah E. Baker** University of Notre Dame, Notre Dame, IN, USA
- Algirdas Bastys** Vilnius University, Vilnius, Lithuania
- Vishnu Naresh Boddeti** Carnegie Mellon University, Pittsburgh, PA, USA
- Kevin W. Bowyer** University of Notre Dame, Notre Dame, IN, USA
- Mark J. Burge** The MITRE Corporation, McLean, VA, USA
- Rama Chellappa** University of Maryland, College Park, MD, USA
- Antwan D. Clark** University of Maryland, College Park, MD, USA
- Ryan Connaughton** University of Notre Dame, Notre Dame, IN, USA
- John Daugman** Computer Laboratory, University of Cambridge, Cambridge, UK
- Patrick J. Flynn** University of Notre Dame, Notre Dame, IN, USA
- Patrick Grother** National Institute of Standards and Technology, Gaithersburg, MD, USA
- Isom H. Herron** Rensselaer Polytechnic Institute, Troy, NY, USA
- Karen P. Hollingsworth** University of Notre Dame, IN, USA
- Xiaofei Hu** Wake Forest University, Winston-Salem, NC, USA
- Raghavender Jillela** West Virginia University, Morgantown, WV, USA
- Mohamed Kamel** University of Waterloo, Waterloo, ON, Canada
- Josh Klontz** The MITRE Corporation, McLean, VA, USA
- Adams Wai Kin Kong** Nanyang Technological University, Singapore, Singapore
- Justas Kranauskas** Vilnius University, Vilnius, Lithuania

- Volker Krüger** Aalborg University, Aalborg, Denmark
- Scott A. Kulp** Rutgers University, Piscataway, NJ, USA
- Matthew Monaco** Noblis, Falls Church, VA, USA
- Francesco Nicolo** West Virginia University, Morgantown, WV, USA
- Paúl Pauca** Wake Forest University, Winston-Salem, NC, USA
- Vishal M. Patel** University of Maryland, College Park, MD, USA
- P. Jonathon Phillips** National Institute of Standards and Technology, Gaithersburg, MD, USA
- Robert Plemmons** Wake Forest University, Winston-Salem, NC, USA
- Jaishanker K. Pillai** University of Maryland, College Park, MD, USA
- Hugo Proença** University of Beira Interior, IT-Instituto de Telecomunicações, Covilhã, Portugal
- George Quinn** National Institute of Standards and Technology, Gaithersburg, MD, USA
- Nalini K. Ratha** IBM Watson Research Center, Hawthorne, NY, USA
- Arun A. Ross** West Virginia University, Morgantown, WV, USA
- Marios Savvides** Carnegie Mellon University, Pittsburgh, PA, USA
- Natalia A. Schmid** West Virginia University, Morgantown, WV, USA
- Jonathon M. Smereka** Carnegie Mellon University, Pittsburgh, PA, USA
- Elham Tabassi** National Institute of Standards and Technology, Gaithersburg, MD, USA
- Jason Thornton** MIT Lincoln Laboratory, Lexington, MA, USA
- Shreyas Venugopalan** Carnegie Mellon University, Pittsburgh, PA, USA
- B.V.K. Vijaya Kumar** Carnegie Mellon University, Pittsburgh, PA, USA
- Harry Wechsler** George Mason University, Fairfax, VA, USA
- David Zhang** The Hong Kong Polytechnic University Kowloon, Kowloon, Hong Kong
- Jinyu Zuo** West Virginia University, Morgantown, WV, USA

Chapter 1

Introduction to the Handbook of Iris Recognition

Kevin W. Bowyer and Mark J. Burge

Abstract Iris biometrics is both a technology already in use in ambitious nation-scale applications and an active research area with many exciting unsolved problems. This chapter gives brief introduction to the area and an overview of the contents of the rest of the book.

1.1 Introduction

The use of iris texture analysis for biometric identification is clearly well established. The United Arab Emirates has been using iris biometrics for border control for the past decade [2] and famously claims “some 2.7 billion iris cross-comparisons being done every day” [7]. India is using iris as part of its Unique ID program, which recently passed 300 million persons enrolled [1]. Airports in London, Amsterdam, and elsewhere continue to use iris biometrics to speed people through border control. The various high-profile uses of iris biometrics, along with the fact there are now multiple competing commercial technologies, might lead some casual observers to think that iris biometrics is a “solved problem” and no longer an active research area. This would be quite wrong. The transition of iris biometrics from laboratory technology to nation-scale applications has actually brought about the recognition of many interesting and difficult research issues. The purpose of this book is to give an introduction to the current state of the art in iris biometrics, as well as some of the important current areas of active research.

K.W. Bowyer (✉)
University of Notre Dame, Notre Dame, IN, USA
e-mail: kwb@cse.nd.edu

M.J. Burge
The MITRE Corporation, McLean, VA, USA
e-mail: burge@ieee.org



Fig. 1.1 Example iris image exhibiting texture from cosmetic contact lens. Note slivers of natural iris texture by sclera boundary near the nose and by pupil boundary away from the nose

There are many current open research issues in iris biometrics. One cluster of such issues is centered around contact lenses. Cosmetic contact lenses present challenges that are more complex than is widely appreciated, and even plain, clear, prescription contact lenses do have some effect of iris biometric performance.

There is a known approach, due to Daugman [6], for detecting whether or not a person is wearing a cosmetic contact lens created by a dot-matrix type process. However, there are other means of manufacturing cosmetic contact lenses that are not detected by this approach. An example image representing a non-dot-matrix type of cosmetic lens is shown in Fig. 1.1.

The problem of automatically detecting cosmetic contact lenses is an important one because cosmetic contacts effectively obscure the natural iris texture. A person wearing a cosmetic contact lens in an image in which they are to be recognized has a very small chance of matching an enrollment image in which they are not wearing contacts.

The issue of contact lenses effecting the accuracy of iris matching actually applies across all types of contacts, not just cosmetic lenses. Figure 1.2 shows an example image in which the “AV” logo printed on the lens is visible. In general, the logo will appear in a different place on the iris in different images, leading to a generally degraded match between two images of the same eye wearing the same contact.

The typical “soft contacts” are the most popular type of contacts, but not the only type. There are also rigid gas permeable lenses or “hard lenses.” These can cause some large artifacts in the iris image, an example of which is shown in Fig. 1.3.

It is not only lenses that have printing on them or hard lenses that can cause problems. Clear, prescription contacts that do not properly fit the eye can cause




Fig. 1.2 Example iris image exhibiting “AV” lettering printed on contact lens. Note the “AV” lettering near the iris-sclera boundary on the side near the nose




Fig. 1.3 Example iris image exhibiting a rigid gas permeable (“hard”) contact lens

dramatic effects in the iris image. An example of this appears in Fig. 1.4. In this case, the bright and dark arcs that appear in the image are an artifact of the poorly fitting contact lens.

Even if the absence of all the special circumstances mentioned above, Baker et al. [4] have shown that plain, clear prescription contacts still have an effect of slightly degrading the match quality. For example, the lens shown in Fig. 1.5 shows some circular banding artifact that can affect the match quality.



Fig. 1.4 Example iris image exhibiting bright and dark arcs due to poor contact lens fit




Fig. 1.5 Example plain contact lens image exhibiting artifacts in the image. Note the circular banding, most prominent near the iris-sclera boundary on the bottom of the eye

In addition to the various types of image artifacts that may be caused by contact lenses, there are also a variety of special conditions that may result from medical or biological causes. The iris in the image in Fig. 1.6 is an example of an iris that was damaged in an accident. It shows a boundary between the pupil and iris that



Fig. 1.6 Example iris image showing accidental damage. Note the irregular shape of the boundary of the pupil




Fig. 1.7 Example iris image showing effects of persistent pupillary membrane

is naturally distinctly noncircular. This boundary is likely to cause problems for most iris segmentation routines. Similar such noncircular boundaries can occur from various medical conditions as well.

Figure 1.7 shows an image that represents a condition called *persistent pupillary membrane*. This condition does not present problems with a person's sight. But the



Fig. 1.8 Example iris image showing an artificial eye. Note the scattered specular highlights

strands that connect from the pupil out over the iris are the remains of a fetal membrane and not a part of the actual iris texture. The persistent pupillary membrane strands can potentially complicate iris segmentation and can (rarely) retract and so change the imaged iris texture.

Figure 1.8 shows an example iris image in which the person actually has an artificial eye. The specular highlights can be very different across different images of this eye, potentially complicating the segmentation.

The above examples are just a few of the many different special conditions that an iris biometric system must encounter and deal with in a large-scale application. There are also many smaller, more subtle “everyday” issues that make the analysis of iris images challenging. Figure 1.9 shows an image in which there is a specular highlight in the iris region due to reflection of the illuminator off the side of the nose. Such interreflections are often visible on the side of the iris near the nose.

1.1.1 A Survey of Iris Biometrics Research: 2008–2010

The field of iris biometrics research is growing and fast changing. In this chapter, Bowyer and coworkers update their 2008 survey paper [5] to cover selected subareas of iris biometrics research through approximately 2010. The growth in iris biometrics research has been so explosive that, even covering this relatively short time period, the number of references in this chapter exceeds that in the 2008 survey. This survey should be useful to any researcher who wishes to acquire a “big picture” view of the current state of iris biometrics research, or who wishes to delve into greater detail about a specific line of iris biometrics research.




Fig. 1.9 Example iris image showing an artificial eye. Note the small, elongated specular highlight in the iris region on the side near the nose. This is caused by reflection of the illumination off of the nose. It may appear in different places in different images of the iris, depending on the geometric relation between illuminator, nose and iris, and on the specularity of the nose surface

1.1.2 Standard Iris Storage Formats

Iris biometric technology is becoming widely used, as evidenced, for example, in India's Unique ID program. As the iris biometric industry grows, standards become increasingly important. Quinn and coworkers from the National Institute of Standards and Technology present an overview of work in the development of standards for data storage formats in iris biometrics. They discuss issues of storage size constraints, standard storage formats, compression, and its effects on matching accuracy. It is important to see the distinction between standard data formats that allow interchange of iris data and potentially nonstandard biometric templates. As the authors state, "Standardized iris images (i.e. iris records), are not iris templates. Rather, they are specialized interoperable images designed for the efficient storage and transmission of iris data." One interesting conclusion of work in this area is that biometrically useful iris image data can be stored in as little as 2 KB.

1.1.3 Iris Quality Metrics for Adaptive Authentication

The topic of quality metrics for iris images is one that has attracted substantial attention in the research community. It seems even more important for systems that are meant to allow image acquisition under less constrained conditions. Schmid and

coworkers discuss both commonly used quality metrics and some of their own more novel metrics. These include metrics for quality of iris segmentation, interlacing, illumination, lighting, occlusion, area in pixels, dilation ratio, off-angle, and blur. They show results using the ICE 2005 dataset that demonstrate that matches between images with better values of a metric result in a better ROC curve. They also present a method to combine quality metrics for an image into an overall score for that image, a method to use the scores for two images to give a confidence level for their match, and a method to use the quality metrics to improve the performance the matching algorithm.

1.1.4 Quality and Demographic Investigation of ICE

Phillips and Flynn present the results of “mining” the quality scores submitted for iris images as part of the Iris Challenge Evaluation 2006. ICE 2006 [14] was a relatively large-scale evaluation of iris biometric performance, done using images acquired with the LG 2200 system. Three participants in the ICE 2006 evaluation submitted the optional quality metric for images. The quality score was required to be an integer in the range of 0–100, with 100 representing highest quality. The authors present results looking at correlation of quality scores between left and right irises, correlation between quality scores and subject demographics, and between quality scores and biometric performance. They note that the correlation coefficients for the pairs of quality scores are rather low, all less than 0.35. They also note that the distributions of three different quality scores over the allowed 0–100 range are qualitatively different. These two facts argue for the quality metrics in the three different systems being focused on different phenomena.

1.1.5 Iris Recognition with Taylor Expansion-Based Features

Bastys and coworkers describe two alternative descriptions of iris texture and their use in iris matching. The texture descriptions are based on using the first- and second-order Taylor series expansion. One representation is a phase-based description based on binarization of the Taylor expansion at two scales. The other representation is based on local extrema in the first two Taylor series expansion coefficients. A sector-based elastic similarity metric is used for comparing two irises, to allow for minor segmentation inaccuracies. The two proposed texture representations show different relative performance on three datasets used for evaluation: CASIA II, ICE 2005, and Multiple Biometric Grand Challenge (MBGC) [15] Portal Challenge. However, the fusion of the two approaches is consistently better than either one on all three datasets. The authors point out that “despite the intuitive similarity to the Gabor features, the Taylor coefficients have greater localization in the space domain,” and although this is “at the expense of poor feature resolution

in the frequency domain,” this problem may be less important because “periodic texture patterns that are well localised in frequency domain are rare” in the iris.

1.1.6 A Theoretical Model for Describing Iris Dynamics

It was once believed that the degree of pupil dilation did not impact the performance of iris biometrics. However, Hollingsworth et al. [10] showed experimental evidence that the difference in pupil dilation between the images being matched does in fact have a significant effect. The larger the difference in dilation ratio between two images of the same iris, the greater the likelihood that a false non-match will occur. This same basic result was also found later in the NIST IREX study [9]. This has led to a desire for a better understanding of the basic process of pupil dilation. Clark and coworkers give an overview of work related to the effects of pupil dilation on iris biometrics and present a theoretical study of the nonlinear dynamics involved in iris deformation during changes in dilation. They give a detailed mathematical development of the dynamics of iris deformation and present finite element model simulation results for the case of the iris being isotropic material and for the case of the iris being orthotropic material. The authors acknowledge that there is much more work still to be done in this area, as “there is a need to explore testing our model extensively on actual data to draw a more accurate comparison” and “there is a need to build a mathematical model that also takes into consideration the dynamics in the thickness of the iris.”

1.1.7 Iris Recognition in the Visible Wavelength

All commercial iris biometric systems that we are aware of, and the vast majority of all academic research on iris biometrics, are based on using iris images acquired with near-infrared illumination. Under appropriate near-infrared illumination, there effectively are no “light” and “dark” irises, and all irises have readily visible texture. Hugo Proença is perhaps the most vocal proponent of performing iris biometrics using images acquired with visible-wavelength illumination rather than near-infrared illumination. In this chapter, he presents the argument for using visible-wavelength illumination, describes the UBIRIS dataset of visible-wavelength iris images, and discusses some results of experiments on UBIRIS data. The UBIRIS dataset is freely available to the biometrics research community and has been used as the basis for the Noisy Iris Challenge Evaluation (NICE) competitions. In general, under visible-wavelength illumination, “light” eyes are easier to recognize than “dark” eyes. Also, performance overall is not as good as with commercial near-infrared illumination systems. While it does not yet appear that visible-wavelength iris biometrics is ready to overtake near-infrared illumination iris biometrics, it presents a significant and challenging research problem.

1.1.8 Multispectral Iris Fusion and Cross-spectrum Matching

Burge and Monaco attack the difficult problem of cross-spectrum iris matching. They report on methods of matching iris images acquired in near-infrared illumination to an iris image acquired under normal visible illumination. They approach this problem by creating an estimated NIR iris image from the color iris image using features derived from the color and structure of the visible light iris image. They report on experiments performed with images acquired using a custom camera system that is capable of simultaneously acquiring registered images at four wavelengths (i.e., 460, 550, 670, and 800 nm). Their results show that, for lighter-colored portions of the iris, texture content increases with the frequency of the illumination and that similarly, for darker-colored portions of the iris, texture content decreases.

1.1.9 Robust and Secure Iris Recognition

The area of cancelable biometrics deals with the study of systems that enroll a template based on a transformation of the original biometric. This is an area that was pioneered by researchers at IBM Research. In this approach, if the enrolled template is compromised and used by an impersonator, the enrollment can be canceled and a different transformation used to create a new enrollment. Pillai and coworkers develop techniques in sparse representation and random projections to create “a unified framework for image quality estimation, recognition and privacy in iris biometrics.” They also exploit the approach of considering an iris image in sectors, matching sectors individually, and then fusing results. In the context of recognition using iris videos, the sectors may come from different frames of the video, depending on how quality varies in different frames. They present results in this area based on the MBGC [15] iris video dataset.

1.1.10 Template Aging in Iris Biometrics

Mansfield and Wayman, in their 2002 report on best practices in biometric testing [12], define template aging as follows: “Template ageing refers to the increase in error rates caused by time related changes in the biometric pattern, its presentation, and the sensor.” One claim that has been widely accepted since the inception of iris biometrics is that iris texture is somehow stable over a person’s lifetime in a way that results in iris biometrics not being subject to template aging. Baker et al. published a paper [3] in the 2009 International Conference on Biometrics with experimental results that demonstrate that template aging does occur for iris biometrics. In particular, their results showed that the false non-accept rate increases for comparison of images taken roughly 4 years apart compared to images taken

only a few months apart. In 2011, Fenker and Bowyer published a study on a different dataset that reaches essentially the same conclusions [8]. In this chapter, Baker and coworkers present an extended analysis on a larger dataset than used in their original study. The conclusions of this study may be controversial, since they challenge an element of conventional wisdom about iris biometrics.

1.1.11 Fusion of Face and Iris Biometrics

Connaughton and coworkers report the results of a multi-biometric face and iris study performed using video clips acquired with an Iris on the Move (IOM) system [13]. The dataset includes 1,886 videos covering 363 different subjects. They explicitly include several different senses of the term multi-biometrics. One, they combine results of iris and face, since both are available in the near-infrared video obtained with the IOM system. Second, they use multi-sample approaches, combining information from individual frames of a video. Third, they use multi-algorithm, combining results from multiple face matchers and from multiple iris matchers. Thus, they have a large number of results of varying types to consider at the fusion stage. They find that fusion of iris across a video segment performs better than fusion of face across a video segment, but that fusion of both performs better still.

1.1.12 Methods for Iris Segmentation

Jillela and Ross discuss the problem of accurately segmenting the iris region and provide an overview of progress that has been made in iris segmentation. They begin by showing examples of iris images that present problems representing a wide range of image quality factors. Then Daugman's integrodifferential segmentation operator and the Hough transform approach proposed by Wildes are summarized in some detail. Then more modern approaches, such as geodesic active contours, variational level sets, and Fourier-based approximation, are introduced. The effect of image acquisition conditions on iris segmentation is considered, in particular including images acquired with the Iris on the Move (IOM) system. Finally, various techniques to refine the results of an initial iris segmentation are described.

1.1.13 Methods for Periocular Segmentation

Jillela and Ross extend classic iris segmentation algorithms to the periocular region. Definitions of the periocular region vary, but typically encompass the skin covering the orbit of the eye. The performance of five segmentation algorithms on the

periocular images of the Face and Ocular Challenge Series (FOCS) database are given. FOCS images are characterized by large variations in illumination, eyelid and eyelash occlusion, defocus blur, motion blur, and low resolution.

1.1.14 Periocular Recognition from Low Quality Iris Images

Klontz and Burge show that especially in cases where the iris has not been acquired with sufficient quality to compute an IrisCode, the periocular region can provide additional discriminative information for biometric identification. They investigate periocular recognition on the FOCS dataset using three distinct classes of features: photometric, keypoint, and frequency-based. They examine the performance of these features alone, in combination, and when fused with classic IrisCodes.

1.1.15 Introduction to the IrisCode Theory

Kong and coworkers present an updated version of their theoretical analysis of the IrisCode approach [11]. In this context, IrisCode refers specifically to the Daugman-style approach to iris biometrics, in which a binary code is used that summarizes the features of the biometric. But the theoretical analysis is aimed at illuminating issues that are deeper than just the use of a binary code. Are the 0 and 1 bits in the iris code equally probable? Is the impostor distribution truly binomial? The authors summarize their theoretical analysis as follows: “we will prove that the IrisCode is a clustering algorithm with four prototypes; the locus of a Gabor function is a two-dimensional ellipse with respect to a phase parameter and can be approximated by a circle in many cases; the Gabor function can be considered as a phase-steerable filter, and the bitwise hamming distance can be regarded as a bitwise phase distance.” As an illustration, they present a method that encodes $n > 2$ bits from a single texture filter result. This chapter is an important contribution in a relatively understudied area of iris biometrics.

1.1.16 Application of Correlation Filters for Iris Recognition

In this chapter, Kumar and coworkers present an approach to iris biometrics developed around a different technical core than that used in Daugman’s approach. Daugman’s approach revolves around creating an iris code based on applying a Gabor filter at a grid of locations on the iris image. Kumar and coworkers develop an approach that uses correlation filters as the basis for iris segmentation and matching. The use of correlation filters is motivated by the desire to handle iris images of degraded quality, of the type that might occur when the iris is imaged from a greater and less controlled distance than is typical in current commercial iris biometric

systems. They present an overview of major concepts in the design of correlation filters, show how to use a cross-correlation method for iris segmentation, show how to perform correlation-based matching of unwrapped iris images, and show how to use Bayesian graphical models to improve iris matching in the presence of nonlinear deformations. In dealing with nonlinear deformations, they consider the unwrapped iris image as a grid of sub-images that can deform differently from each other.

1.1.17 Iris Spoofing

Venugopalan and Savvides outline an approach to “spoofing” in iris biometrics. Their approach assumes that a particular iris code to be spoofed is known and that the method of analyzing iris texture to produce an iris code is also known. Based on these, they are able to demonstrate a method to synthesize an iris texture that will result in an iris code similar to the one to be spoofed. They demonstrate experimental results using images from the ICE 2005 iris image dataset [14], showing ROC curves for verification of spoof iris textures matched to genuine iris textures. Iris biometric spoofing is overall a relatively understudied area, and this work is an important contribution. As the authors state, “Exploring methods such as those presented in this chapter are essential to determine loopholes in state of the art systems that have been deployed at high security facilities and in order to develop counter-measures for the same.”

1.1.18 Optics of Iris Imaging Systems

Ackerman presents a detailed overview of the concerns that enter into the design of a system for imaging the iris. He begins with a statement of basic assumptions about the iris itself: “An iris. . . is about 11 ± 1.5 mm in diameter and sits behind the partially reflective cornea. We note that the diffuse reflectivity of an iris, referred to as its albedo, is dependent on illumination wavelength and is typically low, around 10%, in the near infrared (NIR) band.” Ackerman reviews basic concepts of optic design in the context of imaging the iris and discusses how considerations of lens quality, focal length, depth of field, spatial resolution, signal and noise, and currently available imaging technology enter into the design of a system to image the iris. This leads to a comparison of the trade-offs in designing a system to image the iris at 0.3 m versus imaging the iris at 3 m. The issues involved in the optics of imaging the iris are nontrivial, given the need to image small detail within an already small object with sufficient contrast to support reliable texture analysis. While this chapter does not discuss segmentation of the iris region in an image, or the matching of iris texture patterns, the imaging details summarized are fundamental information for all of the higher-level processing that follows image acquisition.

References

1. Aadhaar gains momentum, clocks two crores in October. *The Times of India* (2011). articles.timesofindia.indiatimes.com/2011-11-02/india/30349953_1_aadhaar-numbers-uidai-director-general-enrollment-stations
2. Al-Raisi, A.N., Al-Khoury, A.M.: Iris recognition and the challenge of homeland and border control security in UAE. *Telemat. Inform.* **25**(2), 117–132 (2008)
3. Baker, S.E., Bowyer, K.W., Flynn, P.J.: Empirical evidence for correct iris match score degradation with increased time-lapse between gallery and probe matches. In: *Advances in Biometrics: LNCS #5558*, pp. 1170–1179. Springer-Verlag Berlin, Heidelberg (2009)
4. Baker, S.E., Hentz, A., Bowyer, K.W., Flynn, P.J.: Degradation of iris recognition performance due to non-cosmetic prescription contact lenses. *Comput. Vis. Image Underst.* **114**(9), 1030–1044 (2010)
5. Bowyer, K., Hollingsworth, K., Flynn, P.: Image understanding for iris biometrics: a survey. *Comput. Vis. Image Underst.* **110**(2), 281–307 (2008)
6. Daugman, J.: Demodulation by complex-valued wavelets for stochastic pattern recognition. *Int. J. Wavelets Multiresolut. Inf. Process.* **1**(1), 1–17 (2003)
7. Daugman, J., Malhas, I.: Iris recognition border-crossing system in the UAE. *Biometrics* **44**(2), 49–53 (2004)
8. Fenker, S., Bowyer, K.W.: Experimental evidence of a template aging effect in iris biometrics. *IEEE Workshop on Applications of Computer Vision (WACV 2011)*. pp. 232–239 (2011)
9. Grother, P., Tabassi, E., Quinn, G.W., Salamon, W.: IREX I performance of iris recognition algorithms on standard images NIST interagency report 7629. biometrics.nist.gov/cs_links/iris/irex/irex_summary.pdf. Gaithersburg, Maryland (2009)
10. Hollingsworth, K., Bowyer, K., Flynn, P.: Pupil dilation degrades iris biometric performance. *Comput. Vis. Image Underst.* **113**(1), 150–157 (2009)
11. Kong, A.W.K., Zhang, D., Kamel, M.S.: An analysis of iriscode. *IEEE Trans. Image Process.* **19**(2), 522–532 (2010)
12. Mansfield, A.J., Wayman, J.L.: Best practices in testing and reporting performance of biometric devices ver 2.01. *Natl. Phys. Lab.* **14**(02), 1–36 (2002)
13. Matey, J.R., Naroditsky, O., Hanna, K., Kolczynski, R., Lolocono, D.J., Mangru, S., Tinker, M., Zappia, T.M., Zhao, W.Y.: Iris on the move: acquisition of images for iris recognition in less constrained environments. *Proc. IEEE* **94**(11), 1936–1947 (2006)
14. Phillips, P.J., Scruggs, W.T., O’Toole, A.J., Flynn, P.J., Bowyer, K.W., Schott, C.L., Sharpe, M.: FRVT 2006 and ICE 2006 large-scale experimental results. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(5), 831–846 (2010)
15. Phillips, P.J., Flynn, P.J., Beveridge, J.R., Scruggs, W.T., O’Toole, A.J., Bolme, D., Bowyer, K.W., Draper, B.A., Givens, G.H., Lui, Y.M., Sahibzada, Scallan, J.A., Weimer, S.: Overview of the multiple biometrics grand challenge. In: *Advances in Biometrics: LNCS #5558*, pp. 75–714. Springer-Verlag Berlin, Heidelberg (2009)

Chapter 2

A Survey of Iris Biometrics Research: 2008–2010

Kevin W. Bowyer, Karen P. Hollingsworth, and Patrick J. Flynn

Abstract A recent survey of iris biometric research from its inception through 2007, roughly 15 years of research, lists approximately 180 publications. This new survey is intended to update the previous one, and covers iris biometrics research over the period of roughly 2008–2010. Research in iris biometrics has expanded so much that, although covering only 3 years and intentionally being selective about coverage, this new survey lists a larger number of references than the inception-through-2007 survey.

2.1 Introduction

Iris biometrics research is an exciting, broad, and rapidly expanding field. At the same time that there are successful practical applications that illustrate the power of iris biometrics, there are also many fundamental research issues to be solved on the way to larger scale and more complex applications.

A survey that appeared in 2008 covered the field from its inception in the early 1990s through roughly the end of 2007 [21]. This new survey is intended to update the previous one, covering roughly the period 2008–2010. However, as illustrated in Fig. 2.1, there has been tremendous growth in the literature in this area. Due to this growth, this new survey does not attempt as exhaustive a coverage of the field as the previous survey. We focus primarily on papers that appeared in SpringerLink or in IEEE Xplore, as these appear to currently be the two major sources of publications in this field. We also omit coverage of some subareas of work judged to be of less importance. These omissions are explained at the appropriate points in the survey.

K.W. Bowyer (✉) • K.P. Hollingsworth • P.J. Flynn
University of Notre Dame, IN, USA
e-mail: kwb@cse.nd.edu; kholling@cse.nd.edu; flynn@cse.nd.edu

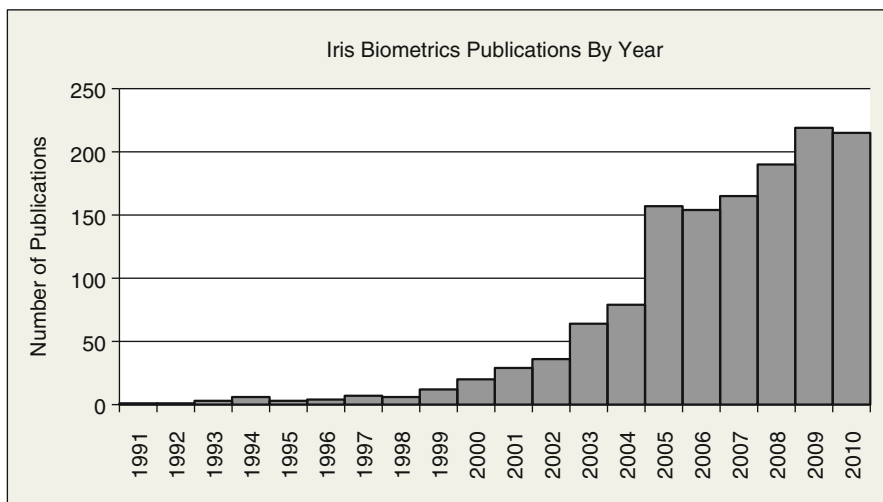


Fig. 2.1 Iris biometrics papers in Google Scholar from 1990 through 2010 (This data was taken using Google Scholar’s “advanced search” facility, searching for “iris biometrics pupil” appearing in articles, excluding patents, in the Engineering, Computer Science, and Mathematics literature)

The main body of this survey is organized into the following sections:

2. Iris image acquisition
3. Iris region segmentation
4. Texture coding and matching
5. Multi-biometrics involving the iris
6. Privacy and security
7. Datasets and evaluations
8. Performance under varying conditions
9. Applications
10. Theoretical analyses

Papers are grouped into a section according to their perceived main area of contribution. In some instances, a paper is mentioned in more than one section. The survey ends with a short discussion and a list of recommended readings.

There are several overview or introductory type articles that can be mentioned in this section. Gorodnichy [59] gives a good overview/introduction to biometrics, emphasizing evaluation of biometric system performance based on a dynamic, or life cycle view of operational systems. Bhattacharyya et al. [14] give a short, high-level overview of biometrics, primarily emphasizing iris biometrics. Phillips and Newton [143] present a short “point of view” type article on biometric evaluation, emphasizing issues such as the number of persons represented in the dataset and the longitudinal time over which biometric samples are collected. Each of these articles contains important elements for anyone new to the field of biometrics.

2.2 Iris Image Acquisition

There are still major research issues in the area of iris image acquisition. One issue involves imaging the iris with a sensor system that allows the person to be more “at a distance” and “on the move.” Matey and Kennell [119] present a comprehensive tutorial on the issues involved in acquiring iris images at a distance of greater than 1 m. The tutorial includes a partial list of commercial iris recognition devices released between 1995 and 2008 and a description of several successful applications of iris biometrics. The authors describe acquisition issues including the wavelength of light used, the type of light source, the amount of light reflected by the iris back to the sensor, required characteristics of the lens, signal-to-noise ratio, eye safety, and image quality. Capture volume, residence time, and sensitivity to subject motion are also discussed.

Wheeler et al. [207] describe a prototype “standoff” iris recognition system designed to work at sensor-to-subject distances of up to 1.5 m. The system uses two wide-field-of-view cameras to perform face location in the scene and an iris camera and illuminator to image the iris. Dong et al. [37] discuss the design of a system to image the iris “at a distance,” allowing a standoff of 3 m. Although current commercial iris biometrics systems all use near-infrared (NIR) illumination and most research assumes NIR imaging similar to that used in current commercial sensors, Proenca [152] argues for visible-wavelength imaging as the more appropriate means to achieve “at a distance” and “on the move” imaging.

Boddeti and Kumar [18] investigate the use of wavefront-coded imagery for iris recognition. This topic has been discussed in the literature before, but Boddeti and Kumar use a larger dataset and present experiments to evaluate how different parts of the recognition pipeline (e.g., segmentation, feature extraction) are affected by wavefront coding. They propose using unrestored image outputs from the wavefront-coded camera directly and test this idea using two different recognition algorithms. They conclude that wavefront coding could help increase the depth of field of an iris recognition system by a factor of 4 and that the recognition performance on unrestored images was only slightly worse than the performance on restored images.

There is little published work dealing with imaging the iris under different wavelength illumination. Ross et al. [170] look at imaging the iris with illumination in the 950–1,650 nm range, as opposed to the 700–900 nm range typically used in commercial systems. They suggest that it is possible to image different iris structure with different wavelength illumination, raising the possibility of multispectral matching as a means to increased recognition accuracy.

Grabowski et al. [61] describe an approach to iris imaging that is meant to allow characterization of structures in the iris tissue over changes in pupil dilation. They use side-illumination, fixed to glasses frames worn by the subject, with imaging resolution that allows an 800-pixel iris diameter. This is many more “pixels on the iris” than in current commercial sensors.

Chou et al. [32] describe an iris image acquisition system meant to handle off-angle views of the iris and to make iris segmentation easier and more reliable.

Their system uses a dual-CCD camera to acquire a color RGB image with one CCD and a near-infrared image with the other. The color image is exploited to improve the reliability of the segmentation. The non-orthogonal-view iris image is rectified to an orthogonal-view iris image using the pupillary boundary.

He et al. [70] design their own iris camera with the goal of being more economical than commercial alternatives while still acquiring high-quality images. They use a CCD sensor with resolution of 0.48 M pixels and add a custom glass lens with a fixed focus at 250 mm and NIR-pass filters that transmit wavelengths between 700 and 900 nm. The illumination unit consists of NIR LEDs of 800 nm wavelength, which they arrange to try to minimize specular reflections on the iris.

McCloskey et al. [121] explore a technique termed “flutter shutter” as a means to acquire sharply focused iris images from moving subjects. The idea is that the camera shutter “flutters” between open and closed while the sensor accumulates an image, from which an appropriately designed deblurring algorithm can then recover an in-focus image.

2.2.1 Non-ideal Images and Quality Metrics

As mentioned earlier, one important current research emphasis is acquisition of images under less-constrained conditions. As iris images are acquired under less-constrained conditions, the issue of image quality becomes more important and complex. Another element of this is the design of algorithms meant to handle “nonideal” or “noisy” images. For our purposes, “nonideal” means something more than just the presence of specular highlights or occlusion by eyelashes or eyelids.

While it is not part of the image acquisition step per se, iris biometric systems typically evaluate the focus quality and possibly other factors, of each candidate image in order to select usable images. Ren and Xie [167, 168] propose approaches to evaluating image focus quality that involve finding the iris region before computing the focus value. While iris biometric systems select images based in part on focus quality, there are few publications dealing with deblurring of iris images. Huang et al. [82] investigate image deblurring algorithms that exploit context specific to iris imagery. He et al. [72] estimate the user distance from the sensor in order to estimate the appropriate point spread function (PSF) for image restoration. They measure the distance between two specular highlights on the iris. Using this information, plus knowledge about the positions of the two infrared LEDs, they get the user’s distance from the camera without using a special distance sensor. The knowledge of the distance from the sensor is used in estimating the PSF.

Belcher and Du [10] combine percent occlusion, percent dilation, and “feature information” to create an iris image quality metric. To compute “feature information,” they calculate the relative entropy of the iris texture when compared with a uniform distribution. To fuse the three types of information into a single score, they

first compute an exponential function of occlusion and an exponential function of dilation. The final quality score is the product of the three measures.

Kalka et al. [87] investigate a number of image quality factors, including percent occlusion, defocus, motion blur, gaze deviation, amount of specular reflection on the iris, lighting variation on the iris, and total pixel count on the iris. In evaluating various datasets, they found that the ICE data had more defocused images, the WVU data had more lighting variation, and the CASIA data had more occlusion than the other sets.

Schmid and Nicolo [181] evaluate iris image quality metrics in terms of how well they predict recognition performance. The quality metric is applied to each of a pair of images being matched, and the metrics mapped to a predicted matching score. The metric(s) can then be evaluated by how well the predicted matching score is correlated with the calculated matching score. Schmid and Nicolo experiment with both iris and face image data.

Zhou et al. [217, 218] propose adding four modules to the traditional iris biometrics system in order to handle nonideal images. A “quality filter unit” eliminates images that are too poor quality to be useful. A “segmentation evaluation unit” evaluates the quality of the segmentation. A “quality measure unit” determines if there is sufficient iris area available to generate features. A “score fusion unit” combines a segmentation score and a quality score. Experiments are shown using the MBGC dataset [146] and their own IUPUI near-field iris video dataset.

Zuo and Schmid [219] present both a global quality metric for selecting individual frames from an iris video or image sequence and multiple local quality metrics for the iris in a given frame. The global quality metric experiments use the Iris On the Move [120] videos distributed as part of the Multiple Biometric Grand Challenge [146]. The local quality metrics look at segmentation quality, interlacing, illumination contrast, illumination evenness, percent occlusion, pixel count, dilation, off-angle view, and blur and are evaluated using images from the ICE 2005 dataset [20, 145].

Breitenbach and Chawdhry [23] perform experiments looking at quality factors for an image and how they predict performance of face and iris recognition. They synthetically vary image factors such as defocus, contrast, and resolution. They find that the factors considered are better predictors of iris biometric performance than face recognition performance.

Proenca [154] presents an approach to quality assessment of iris images acquired in the visible-light domain. Factors considered in the quality assessment include focus, motion, angle, occlusions, area, pupillary dilation, and levels of iris pigmentation. The claim is that by using the output of the segmentation phase in each assessment, the method is able to handle severely degraded samples.

Phillips and Beveridge [142] present a challenging view on the topic of using quality metrics in biometric matching. By analogy to AI-completeness in artificial intelligence and completeness in the theory of algorithms, they introduce the concept of biometric-completeness. The idea is that a problem in biometrics is biometric-complete if it can be shown to be equivalent to the biometric recognition

problem and “the key result in this paper shows that finding the perfect quality measure for any algorithm is equivalent to finding the perfect verification algorithm.”

2.2.2 *Image Compression*

Daugman and Downing [36] present a detailed study of the effects of compression of the original iris image on the performance of iris biometrics. They present schemes that combine isolation of the iris region with JPEG and JPEG 2000 compression, evaluate their approach on images from the Iris Challenge Evaluation (ICE) 2005 dataset [20, 145], and conclude that it is “possible to compress iris images to as little as 2,000 bytes with minimal impact on recognition performance.”

Ives et al. [83] explore the effect of varying levels of JPEG 2000 compression, using the ICE 2005 dataset [20, 145], and find that the false reject rate increases with increasing level of compression, but that the false accept rate is stable.

Konrad et al. [94, 95] aim to compress iris data without degrading matching results. They use JPEG compression on unwrapped polar iris images. They design and compare different quantization tables to use with the JPEG compression. Two of their tested Q-tables are designed to preserve more angular iris texture than radial iris texture (i.e., the horizontal texture in the unwrapped image). The other two Q-tables are derived from the first two through genetic optimization. There is no clear winner among their tested Q-tables, and they conclude that custom Q-tables for iris recognition should be optimized to a specific target bitrate for best performance.

Kostmayer et al. [96] apply compression to the original, rectilinear iris images. They propose custom JPEG quantization tables for iris recognition. Their theory is that the highest and medium frequencies are not essential for iris recognition performance because of the coarse quantization used in template generation. Therefore, they test four custom compression tables, each one with an increasing number of high frequencies suppressed. In most of their tests, their proposed tables outperform the standard JPEG quantization table. Based on their experiments, they conclude that iris compression is not realistic at compression rates greater than 20. On the other hand, their experiments show that compression does not noticeably affect the impostor distribution.

Hämmerle-Uhl et al. [64] use JPEG 2000 compression on original iris images. They aim to improve compression performance by using region-of-interest coding. They detect the iris using edge detection and a Hough transform, then set the ROI to the detected candidate circle with largest radius inside a certain allowed range. They compare compression with and without ROI coding and find that match scores improve and error rates decrease when using the ROI coding.

Carneiro et al. [26] examine the performance of different iris segmentation algorithms in the presence of varying degrees of fractal and JPEG 2000 image compression, using the UBIRIS dataset [155].

2.3 Iris Region Segmentation

Publications related to segmenting the iris region constitute a significant fraction of the published work in iris biometrics. Many of these publications can be grouped as tackling similar versions of the traditional iris segmentation problem, for example, given one still image, find the pupillary and limbic boundaries. However, there are also a variety of approaches being explored to find occlusion by specular highlights and eyelashes, to segment the iris using less-constrained boundaries, and to refine initial segmentation boundaries.

Iris segmentation algorithms that assume circular boundaries for the iris region continue to appear in some conferences. We have chosen not to cover this subarea of work here, as the current frontier in iris segmentation is generally now focused on removing the assumption of circular boundaries [27, 75, 183] and on refining the segmentation to account for various occlusions and distortions of the iris texture.

Publications also continue to appear that propose iris segmentation techniques that are evaluated on the CASIA version 1 dataset. Again, we have chosen not to cover this subarea of work in this survey. The use of the CASIA v1 dataset to evaluate iris segmentation algorithms is inherently problematic. This is because the images have been edited to have a circular region of constant-intensity value for the region of each iris [144]. Therefore, any segmentation algorithm built around the assumption of a circular region of constant dark intensity value should naturally meet with great success on this dataset, even though these conditions are generally not present in the iris region of real images.

A number of researchers have considered various approaches to segmenting the iris with boundaries not constrained to be circles. Wibowo and Maulana [208] evaluate an approach using the CASIA v1 data and their own dataset of 30 visible-light iris images. Labati et al. [100, 101] propose methods to find the pupil center and then to find the inner and outer iris boundaries, presenting experimental results on CASIA v3 and UBIRIS v2 images. Kheirolahy et al. [92] propose a method of finding the pupil in color images, with experiments on the UBIRIS dataset. Chen et al. [30, 31] consider an approach to segmenting the iris region under less-constrained conditions, experimenting with the UBIRIS v2 visible-light iris image dataset, and placing in the top six in the NICE competition. Broussard and Ives [24] train a neural net to classify pixels in an iris image as either being on an iris boundary or not, selecting the most useful eight features from a pool of 322 possible features. Subjective visual evaluation of results indicates improvement over methods that assume circular boundaries. Zuo and Schmid [220] present an approach to segmenting the iris using ellipses for the pupillary and the limbic boundaries, with experiments on CASIA, ICE, and WVU datasets. Pan et al. [137] detect edge points using “phase congruency analysis” and fit ellipses to the detected edge points. They test their method on CASIA v2 and CASIA v3 twins datasets. Roy and Bhattacharya [174–176] suggest a segmentation method using geometric active contours. They apply opening operators to suppress interference from eyelashes [175]. Next, they approximate elliptical boundaries for the pupil

and limbic boundaries. They refine the detected boundary using geometric active contours (i.e., active contours implemented via level set) to a narrow band over the estimated boundary. They fit parabolic curves to the upper and lower eyelids. To isolate the eyelashes, they use 1D Gabor filters and variance of intensity. Roy and Bhattacharya [177] also describe a level set style active contour method for finding the pupil and iris boundaries in nonideal iris images, presenting results on the UBIRIS v2, ICE 2005, and WVU nonideal iris datasets.

Ryan et al. [179] present the “starburst method” for segmenting the iris. They preprocess the image using smoothing and gradient detection, and then they find a pupil location as a starting point for the algorithm. To do so, they set the darkest 5% of the image to black and all other pixels to white. Then they create a chamfer image: the darkest pixel in the chamfer image is the pixel farthest from any white pixel in a thresholded image. They use the darkest point of the chamfer image as a starting point. Next, they compute the gradient of the image along rays pointing radially away from the start point. The two highest gradient locations are assumed to be points on the pupillary and limbic boundaries and are used to fit several ellipses using randomly selected subsets of points. An average of the best ellipses was reported as the final boundary. The eyelids were detected using active contours.

Pundlik et al. [156] treat the image as a graph where pixels are nodes and neighboring pixels are joined with edges. Their first goal is to assign a label – either “eyelash” or “non-eyelash” – to each pixel. After removing specular reflections, they use the gradient covariance matrix to find intensity variation in different directions for each pixel. Then they create a probability map, P , that assigns the probability of each pixel having high texture in its neighborhood. The “energy” corresponding to a particular labeling of the images is written as a function of a smoothness term and a data term. The data term is based on a texture probability map. The second goal was to assign each pixel one of four labels: eyelash, pupil, iris, or background. They use a method similar to the initial eyelash segmentation; however, this time they use an alpha-beta swap graph-cut algorithm. Finally, they refine their labels using a geometric algorithm to approximate the iris with an ellipse.

Vatsa et al. [196] improve the speed of active contour segmentation by using a two-level hierarchical approach. First, they find an approximate initial pupil boundary, modeled as an ellipse with five parameters. The parameters are varied in a search for a boundary with maximum intensity change. For each possible parameter combination, the algorithm randomly selects 40 points on the elliptical boundary and calculated total intensity change across the boundary. Once the pupil boundary is found, the algorithm searches for the iris boundary in a similar manner, this time selecting 120 points on the boundary for computing intensity change. The approximate iris boundaries are refined using an active contour approach. The active contour is initialized to the approximate pupil boundary and allowed to vary in a narrow band of ± 5 pixels. In refining the limbic boundary, the contour is allowed to vary in a band of ± 10 pixels.

Although there are relatively few papers devoted specifically to this topic, better detection of specular highlights in the iris image is still an area of current research [182, 210]. He et al. [71] acknowledge the difficulty of detecting and

removing specular highlights in the iris image and present an interesting multi-sample approach to this problem. They assume that multiple images of the same iris are available, with the specular highlights appearing in different places on the iris in different images. The segmentation of the iris region in the images is simple and assumes concentric circular boundaries for the pupil and iris. The rectangular iris images from the multiple images are then registered, bright spots detected, and the bright spots replaced with values from a different image.

Liu et al. [111] propose a method for eyelid detection in UBIRIS v2 (visible light) images. Their method uses a parabolic integro-differential operator similar to the operator described by Daugman for iris localization. They find that their proposed method has lower pixel error compared to algorithms involving the IDO alone, using detected edge pixels alone, or an algorithm using Canny edge detection and a Hough transform.

While most publications assume a single still image as the input to the segmentation stage, Du et al. [43] propose a method of using multiple thresholds on the intensity value in an image to achieve a rough segmentation of the iris in frames of a video sequence. Du et al. [44] also propose an approach to segmentation of iris images obtained in a context in which the subject is not explicitly cooperative. They filter to drop video frames in which the iris is not visible, fit ellipses for the iris boundaries, and develop a method to remove noise in the iris region.

Several researchers have considered the problem of evaluating the quality of an iris segmentation. Kalka et al. [86] tackle the problem of predicting or detecting when iris segmentation has failed, with experiments on the WVU and ICE datasets and on two iris segmentation algorithms. Li and Savvides [107, 108] present work on taking an existing iris segmentation mask, in principle from any algorithm, and automatically refining it to produce a better segmentation.

Proenca [151, 153] observes that images acquired in the visible wavelength in less-constrained environments tend to have noise that results in severely degraded images. Whereas many iris biometric segmentation algorithms key on the pupil to anchor the segmentation, he proposes to anchor the segmentation on the sclera as much more naturally distinguishable than any other part of the eye. The sclera also provides a useful constraint, in that it must be immediately adjacent on both sides of the iris. One of the differences in iris biometrics processing for visible-light versus near-IR images is that the pupillary boundary tends to be more distinct in near-IR, whereas the limbic boundary appears to be more distinct in visible light.

Lee et al. [103] describe a way to locate and analyze eyes in the MBGC portal videos. They use the Viola-Jones detector that comes with OpenCV and is trained to detect eye pairs. They measure the edge density in an image to determine the focus level and select appropriate frames from the video. The IrisBEE algorithm [147] is used for segmentation and feature extraction. Eyes from the MBGC portal videos are compared to higher-quality still iris images. The two-eye detection rate in the videos was 97.7%. The segmentation rate was 81.5%, and the matching rate was 56.1%. This matching rate is low compared to typical iris recognition systems, likely reflecting the low level of iris image quality in the MBGC portal videos.

Munemoto et al. [134] suggest that “it is important to not only exclude the noise region, but also estimate the true texture patterns behind these occlusions. Even though masks are used for comparison of iris features, the features around masks are still affected by noise. This is because the response of filters near the boundary of the mask is affected by the noisy pixels.” They used an image-filling algorithm to estimate the texture behind the occlusions. This algorithm iteratively fills 9×9 patches of the occluded region with 9×9 patches from unoccluded regions. It estimates textures at the boundary of the region first, selecting 9×9 source patches from the unoccluded iris that closely match the iris texture near the boundary of the area to be filled.

Thompson and Flynn [194] present a method of improving the recognition performance of iris biometrics by perturbing parameters of the iris segmentation. The perturbations generate a set of alternate segmentations, and so also alternate iris codes, which effectively result in an improved authentic distribution.

2.4 Texture Coding and Matching

Performing texture analysis to produce a representation of the iris texture, and the matching of such representations, is at the core of any iris biometric system. A large fraction of the publications in iris biometrics deal with this area. It is not necessarily straightforward to organize these publications into well-defined and meaningful categories. Nevertheless, they are grouped here in a way intended to represent important common themes.

2.4.1 Experiments Using the CASIA v1 Dataset

One cluster of publications compares different texture filter formulations and presents experimental results on the CASIA v1 dataset. The issue with the CASIA v1 dataset that was mentioned earlier – artificial, circular, constant-intensity pupil regions – does not necessarily compromise the use of this dataset in evaluating the performance of algorithms for texture analysis and matching. However, the small size of the dataset and the many papers in the literature that report near-perfect performance on this dataset make it nearly impossible to use it to document a measurable improvement over the state of the art. Therefore, for space considerations, we do not cover this subarea of publications in this survey. Fatt et al. [49, 50] implement a fairly typical 1D log-Gabor iris biometric system on a digital signal processor (DSP) and show results on CASIA v1 dataset. Showing the relative speed of software versus DSP implementations of an algorithm is an example of a context where using the CASIA v1 dataset may be reasonable.

2.4.2 “Eigen-Iris” Approaches

One group of papers might be characterized, by analogy to “eigen-faces” in face recognition, as using an “eigen-iris” approach. Chowhan and Sihinde [33] propose using PCA for iris recognition, in an eigen-face style of approach. Moravec et al. [132] also use a PCA-based approach, with color images of 128 irises. Zhiping et al. [216] use a 2D weighted PCA approach to extracting a feature vector, showing improvement over plain PCA. Chen et al. [29] use 2D PCA and LDA, on UBIRIS images, showing an improvement over PCA or LDA alone. Eskandari and Toygar [47] explore subpattern-based PCA and modular PCA, achieving performance up to 92% rank-one recognition on the CASIA v3 dataset. Erbilek and Toygar [46] look at recognition in the presence of occlusions, comparing holistic versus subpattern-based approaches, using PCA and subspace LDA for iris matching, with experiments on the CASIA, UPOL, and UBIRIS datasets. Xu and Guo [211] propose to extract iris features from the normalized iris image using a method that they call complete 2D PCA.

2.4.3 Alternative Texture Filter Formulations

Many researchers have looked at different mathematical formulations of filters to use in analyzing the iris texture. Patil and Patilkulkarni [140] use wavelet analysis to create a texture feature vector, with experiments on the CASIA v2 dataset. Velisavljevic [198] experiments with the use of oriented separable wavelet transforms, or directionlets, using the CASIA v3 dataset, and shows that they can give improved performance for a larger-size binary iris code. Sun and Tan [187] propose using ordinal features, which represent the relative intensity relationship between regions of the iris image filtered by multilobe differential filters. Krichen et al. [97] explore using a normalized phase correlation approach to matching, as an alternative to the standard binary iris code. They compare results to the OSIRIS [15] and Masek [118] algorithms, on the ICE 2005 and the CASIA-BioSecure iris datasets.

Al-Qunaieer and Ghouti [4] use quaternion log-Gabor filters to analyze the texture of images in the UBIRIS color image dataset and also [57] use a quaternion Fourier transform and phase correlation to improve performance. Bodade and Talbar [17] use a rotated complex wavelet transform in matching iris textures, with experimental results on the UBIRIS dataset, but do not improve recognition performance over the Gabor wavelet. Tajbakhsh et al. [188] present a method of feature extraction based on Ma et al.’s earlier method of analyzing local intensity variation [117] and propose four improvements to the earlier method to make it work with the noisy images in the UBIRIS dataset. Tajbakhsh et al. [189] use a 2D discrete wavelet transform applied to overlapping 32×32 pixel blocks and achieve 0.66% EER on the UBIRIS data.

The motivation behind Miyazawa’s proposed method [128] is that Daugman-like, feature-based iris recognition algorithms require many parameters and that their proposed algorithm should be easier to train. For each comparison using the proposed method, they take two images and select a region that is unoccluded in both images. They take the discrete Fourier transform of both valid regions, then apply a phase-only correlation (POC) function. The POC function involves a difference between the phase components from both images. They use band-limited POC to avoid information from high-frequency noise. The proposed algorithm requires only two parameters: one representing the effective horizontal bandwidth for recognition and the other representing the effective vertical bandwidth. They achieve better results using phase-only correlation than using Masek’s 1D log-Gabor algorithm.

2.4.4 *Alternative Methods of Texture Analysis*

Another group of papers explores texture representation and matching approaches that do not map directly to the typical texture filter framework.

Gray level co-occurrence matrices (GLCM) can be used to describe texture in an image [67]. A GLCM is formed by counting the co-occurrences of brightness values of pixel pairs in the image at a certain distance and direction. Chen et al. [28] propose a modified GLCM based on looking at triples of pixels instead of pairs. They call their modified method a “3D-GLCM” and use it to describe the texture of iris images in the UBIRIS dataset. Using equal error rate, the 2D-GLCM method performs better, but for a FAR of 0%, the 3D-GLCM performs better.

Kannavara and Bourbakis [91] explore using a local-global graph methodology to generate feature vectors, with experiments on color iris images. Sudha et al. [186] compute a local partial Hausdorff distance based on comparing the edge-detected images of two irises, obtaining 98% rank-one recognition on a UPOL dataset representing 128 irises. Kyaw [99] explores using simple statistical features such as mean, median, mode, and variance within concentric bands of the iris but presents no experimental results. Wu and Wang [209] use intensity surface difference between irises for matching and report relatively low performance on the CASIA v1 dataset. Mehrotra et al. [123] use a Harris corner detector to find interest points, which are paired across images for matching. Tests on BATH, CASIA, and IITK datasets indicate that this method does not perform as well as traditional iris code approaches. To avoid aliasing problems from “unwrapping” an iris image, Mehrotra et al. [122] extract features from the annular iris image. They use the SURF algorithm (Speeded Up Robust Features) to identify rotation-invariant features and report recognition accuracy above 97% on BATH, CASIA3, and IITK databases. Radhika et al. [158] use continuous dynamic programming to extract iris texture information. They test their method on CASIA v2 and UBIRIS v1 data. Overall, it appears that none of the various different approaches in this category have yet demonstrated any clear performance improvement over the more traditional texture filtering approaches used in iris biometrics.

Patil and Patilkulkarni [138] describe a comparison of different texture analysis methods for iris matching. They compare the use of statistical measures (mean, median, mode, variance), lifting wavelet transform, and gray-level co-occurrence matrices for deriving texture features. They perform experiments using the CASIA v2 dataset and find that the lifting wavelet transform provides the best recognition accuracy. Patil and Patilkulkarni [139] also explore the use of SIFT features for iris biometrics.

Rathgeb and Uhl [163] develop an approach to iris biometrics that uses contrast-limited adaptive histogram equalization and traces pixel intensity variations along rows of the normalized image (concentric circles of the iris region). These are termed “pixel paths.” They achieve an EER on experiments with the CASIA v3 dataset in the range of 1–2%. They also show how this approach lends itself to cancelable biometrics.

2.4.5 Algorithms that Analyze the Iris in Parts

Several researchers have proposed approaches that analyze the iris region in multiple parts and combine the results. One motivation for this type of approach is to reduce the impact of segmentation errors and noise in the imaging process.

Adam et al. [1] analyze iris texture in eight subregions of the iris and fuse the distances from these local windows, with experiments on data from the CASIA v3 dataset. Bastys et al. [9] divide the iris into sectors and calculate a set number of local extrema in each sector at a number of scales. They achieve perfect separation between genuine and impostor scores for CASIA v1 and CASIA v3 interval, an EER of 0.13% for the CASIA v2 data, and 0.25% for the ICE 2005 data. Garg et al. [53] propose a method that uses a grid on the iris image and a vector of the average pixel values in the elements of the grid for representing and matching the iris texture. Eskandari and Toygar [47] explore subpattern-based PCA and modular PCA, achieving performance up to 92% rank-one recognition on the CASIA v3 dataset. Erbilek and Toygar [46] look at recognition in the presence of occlusions, comparing holistic versus subpattern-based approaches, using PCA and subspace LDA for iris matching, with experiments on the CASIA, UPOL, and UBIRIS datasets. Lin et al. [109] divide the iris area into 4 local areas and the face into 16 local areas in their approach to iris and face multi-biometrics.

Campos et al. [25] propose an alternative method of feature extraction. They apply histogram equalization and binarization to the unwrapped iris image and use a self-organizing map neural network to divide the binary image into nodes. From the topological graph of the image, they compute corresponding Voronoi polygons. Next, they calculate the mean, variance, and skewness of the image in each polygonal region. They achieve 99.87% correct recognition on the Bath University iris data.

Rachubinski [157] presents a method of feature extraction using wavelet coefficients based on a wedgelet dictionary. A wedgelet is a division of a square region

into two sections. The wedgelet is parameterized by the distance of the segment from the center of the square and the angle of the segment dividing the two regions. Rachubinski divides the unwrapped iris image into overlapping local regions of 8×8 pixels and determines a wedgelet dictionary for each region. The wedgelet angles are quantized to create a binary iris code, and codes are compared using Hamming distances. Rachubinski achieves 100% rank-one recognition rate (0.15% EER) on the relatively non-challenging CASIA v1 dataset.

Don et al. [38] present what is termed a “personalized iris matching strategy.” A weight map is learned for the features in the image of each given iris, based on training images of that iris. This is conceptually similar to the “fragile bits” work of Hollingsworth. This approach is said to be especially useful in the case of poor-quality iris images.

2.4.6 Approaches to Speed Iris Matching

Hao et al. [66] present a technique to speed up the search of a large database of iris codes, with experiments that use over 600,000 iris codes from the ongoing application for border control in the United Arab Emirates. They use a “beacon-guided search” to achieve a “substantial improvement in search speed with a negligible loss of accuracy” in comparison to an exhaustive search.

Gentile et al. [56] experiment with generating a shorter iris code that maintains recognition power and conclude that it is best to focus on the middle radial bands of the iris and to sample every n -th band. Gentile et al. [55] also use a short-length iris code to index into a large iris dataset to reduce the total number of iris code comparisons to search the dataset, with a small degradation in recognition rate.

Roy and Bhattacharya [172–176] reduce matching time by applying feature selection to choose the most discriminating features. They explore the use of genetic algorithms to select a subset of most useful features for iris matching [172, 173]. In [175], they use support vector machine-recursive feature elimination (SVM-RFE). In [174], they apply a genetic algorithm to select important features and use an iterative algorithm, called the contribution-selection algorithm, from the field of coalitional game theory, to reduce the feature vector dimension.

Mehrotra et al. [124] propose an indexing algorithm to reduce the search time. They divide each unwrapped iris image into subbands using a multiresolution discrete cosine transform. They create a histogram of transform coefficients for each subband using all the images in the database. They use histograms containing about ten bins each. The algorithm forms a key for each image from noting the bin numbers associated with the subbands of the image. The keys are organized into a search tree. To search for a match to a new image, the algorithm computes the key for the new image, retrieves all irises with matching keys from the database, and compares iris templates from the retrieved set. They achieve a bin miss rate of 1.5% with a penetration rate of 41%.

Rathgeb et al. [165] present an approach to “incremental” iris code matching, with the aim of reducing the number of bit comparisons used per recognition result. It is claimed that “the proposed technique offers significant advantages over conventional bit-masking, which would represent binary reliability masks.”

2.4.7 Exploiting “Fragile” Bits in the Iris Code

Hollingsworth et al. [78] describe the concept of “fragile bits” in the traditional Daugman-style iris code. Bits in the iris code can be fragile due essentially to random variation in the texture filter result, causing them to “flip” between 0 and 1. Recognition performance can be improved by masking an appropriate fraction of the most fragile bits. Dozier et al. [39] use a genetic algorithm to evolve a mask for the iris code that best masks out the “fragile” iris code bits. Hollingsworth et al. [80] describe an approach to averaging the iris image through multiple frames of video, prior to generating the iris code, to improve recognition performance. This approach is effectively reducing the fragility of the bits in the iris code. Hollingsworth et al. [79] also describe an approach to using the spatial coincidence of the fragile bits in the iris code to improve recognition performance.

2.4.8 Use of “Sparse Representation” Techniques

Pillai et al. [148] explore the use of sparse representation techniques for iris biometrics. This approach involves having a number of training images per iris, where the images span the range of different appearances that the iris might have. An unknown iris is then recognized by solving a minimization problem that finds a representation of the unknown image in terms of the training images.

2.5 Multi-biometrics Involving the Iris

The term “multi-biometric” is used to refer to techniques that use more than one biometric sample in making a decision. Often the samples are from different sites on the body, for example, iris and fingerprint. Also they might be from different sensing modalities, for example, 3D and 2D. Or they might be repeated samples from the same sensor and site on the body. The motivation for multi-biometrics is to (a) increase the fraction of the population for which some usable sample can be obtained, (b) increase recognition accuracy, and/or (c) make it more difficult to spoof a biometric system. In India’s unique ID program [169], in many ways the most ambitious biometrics application in the world to date, iris and fingerprint are used primarily, it seems, to increase coverage of the population.

Most multi-biometric work involving the iris has looked at combining iris with some other biometric site, rather than multiple sensing modalities for iris or repeated iris samples. Papers have been published looking at almost any combination of iris and some other modality that one can imagine. Often the practical motivation for the particular pairing is not clear. The vast majority of this work has used *chimera* subjects, that is, virtual subjects created by pairing together biometric samples from already existing unimodal datasets. For example, several papers use iris images from a CASIA dataset and face images from the ORL [102] dataset. In general, there is a need for research in this area to progress to using true multi-biometric datasets, to use datasets representing a much larger number of subjects and images than in the ORL face dataset or the CASIA v1 iris dataset, and to compare performance of the multi-biometric approach to performance of state-of-the-art algorithms for the individual biometrics. In the summaries below, we have tried to explicitly note the few instances where the dataset used was not chimeric.

Perhaps naturally, the largest cluster of papers in this area deals with the combination of face and iris. This group of publications is multi-biometric both in the sense of combining iris and face and often also in the sense of using near-infrared illumination (for iris) and visible light (for face). Lin et al. [109] generalize the posterior union model (PUM) to perform face and iris multi-biometrics, constructing chimera subjects from the XM2VTS or AR face datasets and the CASIA iris dataset and dividing the normalized face images into 16 local areas and the iris area into 4 local areas. Gan and Liu [52] apply a discrete wavelet transform to face and iris images and use a kernel Fisher discriminant analysis, with chimera subjects created from the ORL [102] face database and (apparently) the CASIA v1 iris database. Wang et al. [203, 205] use a complex common vector approach to face and iris, using the ORL and Yale face datasets and the CASIA v1 iris dataset. Liu et al. [110] experiment with a 40-person chimera dataset made from ORL face images and CASIA iris images, with relatively low performance. Wang et al. [204] fuse face and iris information at the feature level. They create a complex feature vector from the real-valued iris feature vector and the real-valued face feature vector. Next, they use complex Fisher discriminate analysis (CFDA) to maximize the between-class scatter with respect to the within-class scatter. They test their algorithm on CASIA v1 iris images and ORL and Yale face images. Wang and Han [199] fuse information from face and iris at the score level. The scores from the two different algorithms are normalized using two sigmoid functions, and then they employ a SVM-based fusion rule to obtain a final score. They test their method using faces from the ORL dataset and irises from UBIRIS dataset. Breitenbach and Chawdhry [23] perform experiments looking at image quality factors for an image and how they predict performance of face and iris recognition. Rattani and Tistarelli [166] fuse information from face and iris at the feature level. They divide the images into windows and extract one SIFT feature from each window. They obtain feature vectors of length 128 each from the face, right eye, and left eye images and find that a fusion of face, right iris, and left iris gets better performance than any one or any fusing of two. Morizet and Gilles [133] use data from the FERET face dataset and a CASIA iris dataset in presenting a method that develops a user-specific fusion of scores from the two modalities.

Vatsa et al. [197] consider approaches based on multiple iris samples. They use elements of belief function theory for iris-based multi-biometrics and look at two scenarios: combining results from enrolling one iris with two images and combining results from the left and right iris each enrolled with one image.

A broad variety of other multi-biometric combinations involving the iris have been studied. Several researchers have looked at fingerprint and iris. Baig et al. [6] investigate iris and fingerprint fusion using the Masek algorithm and a SUNY-Buffalo algorithm, respectively, experimenting on a West Virginia University dataset. It is noted that performance is relatively low, due to design for a “small memory footprint real-time system.” Ross et al. [171] explore multi-biometric iris and fingerprint where fusion is used only in certain cases within the Doddington Zoo framework, experimenting with a chimera dataset of fingerprints from a WVU dataset and irises from a CASIA dataset. Elmadani [45] presents the “fingerIris” algorithm for combination of iris and fingerprint. The approach is evaluated on a true multi-biometric dataset representing 200 individuals. The system gets four to five false reject and/or false accept results on this dataset, depending on the setting of the decision threshold.

Wang et al. [201] explore score-level fusion of iris matching and palm print matching using an apparently chimera dataset representing 100 persons. Tayal et al. [13, 192, 193] use a wavelets approach to analyze iris texture and speech samples for multi-biometrics. Sheela et al. [184] experiment with iris and signature, using CASIA v2 and MYCT datasets, respectively, but do not focus on multi-biometric combination. Mishra and Pathak [127] explore wavelet analysis of iris and ear images for multi-biometrics on a chimera dataset representing 128 persons.

Poh et al. [150] report on multi-biometric research involving face, iris, and fingerprint, carried out as part of the BioSecure project. This project particularly looks at quality-dependent fusion at the score-level and cost-sensitive fusion at the score level. A total of 22 fusion systems were evaluated in this project.

Maltoni et al. [116] discuss pros and cons of fusing multiple biometrics. Generally, fusing more classifiers improves performance if the classifiers are not highly correlated. However, extra classifiers can increase cost and throughput time of the system. Maltoni et al. discuss performing fusion at the image, feature, score, rank, or decision level.

Hollingsworth et al. [76] present an approach that uses multiple iris samples taken using the same sensor, taking advantage of temporal continuity in an iris video to improve matching performance. They select multiple frames from an iris video, unwrap the iris into polar form, and then average multiple frames together. They find that this image-level fusion yields better matching performance than previous multi-gallery score fusion methods.

Conti et al. [34] give an overview of concepts and terminology in multi-biometric systems. They also present an approach to using fuzzy logic methods for score fusion in a multi-biometric system.

Zuo et al. [221] investigate the possibility of matching between a visible-light image and a NIR image of the iris. They formulate a method to estimate the NIR

iris image from a color image. It is claimed that this approach “achieves significantly high performance compared to the case when the same NIR image is matched against R (red) channel alone.”

2.5.1 Ocular Biometrics

The papers covered in this section deal with “ocular” biometrics as a possible multi-biometric complement to iris. An ocular biometric is one based on features of the region of the face around the eye. Much of this research uses ocular regions cropped from visible-light images, often from the Face Recognition Grand Challenge (FRGC) face image dataset. Xu et al. [212] use local binary pattern (LBP) texture features computed over the ocular region. In experiments with images from the FRGC dataset, they achieve 61% verification rate at 0.1% false accept rate. Miller et al. [126] also propose a method using LBP texture features, again using images cropped from the FRGC database. They investigate the effects of image blur, resolution of the periocular region, illumination effects, and different color bands. Lyle et al. [114] present an approach to predicting the gender and ethnicity of a person using LBP features and an SVM classifier. In experiments with images from the FRGC dataset, they obtain 93% accuracy on gender classification and 91% on Asian/non-Asian ethnicity classification. Bharadwaj et al. [11] present a method of ocular recognition with experiments using the UBIRIS iris images. Their method uses the GIST global descriptor and LBP texture features. Merkow et al. [125] predict the gender of the subject based on features computed from the ocular region and obtain 85% correct gender prediction using frontal-view color face images taken from the web.

Hollingsworth et al. [81] study how human observers rate the value of different features of the ocular region for recognition. This study was done with NIR images from the LG 2200 iris sensor. Thus, this investigation is more directly relevant to ocular as a complement to iris and less directly relevant to ocular as a subset of face recognition using visible-light images.

2.6 Privacy and Security

This section includes several somewhat different areas of work. The development of privacy-enhancing techniques generally involves rigorous conceptual or mathematical approaches. More general security techniques look at integrating biometrics into encryption schemes in some way. The study of liveness detection, or spoofing and anti-spoofing, often involves clever exploitation of sensor capabilities.

Ratha [161] gives a broad perspective on security and privacy issues in large-scale biometric systems. Taking a system-level view of biometric authentication, he considers the various possible attack points. He also summarizes the concept

of cancelable biometrics as a means to enhance privacy and security. This is a good general article for someone who is not already familiar with basic concepts in this area.

2.6.1 Privacy-Enhancing Techniques

The area of privacy-enhancing techniques for biometrics is challenging and fast moving. Its importance is perhaps not yet fully understood and appreciated by the field as a whole. One can see the importance of this area by considering what would happen in a biometric-enabled application when a person's biometric template is stolen. The application needs some way to protect each individual's biometric template and/or to be able to revoke an enrollment in the application and reenroll a person.

Several authors have proposed encryption methods to protect the privacy of a biometric template. Luo et al. [113] propose to perform anonymous biometric matching, using encryption to protect the probe biometric. Alghamdi et al. [5] propose using the iris code to generate a key for encryption of the iris image or other data. Moi et al. [130] propose using AES encryption of an enrolled iris code to store the key to encrypted documents.

Li and Du [105, 106] propose watermarking the iris image at the time that it is acquired by the sensor, as a means to later determine the authenticity of the image. This would in principle allow detection of an image that did not originate with the particular sensor.

Tan et al. [191] propose an "image hashing" technique, which converts the iris biometric into a short bit string in a manner that is irreversible. That is, given the short bit string, it is not possible to generate the iris biometric.

Agrawal and Savvides [3] describe an approach to hiding an iris biometric template in a host image. Their steganographic approach is designed to cause imperceptible change in the host image and to be robust to JPEG artifacts.

Adjedj et al. [2] describe a way to create a biometric identification scheme while storing only encrypted data. Their method uses symmetric searchable encryption which is a technique allowing a server to return all documents containing a particular keyword without learning anything about the keyword. They also use a family of locality sensitive hashes.

Hämmerle-Uhl et al. [63] propose a cancelable biometrics technique for irises. Cancelable biometrics are transformations of the original biometric that can be used for authentication without revealing the original, unaltered biometric, thus improving privacy for the user. In a cancelable biometric system, if a user's biometric is stolen, it can be canceled and reissued. They suggest two types of transformations. One proposed transformation is to randomly remap blocks of iris texture to create a new signal. A second proposed transformation is to warp the texture along a grid with randomly offset vertices. Färberböck et al. [48] present an approach to transforming rectangular and polar iris images to enable cancelable iris

biometrics. They experiment with block remapping and texture warping techniques for this purpose, using images from the CASIA v3 iris image dataset. Kanade et al. [88, 89] propose a two-factor approach to cancelable biometrics. Their proposed system uses an iris biometric and a password. In addition, their system uses an error-correcting-code technique and a user-specific shuffling key to increase the separation between the genuine and impostor distributions.

2.6.2 Security

Zhang et al. [214] propose a method to bind cryptographic keys to biometric data. During enrollment, they use Reed-Solomon coding and convolutional coding to add error-correcting data to a random key. They XOR the random key with the iris code, and produce helper data that hides the biometric and the key. During verification, the new iris code is XORed with the helper data, and then Reed-Solomon and Convolutional coding is used to decode the bit-string and correct errors, thus unlocking the original cryptographic key. This method is similar to the method proposed by Hao et al. [65].

Rathgeb and Uhl [162] describe how to construct an iris-based fuzzy commitment scheme to hide and retrieve a cryptographic key. Like [65], they use Reed-Solomon and Hadamard error-correcting codes. However, they show how to extend this scheme to an arbitrary iris biometrics algorithm.

Rathgeb and Uhl [163] discuss the problem of generating cryptographic keys from iris biometric samples. Their proposed approach uses an interval mapping technique and does not store biometric data in either raw or encrypted form. On experiments with the CASIA v3 dataset, they are able to obtain key generation rates as high as 95% using five enrollment samples.

Mahmud et al. [115] present a stream-cipher method that uses an iris code as an initial input to seed a linear feedback shift register (LFSR). The LFSR is used to implement a stream cipher. Since biometric templates are not identically repeatable, their system stores the initial biometric key on a smart card, which is programmed to release the key only when a similar biometric template is presented to unlock the smart card. The authors claim that their method is stronger than other ciphers like A5/1 and RC6.

Plaga [149] computes the theoretical maximal achievable information content of biometric keys. A biometric template, such as the iris code proposed by Daugman, may have a length of 2,048 bits. However, there are correlations in the bits, so in actuality, the information content in the template is smaller, for this example, 249 bits. Even so, a cryptographic key must necessarily be even shorter because some number of bits are required for error correction. The number of bits required for error correction is a function of the number of bit errors between two templates from the same biometric feature. Using numbers provided by Daugman, Plaga determines that the maximum error-free and correlation-free biometric key has length 25 bits. Using numbers from a performance study conducted in the

Frankfurt International Airport, Plaga determines that even fewer bits are available for biometric keys derived from face, fingerprint, and iris systems under airport-type operating conditions. Plaga concludes that “current commercial state-of-the-art biometric systems based on a single biometric feature like one finger or iris create templates from which no more than about 30 bits can be derived.” Therefore, in order to use biometrics to create keys, either the performance of the systems must be substantially improved or the systems must employ multimodal or multi-instance biometrics (e.g., ten-print fingerprints).

Rathgeb and Uhl [164] consider the operation of two-factor authentication systems in which one of the factors is iris biometrics. They illustrate empirically how this helps to increase the separation of the authentic and impostor distributions relative to iris biometrics alone. They point out that the increased recognition accuracy in the two-factor system is based on the assumption “that additional factors are considered to never be stolen, lost, shared or duplicated where in practice the opposite is true” and discuss requirements for performance analysis of two-factor systems where one of the factors is a biometric.

2.6.3 *Liveness Detection (Anti-spoofing)*

Ruiz-Albacete et al. [178] explore “direct attacks” on an iris biometric system, in which a printed image of an iris is presented to an iris biometric system in an attempt to enroll an iris and/or to match an enrolled iris. They find that with appropriate choice of commercial printer, printer paper, and image processing algorithm, they are able to generate printed iris images that are enrolled and/or matched by the iris biometric system with substantial rates of success. The particular iris biometric system used in the experiments is the LG IrisAccess 3000, a model that is no longer marketed. It is not clear that the experience with this system could easily be replicated with current commercial iris biometric systems, as current commercial systems may incorporate some sort of liveness detection that should defeat simple spoof attempts using paper-printed iris images.

Bodade and Talbar [16] propose an approach using multiple images of the same eye to look at variation in pupil dilation in order to detect iris spoofing. Takano and Nakamura [190] describe a neural network approach to iris recognition and to detecting “live” iris versus iris patterns printed on paper with experiments on a limited dataset representing 19 persons.

He et al. [74] aim to detect certain types of spoofs by detecting printed contact lenses. They consider three subregions on the right side of the iris and three on the left. They analyze texture in each subregion using local binary patterns (LBPs) at multiple scales. Gaussian kernel density estimation is applied to complement the insufficiency of counterfeit iris images. They train an AdaBoost classifier and select 85 LBP bins to use in testing. The proposed method achieves lower error rates than previous methods [69, 206].

He et al. [73] research detection of blurry, spoofed images. They note that Daugman’s method of computing the FFT [35] can only detect printed contacts

with high frequency, but it would fail if the spoofed pattern was partially blurred. He et al. use wavelet packet decomposition to perform wavelet packet decomposition and then employ a support vector machine to classify irises as live or spoofed. Their method correctly detects 98.6% of the spoofed images in their dataset.

2.7 Datasets and Evaluations

Datasets and evaluations play a large role in biometrics research. The widespread availability of common datasets has enabled many researchers to enter the field and demonstrate results whose relevance can be more easily understood due to the use of a known dataset. Evaluation programs have given researchers an idea about the current state of the art and helped to focus and shape research to address the interests of sponsoring agencies.

Proenca et al. [155] describe the UBIRIS v2 dataset of visible-light, color iris images, acquired with 4–8 m distance between subject and sensor and with subjects in motion. The dataset represents 261 subjects, with over 11,000 iris images. The purpose of the dataset is to support research on visible-light iris images acquired under far from ideal imaging conditions [155].

Johnson et al. [85] describe the “Q-FIRE” dataset of face and iris videos. These videos represent variations in focus blur, off-angle gaze, and motion blur and are acquired at a range of 5–25 feet. This dataset is potentially useful for research in iris, face, and multi-biometric face + iris.

Fierrez et al. [51] describe a multi-biometrics dataset acquired as part of the BioSecure project. The dataset represents 400 persons, with biometric samples for speech, iris, face, handwriting, fingerprints, palm print, hand contour geometry, and keystroking. The iris images are acquired with an LG IrisAccess EOU 3000 and include four samples per eye with subjects not wearing eyeglasses and the presence of contact lenses recorded.

Ortega-Garcia et al. [136] describe a larger and more varied version of the multi-biometrics dataset resulting from the BioSecure Network of Excellence. This version contains biometric data representing more than 600 individuals. The data represents three different scenarios: “(i) over the Internet, (ii) in an office environment with desktop PC, and (iii) in indoor/outdoor environments with mobile portable hardware.” Again, the iris part of the dataset was acquired using an LG IrisAccess EOU 3000. The total dataset involved the efforts of 11 institutions. The iris portion of the dataset represents 667 persons, with two acquisitions per person and two images of each iris in each session.

Schmid and Nicolo [180] suggest a method of analyzing the quality of an entire database. They compare the capacity of a recognition system to the capacity of a communication channel. Recognition channel capacity can be thought of as the maximum number of classes that can be successfully recognized. This capacity can also be used as a measure of overall quality of data in a database. The authors

evaluate the empirical recognition capacity of biometrics systems that use PCA and ICA. They apply their method to four iris databases and two face databases. They find that the BATH iris database has a relatively high sample signal-to-noise ratio, followed by CASIA-III, then ICE 2005. WVU had the lowest signal-to-noise ratio.

Krichen et al. [98] give a brief introduction to the open-source iris recognition system, OSIRIS. They also describe their BioSecure iris database, which they combine with the CASIA v2 data to create a database with equal numbers of Asian and European subjects. They test the OSIRIS system on the CASIA-BioSecure data and also on the ICE 2005 data and show that the OSIRIS system outperforms the Masek open-source system.

Phillips et al. [147] describe the results of the Face Recognition Vendor Test 2006 and the Iris Challenge Evaluation 2006. These evaluations follow on the Face Recognition Grand Challenge and the Iris Challenge Evaluation 2005. The ICE programs resulted in a dataset of over 64,000 iris images from over 350 subjects, acquired using an LG 2200 iris sensor in 2004 and 2005, being made available to the research community [20]. The dataset contains both “ideal” images and “poor-quality” images. The ICE programs also resulted in the source code of a baseline Daugman-like system being made available to the research community.

Petrovska et al. [141] describe the BioSecure benchmarking methodology for evaluating performance of biometric algorithms. The BioSecure reference system provides open-source software, publicly available biometric databases, and evaluation protocols that allow researchers to conduct reproducible research experiments. The book chapter explains the need for a common benchmarking methodology and summarizes the frameworks. Frameworks for eight different biometric modalities are available: iris, fingerprint, signature, hand geometry, speech, 2D face, 3D face, and talking face.

The US government has organized a number of biometrics challenge problems and evaluations to motivate advancements in biometric technology. Phillips et al. [146] describe the data available in the Multiple Biometrics Grand Challenge (MBGC). The MBGC includes three different challenge problems, one of which involves iris recognition: the portal challenge problem. The goal of the portal challenge problem is to recognize people from near-infrared and visible-light video as they walk through a portal. Five different types of data are provided as part of the portal challenge: (1) still iris images from an LG2200 sensor; (2) video iris images from an LG2200 sensor; (3) medium-resolution, still, frontal face images; (4) high-resolution NIR video acquired from a Sarnoff Iris on the Move (IOM) system; and (5) high-definition, visible-light video acquired at the same time as the IOM videos. MBGC version 1 data was released in May 2008. MBGC version 2 data was released in February 2009.

Newton and Phillips [135] present a meta-analysis of three iris biometric evaluations: the Independent Testing of Iris Recognition Technology performed by the International Biometric Group, the Iris Recognition Study 2006 conducted by Authenti-Corp, and the Iris Challenge Evaluation 2006 conducted by the National

Institute of Standards and Technology. The meta-analysis looks at the variation across the three studies in the false non-match rates reported for a false match rate of 1 in 1,000.

2.8 Performance Under Varying Conditions

Some early folklore of the iris biometrics field held that pupil dilation, contact lenses, and template aging do not negatively impact iris biometrics. Bowyer et al. [22] test these assertions. They show that iris biometric performance can be degraded by varying pupil dilation, by wearing non-cosmetic or cosmetic contact lenses, and by time lapse between enrollment and verification. They also show that using a different sensor between enrollment and verification can degrade performance. These factors primarily affect the match distribution, while the non-match distribution remains stable. Thus, for a verification scenario, the false accept rates are unaffected by these factors. For a watchlist scenario, however, operators should be aware that suspects may attempt to fool the system by, for example, artificially dilating their eyes or wearing contacts.

Baker et al. [8] look at how contact lenses affect iris recognition, with the conclusion that even normal prescription contacts can cause an increase in the false reject rate. The size of the increase in the false reject rate varies greatly across different matching algorithms and different types of contact lenses. In general, the effects of contact lenses on iris biometrics accuracy seem not yet fully understood.

Rankin et al. [160] explore effects of pupil dilation using images from three subjects taken over a period of up to 24 weeks under varying pupil dilation conditions, using a biometric slit lamp. Some unusual results are obtained on applying a version of an early Daugman algorithm and Masek's algorithm to these images. However, results generally agree with those of previous researchers that found that pupil dilation increases the false reject rate [62, 77].

Gonzaga and da Costa [58] propose a method to exploit the "consensual reflex" between a person's irises to illuminate one eye with visible light to control the dilation of both pupil, and image the other eye with NIR illumination. In this way, they can compute features of the iris over dilation.

Baker et al. [7] explore the effects of time lapse on iris biometrics. They compare the average Hamming distance between images taken 4 years apart with the average Hamming distance between images taken within a single semester. They find statistically significant evidence that the distance scores between images taken years apart are greater than the distance scores from images taken within a few months of each other. Using the IrisBEE iris matcher, they observe an approximate 0.018 increase in Hamming distance for matches with a 4-year time lapse. The increased false reject rate for the long-time-lapse matches relative to the short-time-lapse matches indicates that a template aging effect exists for iris biometrics. This was the first study to make any rigorous experimental evaluation of the issue of template aging for iris biometrics.

Borgen et al. [19] investigate the effects of common ocular diseases on iris recognition. They use the UBIRISv1 dataset and simulate different pathologies. All simulated pathologies were validated by ophthalmology and optometry specialists. Changes in iris color, scars from glaucoma surgery, and vessel growths caused only small increases in the false reject rate. Corneal bleaching and scarring caused a false reject rate of 86.8%. The corneal bleaching caused segmentation of the outer iris boundary to fail in many cases. Central keratitis increased the false reject rate of bright-eyed subjects more than dark-eyed subjects. High-density infiltrates caused more problems with dark-eyed subjects. The authors conclude that iris recognition is robust for some pathologies but that others – such as corneal bleaching – can unacceptably damage the false reject rate in just 3 months of disease progression.

2.9 Applications

A small number of publications have appeared which envision the use of iris biometrics in particular application scenarios. One interesting aspect of this group of papers is the very broad range of uses envisioned for biometrics, almost none of which involve national security.

Kadhum et al. propose using iris biometrics to authorize entry through doors to secure areas, an application for which commercial iris biometric systems already exist (e.g., LG Iris). Mondal et al. [131] propose using biometrics for secure access to home appliances over the network. Iris biometrics is used in this chapter, but the approach can potentially be extended to other biometrics. Garg et al. [54] propose a vision system that will recognize a set of hand gestures to control devices and use iris biometrics to authenticate the user identity. Leonard et al. [104] propose using fingerprint, iris, retina, and DNA (“FIRD”) to distinctively identify a patient to his or her complete electronic health care record. Mohammadi and Jahanshahi [129] propose an architecture for a secure e-tendering (offering and entering into a contract) system, with iris as the example biometric for identity verification. Wang et al. [202] propose using Daugman-like iris biometrics “to make the large animals be recognizable and traceable from the farm to the slaughterhouse,” furthering the goal of food chain safety. Wang et al. [200] propose to use face and iris multi-biometrics as part of a scheme to enforce digital rights management, which would allow only authorized remote users to access content. Hassanien et al. [68] show how an iris template can be embedded in a digital image to prove ownership of the image.

Dutta et al. [40–42] propose embedding the iris code of a person in an audio file as a watermark to prove ownership of the audio file. They apply Haar wavelets at four levels of decomposition to create a feature vector from an iris image. Next, they binarize the feature vector by comparing each element of the vector to the median value in the vector. This process creates a bio-key with power evenly distributed throughout the audio spectrum, thus allowing the key to be embedded in the audio signal without affecting listeners. They test their method by embedding bio-keys in five different musical samples, then comparing the embedded keys with all iris

keys in their database. A high correlation between the embedded key and the stored key is evidence of a match. Their method is robust to various types of attack on the audio signal.

2.9.1 Hardware Implementations

Liu-Jimenez et al. [112] and Rakvic et al. [159] describe the implementation of iris biometric algorithms on FPGAs. Zhao and Xie [215] describe an implementation of an iris biometric system on a DSP. Vandal and Savvides [195] present results of iris matching parallelized for execution on graphics processing units and report a 14-times speedup relative to state-of-the-art single-core CPUs.

Jang et al. [84] describe the design and implementation of a “portable” or handheld iris biometric sensor. The heart of the system is an “ultra-mobile personal computer,” the Sony model VGN-UX17LP. The system uses a near-infrared illuminator and a CCD camera with a fixed-focus zoom lens. An image restoration algorithm is used to increase the effective capture volume, which is claimed to exceed that of the PIER 2.4 and the HIDE systems.

Kang and Park [90] describe an iris biometrics system implemented to operate on a mobile phone. The system repeatedly takes images of both eyes and performs a quality assessment until at least one image passes the quality assessment check. Then it performs authentication either with one image or with score-level fusion of two images.

2.10 Theoretical Analyses

There are relatively few studies that might be considered theoretical analyses of fundamental issues in the field. Bhatnagar et al. [12] develop a theoretical model for estimating the probability of random correspondence of two iris codes and compare this with the analogous value for a pair of palm prints. Kong et al. [93] undertake a theoretical analysis of the Daugman-style iris code representation of iris texture. One interesting element of this is a discussion of the impostor distribution as an instance of the binomial distribution.

Gorodnichy and Hoshino [60] develop a score calibration function that can convert match scores into probability-based confidence scores. They present a theoretical argument and also supporting experimental results to show that this approach results in the best possible detection error tradeoff curves. The calibration that is effected is meant to ensure that “...the statement ‘I am 60% sure that this person is Alice’ is correct exactly 60% of the time.”

Yager and Dunstone [213] tested for the existence of “Doddington Zoo” animals in a number of different biometric databases, using a number of biometric algorithms. Each of the animal types was present in some of the experiments and absent

in others. The authors note that “The reasons that a particular animal group exist are complex and varied. They depend on a number of factors, including enrollment procedures, feature extraction and matching algorithms, data quality, and intrinsic properties of the user population” [213]. Their analysis also leads the authors to assert that people are rarely “inherently hard to match.” Instead, they suggest that matching errors are more likely due to enrollment issues and algorithmic weaknesses rather than intrinsic properties of the users.

Stark et al. [185] conduct experiments in which human observers view iris images and categorized them into groups of similar-appearing texture pattern. The results suggest that there are a small number of generally agreed-upon texture categories. The results also suggest that texture categories may be correlated with ethnicity, although the iris textures in the experiment all represent either Asian or Caucasian ethnicity, and so greater variation in ethnicity remains to be examined.

2.11 Discussion

In this section, we give eight “recommended reading” suggestions. This is not meant as a best papers list but rather as a list of papers representing interesting and/or unusual viewpoints and directions in iris biometrics.

Gorodnichy’s paper “Evolution and evaluation of biometric systems” [59] is a worthwhile read for those who want to get a sense of how biometric technology is evolving, how the performance of biometric technology is evaluated, and an introduction to much basic biometric terminology. Gorodnichy is Senior Research Scientist with the Canadian Border Services Agency, and so he brings a systems and application-oriented viewpoint to the task of evaluating biometric technology. He particularly makes the point that biometric systems are not fielded in a static context, but that the mix of data and challenges that they must handle naturally evolve over time, and so the biometric technology must evolve as well.

Current commercial iris biometric systems all, to our knowledge, use near-infrared illumination in the 700–900 nm wavelength range. There is also a body of iris biometric research based on visible-wavelength images. But there is almost no published work on imaging the iris outside of the 700–900 nm range. For this reason, the paper by Ross et al. [170], “Exploring multispectral iris recognition beyond 900 nm,” is unique. It remains to be seen whether or not it will be technically and economically viable to image the iris at multiple wavelengths and/or to match iris texture across wavelengths. For those who are intrigued by the topic, this paper is a good introduction. This is likely an area that will see increased attention in the future.

To our knowledge, the paper by Chou et al. [32], “Non-Orthogonal View Iris Recognition System,” is the only system proposed to simultaneously acquire both a visible-light image and a near-infrared image of the iris. They exploit the two images in a complementary manner in the segmentation stage, using the color image to aid in finding the limbic boundary. For anyone interested in multi-biometrics, the

relative simplicity of the sensor design and the method of exploiting the two images should be interesting and suggest additional possibilities.

Proenca's paper, "On the Feasibility of the Visible Wavelength, At-a-Distance and On-the-Move Iris Recognition" [152], is interesting because it argues that visible-light imaging is the way to go, especially for imaging "at a distance" and "on the move." This argument runs counter to the approach used by all commercial systems that we are aware of and also counter to the majority of academic research. However, because it does represent a "contrarian" sort of approach, those interested in the illumination issue for iris biometrics should find this chapter worthwhile.

The paper by Pillai et al., "Sparsity inspired selection and recognition of iris images" [148], is the first that we know of to try to transfer the excitement about sparse representation techniques in the face recognition community over to iris recognition. Extraordinary recognition performance has been claimed for face recognition systems using sparse representation techniques. A potential weakness of using a sparse representations approach is the requirement for a large number of training images per iris and that the images should span the range of different possible appearances. It remains to be seen whether or not sparse representation techniques will revolutionize either face or iris recognition in practice, but this chapter is a good starting point for how the concepts could be applied in iris recognition.

The paper by Vatsa et al. [197], "Belief Function Theory Based Biometric Match Score Fusion: Case Studies In Multi-instance and Multi-unit Iris Verification," is interesting as an example for what it terms "multi-instance" and "multiunit" iris biometrics. Multi-instance refers to using multiple images of the same iris, either to enroll a person in the system and/or as a probe to be matched for recognition. Multiunit refers to using an image of both irises rather than a single iris. Early iris biometric systems seem to have all enrolled a person using a single iris biometric template formed from a single image. This chapter shows that there are simple ways of increasing recognition performance by using multiple images.

For anyone not already familiar with the concept of cancelable biometrics, the paper by Kanade et al., "Cancelable Iris Biometrics and Using Error Correcting Codes to Reduce Variability in Biometric Data" [88], should be worth reading. In this particular instance, they propose a two-factor approach to cancelable biometrics. The two factors are the biometric and the password. If needed, a person's current enrollment in a biometric system using this scheme can be canceled, and then the person reenrolled with a new password. This particular proposed system also uses the password to effectively increase the separation between the genuine and impostor distributions.

Zuo and Schmid's paper, "Global and Local Quality Measures for NIR Iris Video" [219], provides a good introduction to the complexity of the problem of evaluating the quality of an iris image. For a single iris image, they compute nine different quality metrics, for segmentation quality, interlacing, illumination contrast, illumination evenness, percent occlusion, pixel count, dilation, off-angle view, and blur. Quality metrics concerned with interlacing will presumably not be important in the future, as iris images will be acquired digital rather than digitized from analog

video. But the problem is actually even more complex than it appears here. For example, the focus quality of an image is not necessarily even over the entire iris. Also, it is not only the dilation of a single image that is important but the difference in dilation between two images that are being matched [77].

This group of eight papers that touch on very different topics in the field of iris biometrics research should convey a sense of the breadth of the field. It should also help to convey a sense of the excitement in the field, in that there are many directions being explored that could serve to increase accuracy of, and/or increase the breadth of application of, iris biometrics.

Acknowledgements The authors were supported by the Federal Bureau of Investigation, the Central Intelligence Agency, the Intelligence Advanced Research Projects Activity, the Biometrics Task Force, the Technical Support Working Group under US Army contract W91CRB-08-C-0093, and the Intelligence Community Postdoctoral Fellowship Program.

The opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of our sponsors. The identification of any commercial product or trade name does not imply endorsement or recommendation by the authors, the University of Notre Dame, or the National Institute of Standards and Technology.

References

1. Adam, M., Rossant, F., Mikovicova, B., Amiel, F.: Iris identification based on a local analysis of the iris texture. In: Proceedings of 6th International Symposium on Image and Signal Processing and Analysis (ISPA), pp. 523–528 (2009)
2. Adjedj, M., Bringer, J., Chabanne, H., Kindarji, B.: Biometric identification over encrypted data made feasible. In: Information Systems Security: Lecture Notes in Computer Science #5905, pp. 86–100. Springer-Verlag Berlin Heidelberg (2009)
3. Agrawal, N., Savvides, M.: Biometric data hiding: A 3 factor authentication approach to verify identity with a single image using steganography, encryption and matching. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops), pp. 85–92 (2009)
4. Al-Qunaieer, F.S., Ghouti, L.: Color iris recognition using hypercomplex gabor wavelets. In: Symposium on Bio-Inspired Learning and Intelligent Systems for Security (BLISS '09), pp. 18–19 (2009)
5. Alghamdi, A.S., Ullah, H., Mahmud, M., Khan, M.K.: Bio-chaotic stream cipher-based iris image encryption. In: International Conference on Computational Science and Engineering (CSE '09), pp. 739–744 (2009)
6. Baig, A., Bouridane, A., Kurugollu, F., Qu, G.: Fingerprint-iris fusion based identification system using a single hamming distance. In: Symposium on Bio-inspired Learning and Intelligent Systems for Security (BLISS '09), pp. 9–12 (2009)
7. Baker, S., Bowyer, K., Flynn, P.: Empirical evidence for correct iris match score degradation with increased time-lapse between gallery and probe matches. In: Advances in Biometrics: Lecture Notes in Computer Science #5558, pp. 1170–1179. Springer-Verlag Berlin Heidelberg (2009)
8. Baker, S., Hentz, A., Bowyer, K., Flynn, P.: Degradation of iris recognition performance due to non-cosmetic prescription contact lenses. *Comput. Vis. Image Underst.* **114**, 1030–1044 (2010)

9. Bastys, A., Kranauskas, J., Masiulis, R.: Iris matching by local extremum points of multiscale taylor expansion. In: *Advances in Biometrics: Lecture Notes in Computer Science #5558*, pp. 1070–1079. Springer-Verlag Berlin Heidelberg (2009)
10. Belcher, C., Du, Y.: A selective feature information approach for iris image-quality measure. *IEEE Trans. Inf. Forensics Secur.* **3**(3), 572–577 (2008)
11. Bharadwaj, H., Bhatt, H.S., Vatsa, M., Singh, R.: Periocular biometrics: when iris recognition fails. In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS 10)*. Washington, DC (2010)
12. Bhatnagar, J.R., Patney, R.K., Lall, B.: An information theoretic approach for formulating probability of random correspondence of biometrics. In: *World Congress on Nature and Biologically Inspired Computing (NaBIC)*, pp. 1184–1189 (2009)
13. Bhattacharjee, A., Saggi, M., Tayal, R.B.A., Kumar, A.: Decision theory based multimodal biometric authentication system using wavelet transform. *2009 International Conference on Machine Learning and Cybernetics*, vol. 4, pp. 2336–2342 (2009)
14. Bhattacharyya, D., Ranjan, R., Das, P., Kim, T., Bandyopadhyay, S.K.: Biometric authentication techniques and its future possibilities. In: *Second International Conference on Computer and Electrical Engineering (ICCEE '09)*, vol. 2, pp. 652–655 (2009)
15. Biosecure: A biometric reference system for iris OSIRIS version 2.01 (2009). <http://svnext.it-sudparis.eu/svnview2-eph/refsys/IrisOsiris/doc/howto.pdf>
16. Bodade, R.M., Talbar, S.N.: Dynamic iris localisation: a novel approach suitable for fake iris detection. In: *International Conference on Ultra Modern Telecommunications and Workshops (ICUMT '09)*, pp. 1–5 (2009)
17. Bodade, R.M., Talbar, S.N.: Shift invariant iris feature extraction using rotated complex wavelet and complex wavelet for iris recognition system. In: *Seventh International Conference on Advances in Pattern Recognition (ICAPR '09)*, pp. 449–452 (2009)
18. Boddeti, V.N., Kumar, B.V.K.V.: Extended-depth-of-field iris recognition using unrestored wavefront-coded imagery. *IEEE Trans. Syst. Man Cybern.* **40**(3), 495–508 (2010)
19. Borgen, H., Bours, P., Wolthusen, S.: Simulating the influences of aging and ocular disease on biometric recognition performance. In: *Advances in Biometrics: Lecture Notes in Computer Science #5558*, pp. 857–867. Springer-Verlag Berlin Heidelberg (2009)
20. Bowyer, K.W., Flynn, P.J.: The ND-IRIS-0405 Iris Image Database. <http://www.nd.edu/~cvrl/papers/ND-IRIS-0405.pdf>
21. Bowyer, K.W., Hollingsworth, K., Flynn, P.: Image understanding for iris biometrics: a survey. *Comput. Vis. Image Underst.* **110**(2), 281–307 (2008)
22. Bowyer, K.W., Baker, S., Hentz, A., Hollingsworth, K., Peters, T., Flynn, P.: Factors that degrade the match distribution in iris biometrics. *Identity Inf. Soc.* **2**, 327–343 (2009)
23. Breitenbach, L., Chawdhry, P.: Image quality assessment and performance evaluation for multimodal biometric recognition using face and iris. In: *Proceedings of 6th International Symposium on Image and Signal Processing and Analysis (ISPA)*, pp. 550–555 (2009)
24. Broussard, R.P., Ives, R.W.: Using artificial neural networks and feature saliency to identify iris measurements that contain the most discriminatory information for iris segmentation. In: *IEEE Workshop on Computational Intelligence in Biometrics: Theory, Algorithms, and Applications (CIB)*, pp. 46–51. Nashville, Tennessee (2009)
25. Campos, S., Salas, R., Allende, H., Castro, C.: Multimodal algorithm for iris recognition with local topological descriptors. In: *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications: Lecture Notes in Computer Science #5856*, pp. 766–773. Springer-Verlag Berlin Heidelberg (2009)
26. Carneiro, M., Veiga, A., Silva, S.P., Flores, E.L., Carrijo, G.A.: Analyzing the performance of algorithms used to localize the iris region in eye images submitted to severely compressed images. In: *IEEE International Symposium on Intelligent Signal Processing (WISP)*, pp. 281–285 (2009)
27. Chen, R., Lin, X., Ding, T., Ma, J.: Accurate and fast iris segmentation applied to portable image capture device. In: *IEEE International Workshop on Imaging Systems and Techniques (IST '09)*, pp. 80–84 (2009)

28. Chen, W., Huang, R., Hsieh, L.: Iris recognition using 3D co-occurrence matrix. In: *Advances in Biometrics: Lecture Notes in Computer Science #5558*, pp. 1122–1131. Springer-Verlag Berlin Heidelberg (2009)
29. Chen, W.S., Chuan, C.A., Shih, S.W., Chang, S.H.: Iris recognition using 2D-LDA + 2D-PCA. In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 869–872 (2009)
30. Chen, Y., Adjouadi, M., Barreto, A., Rishe, N., Andrian, J.: A computational efficient iris extraction approach in unconstrained environments. In: *IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS 09)*, pp. 1–7. Washington, DC (2009)
31. Chen, Y., Adjouadi, M., Changan, H., Barreto, A.: A new unconstrained iris image analysis and segmentation method in biometrics. In: *IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI '09)*, pp. 13–16 (2009)
32. Chou, C.T., Shih, S.W., Chen, W.S., Cheng, V.W., Chen, D.Y.: Non-orthogonal view iris recognition system. In: *IEEE Trans. Circuits Syst. Video Technol.* **20**(3), 417–430 (2010)
33. Chowhan, S.S., Shinde, G.N.: Evaluation of statistical feature encoding techniques on iris images. In: *2009 WRI World Congress on Computer Science and Information Engineering*, vol. 7, pp. 71–75 (2009)
34. Conti, V., Milici, G., Ribino, P., F. Sorbello, a.S.V.: Fuzzy fusion in multimodal biometric systems. In: *KES 2007/WIRN 2007: Lecture Notes in Artificial Intelligence #4692*, pp. 108–115 (2007)
35. Daugman, J.: Iris recognition and anti-spoofing countermeasures. In: *7th International Biometrics Conference*, London (2004)
36. Daugman, J., Downing, C.: Effect of severe image compression on iris recognition performance. *IEEE Trans. Inf. Forensics Secur.* **3**(1), 52–61 (2008)
37. Dong, W., Sun, Z., Tan, T.: A design of iris recognition system at a distance. In: *Chinese Conference on Pattern Recognition (CCPR)*, pp. 1–5 (2009)
38. Dong, W., Tan, T., Sun, Z.: Iris matching based on personalized weight map. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(9), 1744–1757 (2011)
39. Dozier, G., Frederiksen, K., Meeks, R.: Minimizing the number of bits needed for iris recognition via bit inconsistency and GRIT. In: *IEEE Workshop on Computational Intelligence in Biometrics: Theory, Algorithms, and Applications, (CIB)*, pp. 30–37. Nashville, Tennessee (2009)
40. Dutta, M., Gupta, P., Pathak, V.: Biometric based unique key generation for authentic audio watermarking. In: *Pattern Recognition and Machine Intelligence: Lecture Notes in Computer Science #5909*, pp. 458–463. Springer-Verlag Berlin Heidelberg (2009)
41. Dutta, M.K., Gupta, P., Pathak, V.K.: Biometric based watermarking in audio signals. In: *International Conference on Multimedia Information Networking and Security (MINES 2009)*, vol. 1, pp. 10–14 (2009)
42. Dutta, M.K., Gupta, P., Pathak, V.K.: Blind watermarking in audio signals using biometric features in wavelet domain. In: *IEEE Region 10 Conference (TENCON)*, pp. 1–5 (2009)
43. Du, Y., Thomas, N.L., Arslanturk, E.: Multi-level iris video image thresholding. In: *IEEE Workshop on Computational Intelligence in Biometrics: Theory, Algorithms, and Applications, (CIB)*, pp. 38–45. Nashville, Tennessee (2009)
44. Du, Y., Arslanturk, E., Zhou, Z., Belcher, C.: Video-based noncooperative iris image segmentation. *IEEE Trans. Syst. Man Cybern.* **41**(1), 64–74 (2011)
45. Elmadani, A.B.: Human authentication using fingeriris algorithm based on statistical approach. In: *Second International Conference on Networked Digital Technologies (NDT 2010)*, pp. 288–296 (2010)
46. Erbilek, M., Toygar, O.: Recognizing partially occluded irises using subpattern-based approaches. In: *24th International Symposium on Computer and Information Sciences (ISCIS)*, pp. 606–610 (2009)

47. Eskandari, M., Toygar, O.: Effect of eyelid and eyelash occlusions on iris images using subpattern-based approaches. In: Fifth International Conference on Soft Computing, Computing with Words and Perceptions in System Analysis, Decision and Control, (ICSCCW), pp. 1–4 (2009)
48. Färberböck, P., Hämmerle-Uhl, J., Kaaser, D., Pschernig, E., Uhl, A.: Transforming rectangular and polar iris images to enable cancelable biometrics. In: Image Analysis and Recognition: Lecture Notes in Computer Science #6112, pp. 276–286. Springer-Verlag Berlin Heidelberg (2010)
49. Fatt, R., Haur, T.Y., Mok, K.M.: Iris verification algorithm based on texture analysis and its implementation on dsp. In: International Conference on Signal Acquisition and Processing (ICSAP), pp. 198–202 (2009)
50. Fatt, R.Y., Tay, Y.H., Mok, K.M.: DSP-based implementation and optimization of an iris verification algorithm using textural feature. In: Sixth International Conference on Fuzzy Systems and Knowledge Discovery (FSKD), vol. 5, pp. 374–378 (2009)
51. Fierrez, J., Galbally, J., Ortega-Garcia, J.: BiosecuRID: a multimodal biometric database. Pattern Anal. Appl. **13**, 235–246 (2010)
52. Gan, J.Y., Liu, J.F.: Fusion and recognition of face and iris feature based on wavelet feature and KFDA. In: International Conference on Wavelet Analysis and Pattern Recognition (ICWAPR), pp. 47–50 (2009)
53. Garg, R., Gupta, V., Agrawal, V.: Efficient iris recognition method for identification. In: International Conference on Ultra Modern Telecommunications and Workshops, (ICUMT), pp. 1–6 (2009)
54. Garg, R., Shriram, N., Gupta, V., Agrawal, V.: A biometric security based electronic gadget control using hand gestures. In: International Conference on Ultra Modern Telecommunications and Workshops (ICUMT '09), pp. 1–8 (2009)
55. Gentile, J.E., Ratha, N., Connell, J.: An efficient, two-stage iris recognition system. In: IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS 09). Washington, DC (2009)
56. Gentile, J.E., Ratha, N., Connell, J.: SLIC: Short-length iris codes. In: IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS). Washington, DC (2009)
57. Ghouti, L., Al-Qunaieer, F.S.: Color iris recognition using quaternion phase correlation. In: Symposium on Bio-inspired Learning and Intelligent Systems for Security (BLISS '09), pp. 20–25 (2009)
58. Gonzaga, A., da Costa, R.M.: Extraction and selection of dynamic features of the human iris. In: XXII Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAPI), pp. 202–208 (2009)
59. Gorodnichy, D.O.: Evolution and evaluation of biometric systems. In: IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA), pp. 1–8 (2009)
60. Gorodnichy, D.O., Hoshino, R.: Score calibration for optimal biometric identification. In: Advances in Artificial Intelligence: Lecture Notes in Computer Science #6085, pp. 357–361. Springer-Verlag Berlin Heidelberg (2010)
61. Grabowski, K., Sankowski, W., Zubert, M., Napieralska, M.: Iris structure acquisition method. In: 16th International Conference Mixed Design of Integrated Circuits and Systems (MIXDES), pp. 640–643 (2009)
62. Grother, P., Tabassi, E., Quinn, G.W., Salamon, W.: Irex i: performance of iris recognition algorithms on standard images. In: NIST Interagency Report 7629 (2009). biometrics.nist.gov/cs_links/iris/irex/irex_summary.pdf
63. Hämmerle-Uhl, J., Pschernig, E., Uhl, A.: Cancelable iris biometrics using block re-mapping and image warping. In: Information Security: Lecture Notes in Computer Science #5735, pp. 135–142. Springer-Verlag Berlin Heidelberg (2009)
64. Hämmerle-Uhl, J., Prähauser, C., Starzacher, T., Uhl, A.: Improving compressed iris recognition accuracy using JPEG2000 RoI coding. In: Advances in Biometrics: Lecture Notes in Computer Science #5558, pp. 1102–1111. Springer-Verlag Berlin Heidelberg (2009)

65. Hao, F., Anderson, R., Daugman, J.: Combining crypto with biometrics effectively. *IEEE Trans. Comput.* **55**(9), 1081–1088 (2006)
66. Hao, F., Daugman, J., Zielinski, P.: A fast search algorithm for a large fuzzy database. *IEEE Trans. Inf. Forensics Secur.* **3**(2), 203–212 (2008)
67. Haralick, R., Shanmugam, K., Dinstein, L.: Textural features for image classification. *IEEE Trans. Syst. Man Cybern.* **3**(6), 610–621 (1973)
68. Hassanien, A., Abraham, A., Grosan, C.: Spiking neural network and wavelets for hiding iris data in digital images. *Soft Comput.* **13**(4), 401–416 (2009)
69. He, X., An, X., Shi, P.: Statistical texture analysis based approach for fake iris detection using support vector machine. In: *Advances in Biometrics: Lecture Notes in Computer Science #4642*, pp. 540–546. Springer-Verlag Berlin Heidelberg (2007)
70. He, X., Yan, J., Chen, G., Shi, P.: Contactless autofeedback iris capture design. *IEEE Trans. Instrum. Meas.* **57**(7), 1369–1375 (2008)
71. He, Y., H. Yang, Y.H., He, H.: An elimination method of light spot based on iris image fusion. *Commun. Comput. Inf. Sci.* **15**(12), 415–422 (2008)
72. He, Z., Sun, Z., Tan, T., Qiu, X.: Enhanced usability of iris recognition via efficient user interface and iris image restoration. In: *15th IEEE International Conference on Image Processing (ICIP)*, pp. 261–264 (2008)
73. He, X., Lu, Y., Shi, P.: A new fake iris detection method. In: *Advances in Biometrics: Lecture Notes in Computer Science #5558*, pp. 1132–1139. Springer-Verlag Berlin Heidelberg (2009)
74. He, Z., Sun, Z., Tan, T., Wei, Z.: Efficient iris spoof detection via boosted local binary patterns. In: *Advances in Biometrics: Lecture Notes in Computer Science #5558*, pp. 1080–1090. Springer-Verlag Berlin Heidelberg (2009)
75. He, Z., Tan, T., Sun, Z., Qiu, X.: Toward accurate and fast iris segmentation for iris biometrics. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**(9), 1670–1684 (2009)
76. Hollingsworth, K., Bowyer, K., Flynn, P.: Pupil dilation degrades iris biometric performance. *Comput. Vis. Image Underst.* **113**, 150–157 (2009)
77. Hollingsworth, K.P., Bowyer, K.W., Flynn, P.J.: Pupil dilation degrades iris biometric performance. *Comput. Vis. Image Underst.* **113**(1), 150–157 (2009)
78. Hollingsworth, K.P., Bowyer, K.W., Flynn, P.J.: The best bits in an iris code. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**(6), 964–973 (2009)
79. Hollingsworth, K.P., Bowyer, K.W., Flynn, P.J.: Using fragile bit coincidence to improve iris recognition. In: *IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS)*, pp. 1–6. Washington, DC (2009)
80. Hollingsworth, K., Peters, T., Bowyer, K.W., Flynn, P.J.: Iris recognition using signal-level fusion of frames from video. *IEEE Trans. Inf. Forensics Secur.* **4**(4), 837–848 (2009)
81. Hollingsworth, K.P., Bowyer, K.W., Flynn, P.J.: Identifying useful features for recognition in near-infrared periocular images. In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS)*. Washington, DC (2010)
82. Huang, X., Ren, L., Yang, R.: Image deblurring for less intrusive iris capture. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1558–1565 (2009)
83. Ives, R.W., Bishop, D.A.D., Du, Y., Belcher, C.: Effects of image compression on iris recognition performance and image quality. In: *IEEE Workshop on Computational Intelligence in Biometrics: Theory, Algorithms, and Applications (CIB)*, pp. 16–21. Nashville, Tennessee (2009)
84. Jang, Y., Kang, B.J., Park, K.R.: A novel portable iris recognition system and usability evaluation. *Int. J. Control Autom. Syst.* **8**(1), 91–98 (2010)
85. Johnson, P.A., Lopez-Meyer, P., Sazonova, N., Hua, F., Schuckers, S.: Quality in face and iris research ensemble (Q-FIRE). In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS)*. Washington, DC (2010)
86. Kalka, N., Bartlow, N., Cukic, B.: An automated method for predicting iris segmentation failures. In: *IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS)*. Washington, DC, pp. 1–8 (2009)

87. Kalka, N., Zuo, J., Schmid, N.A., Cukic, B.: Estimating and fusing quality factors for iris biometrics images. *IEEE Trans. Syst. Man Cybern.* **40**(3), 509–524 (2010)
88. Kanade, S., Petrovska-Delacretaz, D., Dorizzi, B.: Cancelable iris biometrics and using error correcting codes to reduce variability in biometric data. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 120–127 (2009)
89. Kanade, S., Petrovska-Delacretaz, D., Dorizzi, B.: Multi-biometrics based cryptographic key regeneration scheme. In: *IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS)*. Washington, DC (2009)
90. Kang, B.J., Park, K.R.: A new multi-unit iris authentication based on quality assessment and score level fusion for mobile phones. *Mach. Vis. Appl.* **21**, 541–553 (2010)
91. Kannavara, R., Bourbakis, N.: Iris biometric authentication based on local global graphs: an FPGA implementation. In: *IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA)*, pp. 1–7 (2009)
92. Kheirolahy, R., Ebrahimnezhad, H., Sedaaghi, M.H.: Robust pupil boundary detection by optimized color mapping for iris recognition. In: *14th International CSI Computer Conference (CSICC)*, pp. 170–175 (2009)
93. Kong, A., Zhang, D., Kamel, M.: An analysis of IrisCode. *IEEE Trans. Image Process.* **19**, 522–532 (2010)
94. Konrad, M., Stogner, H., Uhl, A.: Custom design of JPEG quantisation tables for compressing iris polar images to improve recognition accuracy. In: *Advances in Biometrics: Lecture Notes in Computer Science #5558*, pp. 1091–1101. Springer-Verlag Berlin Heidelberg (2009)
95. Konrad, M., Stogner, H., Uhl, A.: Evolutionary optimization of JPEG quantization tables for compressing iris polar images in iris recognition systems. In: *Proceedings of 6th International Symposium on Image and Signal Processing and Analysis (ISPA)*, pp. 534–539 (2009)
96. Kostmayer, G., Stogner, H., Uhl, A.: Custom JPEG quantization for improved iris recognition accuracy. In: *Emerging Challenges for Security, Privacy and Trust: IFIP Advances in Information and Communication Technology #297*, pp. 76–86 (2009)
97. Krichen, E., Garcia-Salicetti, S., Dorizzi, B.: A new phase-correlation-based iris matching for degraded images. *IEEE Trans. Syst. Man Cybern.* **39**(4), 924–934 (2009)
98. Krichen, E., Dorizzi, B., Sun, Z., Garcia-Salicetti, S.: Iris recognition. In: *Guide to Biometric Reference Systems and Performance Evaluation*, p. 25. Springer (2009)
99. Kyaw, K.S.S.: Iris recognition system using statistical features for biometric identification. In: *International Conference on Electronic Computer Technology*, pp. 554–556 (2009)
100. Labati, R.D., Piuri, V., Scotti, F.: Agent-based image iris segmentation and multiple views boundary refining. In: *IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS)*, pp. 1–7. Washington, DC (2009)
101. Labati, R.D., Piuri, V., Scotti, F.: Neural-based iterative approach for iris detection in iris recognition systems. In: *IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA)*, pp. 1–6 (2009)
102. AT&T Cambridge Laboratories: The database of faces. <http://www.cl.cam.ac.uk/research/dtg/attarchive/facedatabase.html>
103. Lee, Y., Phillips, P., Micheals, R.: An automated video-based system for iris recognition. In: *Advances in Biometrics: Lecture Notes in Computer Science #5558*, pp. 1160–1169. Springer-Verlag Berlin Heidelberg (2009)
104. Leonard, D.C., Pons, A.P., Asfour, S.S.: Realization of a universal patient identifier for electronic medical records through biometric technology. *IEEE Trans. Inf. Technol. Biomed.* **13**(4), 494–500 (2009)
105. Li, Y., Du, S.: Biometric watermarking based on affine parameters estimation. In: *2nd International Congress on Image and Signal Processing (CISP)*, pp. 1–6 (2009)
106. Li, Y., Du, S.: Biometric watermarking based on affine parameters estimation. In: *International Conference on Multimedia Computing and Systems (ICMCS)*, pp. 123–128 (2009)
107. Li, Y., Savvides, M.: Automatic iris mask refinement for high performance iris recognition. In: *IEEE Workshop on Computational Intelligence in Biometrics: Theory, Algorithms, and Applications (CIB)*, pp. 52–58. Nashville, Tennessee (2009)

108. Li, Y., Savvides, M.: A pixel-wise, learning-based approach for occlusion estimation of iris images in polar domain. In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1357–1360 (2009)
109. Lin, J., Li, J.P., Lin, H., Ming, J.: Robust person identification with face and iris by modified PUM method. In: *International Conference on Apperceiving Computing and Intelligence Analysis (ICACIA)*, pp. 321–324 (2009)
110. Liu, L., Gu, X.F., Li, J.P., Lin, J., Shi, J.X., Huang, Y.Y.: Research on data fusion of multiple biometric features. In: *International Conference on Apperceiving Computing and Intelligence Analysis (ICACIA)*, pp. 112–115 (2009)
111. Liu, X., Li, P., Song, Q.: Eyelid localization in iris images captured in less constrained environment. In: *Advances in Biometrics: Lecture Notes in Computer Science #5558*, pp. 1140–1149 (2009)
112. Liu-Jimenez, J., Sanchez-Reillo, R., Fernandez-Saavedra, B.: Iris biometrics for embedded systems. *IEEE Trans. Very Large Scale Integr. Syst.* **19**(2), 274–282 (2011)
113. Luo, Y., Cheung, S., Ye, S.: Anonymous biometric access control based on homomorphic encryption. In: *IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1046–1049 (2009)
114. Lyle, J.R., Miller, P.E., Pundlik, S.J., Woodard, D.L.: Soft biometric classification using periocular region features. In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS)*. Washington, DC (2010)
115. Mahmud, M., Khan, M.K., Alghathbar, K.: Biometric-gaussian-stream (BGS) cipher with new aspect of image encryption (data hiding). In: *Bio-Science and Bio-Technology*, pp. 97–107. Springer, Berlin/Heidelberg (2009)
116. Maltoni, D.: Biometric fusion. In: *Handbook of Fingerprint Recognition*. Springer-Verlag London (2009)
117. Ma, L., Tan, T., Wang, Y., Zhang, D.: Local intensity variation analysis for iris recognition. *Pattern Recognit.* **37**(6), 1287–1298 (2004)
118. Masek, L.: Iris recognition (2003). http://www.csse.uwa.edu.au/_pk/studentprojects/libor/
119. Matey, J.R., Kennell, L.R.: Iris recognition – beyond one meter. In: *Handbook of Remote Biometrics* (2009)
120. Matey, J.R., Naroditsky, O.: Iris on the move: acquisition of images for iris recognition in less constrained environments. *Proc. IEEE* **94**(11), 1936–1947 (2006)
121. McCloskey, S., Au, A.W., Jelinek, J.: Iris capture from moving subjects using a fluttering shutter. In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS)*. Washington, DC (2010)
122. Mehrotra, H., Majhi, B., Gupta, P.: Annular iris recognition using SURF. In: *Pattern Recognition and Machine Intelligence: Lecture Notes in Computer Science #5909*, pp. 464–469. Springer-Verlag Berlin Heidelberg (2009)
123. Mehrotra, H., Badrinath, G.S., Majhi, B., Gupta, P.: An efficient dual stage approach for iris feature extraction using interest point pairing. In: *IEEE Workshop on Computational Intelligence in Biometrics: Theory, Algorithms, and Applications (CIB)*, pp. 59–62 (2009)
124. Mehrotra, H., Srinivas, B.G., Majhi, B., Gupta, P.: Indexing iris biometric database using energy histogram of DCT subbands. *Contemp. Comput.* **40**(4), 194–204 (2009)
125. Merkow, J., Jou, B., Savvides, M.: An exploration of gender identification using only the periocular region. In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS)*. Washington, DC (2010)
126. Miller, P.E., Lyle, J.R., Pundlik, S.J., Woodard, D.L.: Performance evaluation of local appearance based periocular recognition. In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS)*. Washington, DC (2010)
127. Mishra, R., Pathak, V.: Human recognition using fusion of iris and ear data. In: *International Conference on Methods and Models in Computer Science (ICM2CS)*, pp. 1–5 (2009)
128. Miyazawa, K., Ito, K., Aoki, T., Kobayashi, K.: An effective approach for iris recognition using phase-based image matching. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(10), 1741–1756 (2008)

129. Mohammadi, S., Jahanshahi, H.: A secure e-tendering system. In: IEEE International Conference on Electro/Information Technology (EIT), pp. 62–67 (2009)
130. Moi, S.H., Rahim, N., Abdul, B. et al.: iris biometric cryptography for identity document. International Conference of Soft Computing and Pattern Recognition (SOCPAR), pp. 736–741 (2009)
131. Mondal, A., Roy, K., Bhattacharya, P.: Secure and simplified access to home appliances using iris recognition. In: IEEE Workshop on Computational Intelligence in Biometrics: Theory, Algorithms, and Applications (CIB), pp. 22–29. Nashville, Tennessee (2009)
132. Moravec, P., Gajdos, P., Snasel, V., Saeed, K.: Normalization impact on SVD-based iris recognition. In: International Conference on Biometrics and Kansei Engineering (ICBAKE), pp. 60–64 (2009)
133. Morizet, N., Gilles, J.: A new adaptive combination approach to score level fusion for face and iris biometrics combining wavelets and statistical moments. In: Advances in Visual Computing: Lecture Notes in Computer Science #5359, pp. 661–671. Springer-Verlag Berlin Heidelberg (2008)
134. Munemoto, T., Li, Y.H., Savvides, M.: Hallucinating irises – dealing with partial and occluded iris regions. In: IEEE 2nd International Conference on Biometrics: Theory, Applications, and Systems (BTAS). Washington, DC (2008)
135. Newton, E.M., Phillips, P.J.: Meta-analysis of third-party evaluations of iris recognition. IEEE Trans. Syst. Man Cybern. **39**(1), 4–11 (2009)
136. Ortega-Garcia, J., Fierrez, J., Alonso-Fernandez, F., Galbally, J.: The multi-scenario multi-environment biosecure multimodal database (BMDB). IEEE Trans. Pattern Anal. Mach. Intell. **32**(6), 1097–1111 (2010)
137. Pan, L., Xie, M., Zheng, T., Ren, J.: A robust iris localization model based on phase congruency and least trimmed squares estimation. In: Image Analysis and Processing -ICIAP Lecture Notes in Computer Science #5716, pp. 682–691 (2009)
138. Patil, C.M., Patilkulkarani, S.: A comparative study of feature extraction approaches for an efficient iris recognition system. In: Information Processing and Management, pp. 411–416 (2010)
139. Patil, C.M., Patilkulkarani, S.: An approach to enhance security environment based on SIFT feature extraction and matching to iris recognition. In: Information Processing and Management, pp. 527–530 (2010)
140. Patil, C.M., Patilkulkarani, S.: Iris feature extraction for personal identification using lifting wavelet transform. In: International Conference on Advances in Computing, Control, and Telecommunication Technologies (ACT), pp. 764–766 (2009)
141. Petrovska-Delacretaz, D., Mayoue, A., Dorizzi, B.: The biosecure benchmarking methodology for biometric performance evaluation. In: Guide to Biometric Reference Systems and Performance Evaluation (2009)
142. Phillips, P.J., Beveridge, J.R.: An introduction to biometric-completeness: The equivalence of matching and quality. In: IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS), pp. 1–5. Washington, DC (2009)
143. Phillips, P.J., Newton, E.M.: Biometric systems: the rubber meets the road. Proc. IEEE **97**(5), 782–783 (2009)
144. Phillips, P.J., Bowyer, K.W., Flynn, P.J.: Comment on the CASIA version 1.0 iris dataset. IEEE Trans. Pattern Anal. Mach. Intell. **29**(10), 1869–1870 (2007)
145. Phillips, P.J., Bowyer, K.W., Flynn, P.J., Liu, X., Scruggs, W.T.: The iris challenge evaluation 2005. In: 2nd IEEE International Conference on Biometrics: Theory, Applications and Systems (BTAS). Washington, DC (2008)
146. Phillips, P., Flynn, P., Beveridge, J., Scruggs, W., OToole, A., Bolme, D., Bowyer, K., Draper, B., Givens, G., Lui, Y., Sahibzada, H., Scallan, J., Weimer, S.: Overview of the multiple biometrics grand challenge. In: Advances in Biometrics: Lecture Notes in Computer Science #5558. Springer-Verlag Berlin Heidelberg (2009)

147. Phillips, P.J., Scruggs, W., O'Toole, A., Flynn, P., Bowyer, K.W., Schott, C., Sharpe, M.: FRVT 2006 and ICE 2006 large-scale experimental results. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(5), 831–846 (2010)
148. Pillai, J.K., Patel, V.M., Chellappa, R.: Sparsity inspired selection and recognition of iris images. In: *IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS)*, pp. 1–6. Washington, DC (2009)
149. Plaga, R.: Biometric keys: suitable use cases and achievable information content. *Int. J. Inf. Secur.* **8**, 447–454 (2009)
150. Poh, N., Bourlai, T., Kittler, J., Allano, L., Alonso-Fernandez, F.: Benchmarking quality-dependent and cost-sensitive score-level multimodal biometric fusion algorithms. *IEEE Trans. Inf. Forensics Secur.* **4**(4), 849–866 (2009)
151. Proenca, H.: Iris recognition: a method to segment visible wavelength iris images acquired on-the-move and at-a-distance. In: *Advances in Visual Computing: Lecture Notes in Computer Science #5358*, vol. 32(8), pp. 731–742. Springer-Verlag Berlin Heidelberg (2008)
152. Proenca, H.: On the feasibility of the visible wavelength, at-a-distance and on-the-move iris recognition. In: *IEEE Workshop on Computational Intelligence in Biometrics: Theory, Algorithms, and Applications (CIB)*, pp. 9–15. Nashville, Tennessee (2009)
153. Proenca, H.: Iris recognition: On the segmentation of degraded images acquired in the visible wavelength. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(8), 1502–1516 (2010)
154. Proenca, H.: Quality assessment of degraded iris images acquired in the visible wavelength. *IEEE Trans. Inf. Forensics Secur.* **6**(1), 82–95 (2011)
155. Proenca, H., Filipe, S., Santos, R., Oliveira, J., Alexandre, L.: The UBIRIS.v2: a database of visible wavelength images captured on-the-move and at-a-distance. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(8), 1529–1535 (2010)
156. Pundlik, S.J., Woodard, D.L., Birchfield, S.T.: Non-ideal iris segmentation using graph cuts. In: *IEEE CVPR Workshop on Biometrics* (2008)
157. Rachubinski, M.: Iris identification using geometrical wavelets. In: *Computer Vision and Graphics* (2009)
158. Radhika, K., Sheela, S., Venkatesha, M., Sekhar, G.: Multi-modal authentication using continuous dynamic programming. In: *Biometric ID Management and Multimodal Communication: Lecture Notes in Computer Science #5707*, pp. 228–235. Springer-Verlag Berlin Heidelberg (2009)
159. Rakvic, R.N., Ulis, B.J., Broussard, R.P., Ives, R.W., Steiner, N.: Parallelizing iris recognition. *IEEE Trans. Inf. Forensics Secur.* **4**(4), 812–823 (2009)
160. Rankin, D., Scotney, B., Morrow, P., McDowell, R., Pierscionek, B.: Comparing and improving algorithms for iris recognition. In: *13th International Machine Vision and Image Processing Conference (IMVIP)*, pp. 99–104 (2009)
161. Ratha, N.K.: Privacy protection in high security biometrics applications. In: *Ethics and Policy of Biometrics: Lecture Notes in Computer Science #6005*, pp. 62–69. Springer-Verlag Berlin Heidelberg (2010)
162. Rathgeb, C., Uhl, A.: Systematic construction of iris-based fuzzy commitment schemes. In: *Advances in Biometrics: Lecture Notes in Computer Science #5558*, pp. 940–949. Springer-Verlag Berlin Heidelberg (2009)
163. Rathgeb, C., Uhl, A.: Privacy preserving key generation for iris biometrics. In: *Communications and Multimedia Security: Lecture Notes in Computer Science #6109*, pp. 191–200. Springer-Verlag Berlin Heidelberg (2010)
164. Rathgeb, C., Uhl, A.: Two-factor authentication or how to potentially counterfeit experimental results in biometric systems. In: *ICIAR 2010: Lecture Notes in Computer Science #6112*, pp. 296–305. Springer-Verlag Berlin Heidelberg (2010)
165. Rathgeb, C., Uhl, A., Wild, P.: Incremental iris recognition: a single-algorithm serial fusion strategy to optimize time complexity. In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS)*. Washington, DC (2010)

166. Rattani, A., Tistarelli, M.: Robust multi-modal and multi-unit feature level fusion of face and iris biometrics. In: *Advances in Biometrics: Lecture Notes in Computer Science #5558*, pp. 960–969. Springer-Verlag Berlin Heidelberg (2009)
167. Ren, J., Xie, M.: Evaluation of iris images definition based on pupil's edge kurtosis. In: *2nd International Congress on Image and Signal Processing (CISP)*, pp. 1–4 (2009)
168. Ren, J., Xie, M.: Research on clarity-evaluation-method for iris images. In: *Second International Conference on Intelligent Computation Technology and Automation (ICICTA)*, vol. 1, pp. 682–685 (2009)
169. Ricanek, K.: Dissecting the human identity. *Computer* **44**, 96–97 (2011)
170. Ross, A., Pasula, R., Hornak, L.: Exploring multispectral iris recognition beyond 900nm. In: *IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS)*. Washington, DC (2009)
171. Ross, A., Rattani, A., Tistarelli, M.: Exploiting the doddington zoo effect in biometric fusion. In: *IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems (BTAS)* (2009)
172. Roy, K., Bhattacharya, P.: Improving features subset selection using genetic algorithms for iris recognition. In: *Artificial Neural Networks in Pattern Recognition: Lecture Notes in Computer Science #5064*, pp. 292–304. Springer-Verlag Berlin Heidelberg (2008)
173. Roy, K., Bhattacharya, P.: Optimal features subset selection using genetic algorithms for iris recognition. In: *Image Analysis and Recognition: Lecture Notes in Computer Science #5112*, pp. 894–904. Springer-Verlag Berlin Heidelberg (2008)
174. Roy, K., Bhattacharya, P.: Iris recognition in nonideal situations. In: *Information Security: Lecture Notes in Computer Science #5735*, pp. 143–150 (2009)
175. Roy, K., Bhattacharya, P.: Level set approaches and adaptive asymmetrical SVMs applied for nonideal iris recognition. In: *Image Analysis and Recognition: Lecture Notes in Computer Science #5627*, pp. 418–428 (2009)
176. Roy, K., Bhattacharya, P.: Nonideal iris recognition using level set approach and coalitional game theory. In: *Computer Vision Systems: Lecture Notes in Computer Science #5815*, pp. 394–402 (2009)
177. Roy, K., Bhattacharya, P.: Unideal iris segmentation using region-based active contour model. In: *ICIAR Lecture Notes in Computer Science #6112*, pp. 256–265. Springer-Verlag Berlin Heidelberg (2010)
178. Ruiz-Albacete, V., Tome-Gonzalez, P., Alonso-Fernandez, F.: Direct attacks using fake images in iris verification. In: *Biometrics and Identity Management: Lecture Notes in Computer Science #5372*, pp. 181–190. Springer-Verlag Berlin Heidelberg (2008)
179. Ryan, W.J., Woodard, D.L., Duchowski, A.T., Birchfield, S.T.: Adapting starburst for elliptical iris segmentation. In: *IEEE 2nd International Conference on Biometrics: Theory, Applications, and Systems (BTAS)*. Washington, DC (2008)
180. Schmid, N.A., Nicolò, F.: On empirical recognition capacity of biometric systems under global PCA and ICA encoding. *IEEE Trans. Inf. Forensics Secur.* **3**(3), 512–528 (2008)
181. Schmid, N.A., Nicolò, F.: A method for selecting and ranking quality metrics for optimization of biometric recognition systems. In: *Computer Vision and Pattern Recognition Workshops (CVPR Workshops)*, pp. 126–133 (2009)
182. Scotti, F., Piuri, V.: Adaptive reflection detection and location in iris biometric images by using computational intelligence techniques. *IEEE Trans. Instrum. Meas.* **59**(7), 1825–1833 (2010)
183. Shah, S., Ross, A.: Iris segmentation using geodesic active contours. *IEEE Trans. Inf. Forensics Secur.* **4**(4), 824–836 (2009)
184. Sheela, S., Radhika, K., Venkatesha, M., Vijaya, P.: Iris and signature authentication using continuous dynamic programming. In: *2nd International Congress on Image and Signal Processing (CISP)*, pp. 1–5 (2009)
185. Stark, L., Bowyer, K.W., Siena, S.: Human perceptual categorization of iris texture patterns. In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS)*. Washington, DC (2010)

186. Sudha, N., Puhan, N.B., Xia, H., Jiang, X.: Iris recognition on edge maps. *IET Comput. Vis.* **3**(1), 1–7 (2009)
187. Sun, Z., Tan, T.: Ordinal measures for iris recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**(12), 2211–2226 (2009)
188. Tajbakhsh, N., Araabi, B., Soltanian-zadeh, H.: Noisy iris verification: a modified version of local intensity variation method. In: *Advances in Biometrics: Lecture Notes in Computer Science #5558*, pp. 1150–1159 (2009)
189. Tajbakhsh, N., Misaghian, K., Bandari, N.: A region-based iris feature extraction method based on 2D-wavelet transform. In: *Biometric ID Management and Multimodal Communication: Lecture Notes in Computer Science #5707*, pp. 301–307 (2009)
190. Takano, H., Nakamura, K.: Rotation independent iris recognition by the rotation spreading neural network. In: *IEEE 13th International Symposium on Consumer Electronics (ISCE)*, pp. 651–654 (2009)
191. Tan, F., Ong, T.S., Tee, C., Teoh, A.B.J.: Image hashing enabled technique for biometric template protection. In: *IEEE Region 10 Conference (TENCON)*, pp. 1–5 (2009)
192. Tayal, A., Balasubramaniam, R., Kumar, A., Bahattacharjee, A., Saggi, M.: A multimodal biometric authentication system using decision theory, iris and speech recognition. In: *2nd International Workshop on Nonlinear Dynamics and Synchronization (INDS)*, pp. 1–8 (2009)
193. Tayal, A., Balasubramaniam, R., Kumar, A., Bhattacharjee, A., Saggi, M.: A multimodal biometric system coupling iris recognition and speaker identification systems through decision theory. In: *3rd International Conference on Anti-counterfeiting, Security, and Identification in Communication (ASID)*, pp. 135–137 (2009)
194. Thompson, J.W., Flynn, P.J.: A segmentation perturbation method for improved iris recognition. In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS)*. Washington, DC (2010)
195. Vandal, N.A., Savvides, M.: CUDA accelerated iris template matching on graphics processing units. In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS)*. Washington, DC (2010)
196. Vatsa, M., Singh, R., Noore, A.: Improving iris recognition performance using segmentation, quality enhancement, match score fusion, and indexing. *IEEE Trans. Syst. Man Cybern.* **38**(4), 1021–1035 (2008)
197. Vatsa, M., Singh, R., Noore, A., Singh, S.K.: Belief function theory based biometric match score fusion: case studies in multi-instance and multi-unit iris verification. In: *Seventh International Conference on Advances in Pattern Recognition (ICAPR)*, pp. 433–436 (2009)
198. Velisavljevic, V.: Low-complexity iris coding and recognition based on directionlets. *IEEE Trans. Inf. Forensics Secur.* **4**(3), 410–417 (2009)
199. Wang, F., Han, J.: Multimodal biometric authentication based on score level fusion using support vector machine. In: *Opto-Electronics Review*, vol. 17, pp. 59–64 (2009)
200. Wang, D., Li, J., Memik, G.: Authentication scheme of DRM system for remote users based on multimodal biometrics, watermarking and smart card. In: *WRI Global Congress on Intelligent Systems (GCIS)*, vol. 2, pp. 530–534 (2009)
201. Wang, J., Li, Y., Ao, X., Wang, C., Zhou, J.: Multi-modal biometric authentication fusing iris and palmprint based on GMM. In: *IEEE 15th Workshop on Statistical Signal Processing (SSP)*, pp. 349–352 (2009)
202. Wang, X., Zhao, L., Kong, Q.: Iris recognition system design and development of large animals for tracing source of infection. *Int. Joint Conf. Comput. Sci. Optim.* **1**, 610–613 (2009)
203. Wang, Z.F., Han, Q., Li, Q., Niu, X.M., Busch, C.: Complex common vector for multimodal biometric recognition. *Electronics Lett.* **45**(10), 495–496 (2009)
204. Wang, Z., Han, Q., Niu, X., Busch, C.: Feature-level fusion of iris and face for personal identification. In: *Advances in Neural Networks – ISNN 2009: Lecture Notes in Computer Science #5553*, pp. 356–364. Springer-Verlag Berlin Heidelberg (2009)

205. Wang, Z., Li, Q., Niu, X., Busch, C.: Multimodal biometric recognition based on complex KFDA. In: Fifth International Conference on Information Assurance and Security (IAS), vol. 2, pp. 177–180 (2009)
206. Wei, Z., Qiu, X., Sun, Z., Tan, T.: Counterfeit iris detection based on texture analysis. In: 19th International Conference on Pattern Recognition (2008)
207. Wheeler, F.W., Perera, A.G.A., Abramovich, G., Yu, B., Tu, P.H.: Stand-off iris recognition system. In: IEEE 2nd International Conference on Biometrics: Theory, Applications, and Systems (BTAS). Washington, DC (2008)
208. Wibowo, E.P., Maulana, W.S.: Real-time iris recognition system using a proposed method. In: International Conference on Signal Processing Systems, pp. 98–102 (2009)
209. Wu, D.M., Wang, J.N.: An improved iris recognition method based on gray surface matching. In: Fifth International Conference on Information Assurance and Security (IAS), vol. 1, 247–249 (2009)
210. Xiangde, Z., Qi, W., Hegui, Z., Cuili, Y., Longcheng, G., Xianyan, L.: Noise detection of iris image based on texture analysis. In: Chinese Control and Decision Conference (CCDC), pp. 2366–2370 (2009)
211. Xu, X., Guo, P.: Iris feature extraction based on the complete 2DPCA. In: Advances in Neural Networks: Lecture Notes in Computer Science #5552, pp. 950–958 (2009)
212. Xu, J., Cha, M., Heyman, J.L., Venugopalan, S., Abiantun, R., Savvides, M.: Robust local binary pattern feature sets for periocular biometric identification. In: Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS). Washington, DC (2010)
213. Yager, N., Dunstone, T.: The biometric menagerie. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(2), 220–230 (2010)
214. Zhang, L., Sun, Z., Tan, T., Hu, S.: Robust biometric key extraction based on iris cryptosystem. In: Advances in Biometrics: Lecture Notes in Computer Science #5558. Springer-Verlag Berlin Heidelberg (2009)
215. Zhao, X., Xie, M.: A practical design of iris recognition system based on DSP. In: International Conference on Intelligent Human-Machine Systems and Cybernetics (IHMSC), vol. 1, pp. 66–70 (2009)
216. Zhiping, Z., Maomao, H., Ziwen, S.: An iris recognition method based on 2DWPCA and neural network. In: Chinese Control and Decision Conference (CCDC), pp. 2357–2360 (2009)
217. Zhou, Z., Du, Y., Belcher, C.: Transforming traditional iris recognition systems to work in nonideal situations. *IEEE Trans. Ind. Electron.* **56**(8), 3203–3213 (2009)
218. Zhou, Z., Du, Y., Belcher, C.: Transforming traditional iris recognition systems to work on non-ideal situations. In: IEEE Workshop on Computational Intelligence in Biometrics: Theory, Algorithms, and Applications (CIB), pp. 1–8. Nashville, Tennessee (2009)
219. Zuo, J., Schmid, N.A.: Global and local quality measures for NIR iris video. In: Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 120–125. Nashville, Tennessee (2009)
220. Zuo, J., Schmid, N.A.: On a methodology for robust segmentation of nonideal iris images. *IEEE Trans. Syst. Man Cybern.* **40**(3), 703–718 (2010)
221. Zuo, J., Nicolo, F., Schmid, N.A.: Cross spectral iris matching based on predictive image mapping. In: Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS). Washington, DC (2010)

Chapter 3

Standard Iris Storage Formats

George Quinn, Patrick Grother, and Elham Tabassi

Abstract Iris recognition standards are open specifications for iris cameras, iris image properties, and iris image records. Biometric data standards are a necessity for applications in which a consumer of a data record must process biometric input from an arbitrary producer. The archetype for standard iris storage formats has been the flight of standards already developed for the storage of biometric data on electronic passports.

3.1 Introduction

Iris recognition standards are open specifications for iris cameras, iris image properties, and iris image records. Biometric data standards are a necessity for applications in which a consumer of a data record must process biometric input from an arbitrary producer. The archetype for standard iris storage formats has been the flight of standards already developed for the storage of biometric data on electronic passports. The most widely implemented of these is the ISO/IEC 19794-5 [10], which specifies the format requirements for the storage of facial images in the data group (DG2) container of the smart card chip.¹ In this case, the producer of the record, a national passport issuer, is tasked with collecting and preparing a facial image conformant to the standard. An immigration checkpoint is later tasked with reading and using that image. The extent to which such exchanges

¹For iris recognition, the electronic passport's logical data structure specification includes an analogous container, DG4, that will hold an ISO/IEC 19794-6 iris image record. To date, this has not been widely utilized.

G. Quinn (✉) • P. Grother • E. Tabassi
National Institute of Standards and Technology, Gaithersburg, MD, USA
e-mail: gw@nist.gov

are successful is quantified in terms of performance metrics for identification, verification, and enrollment. To this end, formal, open, and widely implemented documentary standards are critical elements.

3.2 Role and Importance of Iris Standards

Image data interchange standards define parseable records that allow syntactic interoperability. This is a foundational element for a marketplace of off-the-shelf products and is a necessary condition to achieve supplier independence and to prevent vendor lock-in. It is perhaps surprising that while many raster image formats are open and standardized [3–5], it remains common to store images in fully proprietary (i.e., unpublished) formats. Such practices may be acceptable *within* an application but are fatal once cross-organizational interchange of data is required. Data format standards allow modular integration of products without compromising architectural scope. Additionally, they facilitate the upgrade process, thereby preventing obsolescence.

There are several business implications of standards development. A well-structured and widely implemented standard can create entirely new markets for companies (e.g., e-passports can include face, fingerprint, and iris records). On the other hand, robust standards also promote competition, which leads to reduced profit margins. This process, commoditization, is an inhibitory factor for many technology companies that balance the promise of new or expanded marketplaces against reduced barriers to entry for competitors. The decision to enter the marketplace is influenced by the amount of intellectual property that a standard allows suppliers to hide behind their implementations. From the user's perspective, standards serve to promote competition, which leads to enhanced performance. For example, the National Institute of Standards and Technology (NIST) conducted the Iris Exchange test (IREX I [11]), to offer quantitative support to iris recognition standards. Ten providers implemented the software modules necessary to create a standard record and to recognize it. This commercial support constituted an order of magnitude expansion of the iris recognition industry compared to 5 years earlier.

Standards do not in and of themselves assure interoperability. Specifically, when a standard is not fully prescriptive or does not allow for optional content, multiple implementations of the standard may still lack interoperability with one another. This problem can be averted by applying further constraints on the application of the standard. "Application profile" standards identify the needed base standards and refine their optional content and interpretations. This was done recently by the US government for its PIV employee credential [21].

Iris interoperability requires that data records are both syntactically and semantically understood by the receiving system. The syntactic requirements of a standard allow parsing software to find the correct information. The semantic requirements are present to regulate the appearance of the iris image (i.e., the array of pixel values) so that an iris recognition algorithm is able to *accurately* compare or identify

the input image. For example, a blurred image will often not be recognizable. The notion that syntactic correctness is not in itself sufficient for interoperability is an important aspect of biometrics that is not always present in other fields.

3.3 History of Iris Standardization

The current iris standards descend from standardization efforts beginning around 2002. At that time, the only commercial entity in the field, Iridian, maintained a certification program for iris camera manufacturers to establish minimum standards for image quality and functionality. While some cameras produced an image, the closed-format IrisCode template was effectively the data interchange format because the only recognition provider was Iridian. After the September 11 attack, there was considerable impetus behind secure credentialing in a number of applications. Foremost of these was border management, and the International Civil Aviation Organization's (ICAO's) standardization of the electronic passport was paralleled by the development of biometric data interchange standards in the newly formed ISO/IEC JTC 1 subcommittee 37 for biometrics [1]. The iris recognition industry at the time, essentially the Iridian company, volunteered the first draft of a data interchange record format, which was progressed in parallel efforts in the United States and internationally. The results, respectively, were INCITS 379 and the almost identical international standard ISO/IEC 19794-6 [2] (see Table 3.1). Both standards included essentially two iris image formats. The first was a generic rectilinear format (see Fig. 3.1), and the second was a polar-sampled version of the iris texture lying between nonconcentric limbic and pupil circles. The latter format was intended to support the compact storage of iris imagery on identity credentials.

The rectilinear version of the ISO standard was adopted by the law enforcement community in the ANSI/NIST Type 17 record. The polar format was not considered because compact size was not an imperative. This decision was fortuitous in light of subsequent technical contributions to SC37 working group 3 (Germany, N2059;

Table 3.1 The evolution of contemporary iris image standards

Date	Title of standard
05/2004	INCITS 379 – Iris image interchange format
06/2005	ISO/IEC 19794 – Biometric data interchange formats – Part 6: Iris image data
04/2007	ANSI/NIST-ITL 1-2007 Data format for the interchange of fingerprint, facial, and other biometric information Type 17
–/2011	ISO/IEC 19794 – Biometric data interchange formats – Part 6: Iris image data, second edition.
–/2011	ANSI/NIST-ITL 1-2007 data format for the interchange of fingerprint, facial, and other biometric information Type 17, second edition
–/2013	ISO/IEC 29794 – Biometric sample quality – Part 6: Iris image data
–/2013	ISO/IEC 29109 – Conformance testing methodology for biometric data interchange and formats defined in ISO/IEC 19794 – Part 6: iris image data

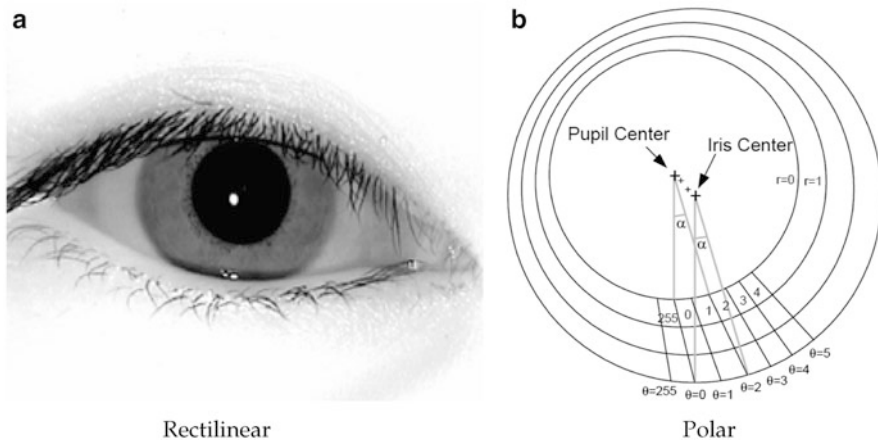


Fig. 3.1 ISO/IEC 19794-6 supports both rectilinear and polar storage formats. The *right* figure illustrates how pixels are sampled in the polar format. **(a)** Rectilinear. **(b)** Polar

Great Britain, N2124) that asserted the accuracy of the polar format is critically sensitive to the consistency of localization and subject to sampling problems [17]. The two documents advocated removal of the polar format from the standard. A useful GB contribution [8] suggested the compact size of the polar format (around 2 kB) could alternatively be achieved via cropping and compression of the rectilinear format. While interoperability problems were never formally demonstrated, the SC37 working group solicited quantitative evidence for the newly proposed formats. In the United States, the National Institute of Standards and Technology (NIST) was concurrently interested in establishing a set of specifications for the efficient transmission of iris images across networks and for storage on ISO/IEC 7816 cryptotokens [12]. To this end, NIST initiated the IREX study to evaluate the proposed image formats. The results of the IREX I evaluation lead to the adoption of compact iris storage formats in the ISO/IEC 19794-6 standard. The same formats were subsequently incorporated into the ANSI/NIST-ITL 1-2007 standard for law enforcement due to their potential for supporting network-based and mobile identification.

In 2009, NIST initiated the IREX II Iris Quality Calibration and Evaluation (IQCE) to support the development of a standard addressing iris image quality. Prior studies on the iris biometric [6, 15, 20] revealed that problems such as poor focus, motion-blur, occlusion by eyelids, and off-angle gazes negatively impact recognition accuracy. The new standard, ISO/IEC 29794-6: iris image quality, will formally define a vector of quality components for iris images, where each component will be a quantitative measure of a subject-specific or image-specific covariate. The evaluation also aims to promote the development of automated quality assessment algorithms. The ability to identify poor quality images quickly and effectively is particularly useful for acquisition systems that are tasked with identifying the best quality iris image among several potential candidates.

3.4 Storage Size Constraints

Many biometric systems operate with restrictions on the size of stored samples. For example, smart cards typically have a maximum storage capacity for their biometric samples. Additionally, smaller samples reduce the to-card and from-card transfer times, which can influence the selection of modality (e.g., fingerprint, iris) as well as the number of instances (e.g., two fingers, two irises). Centralized systems operating over limited bandwidth connections would also experience a benefit from reduced transfer times. Figure 3.2 depicts a network-centric application where field units (at left) pass compact KIND 3 and KIND 7 iris records across narrow bandwidth connections to a central server for enrollment. Compact iris records would transfer more quickly and require less storage space on the central server.

Broadly, there are two operational scenarios where compact storage formats are advantageous:

- Identity credential: A compressed standard iris image is stored on, for example, an ISO/IEC 7816 smart card, and is compared during authentication against a newly collected uncompressed sample. This scenario is representative of cooperative physical or logical access control situations in which the first sample is collected and prepared in an attended formal enrollment session, and the authentication sample exists only for the duration of the attempt. In this scenario, matching is performed between uncompressed verification samples and compressed enrollment samples.

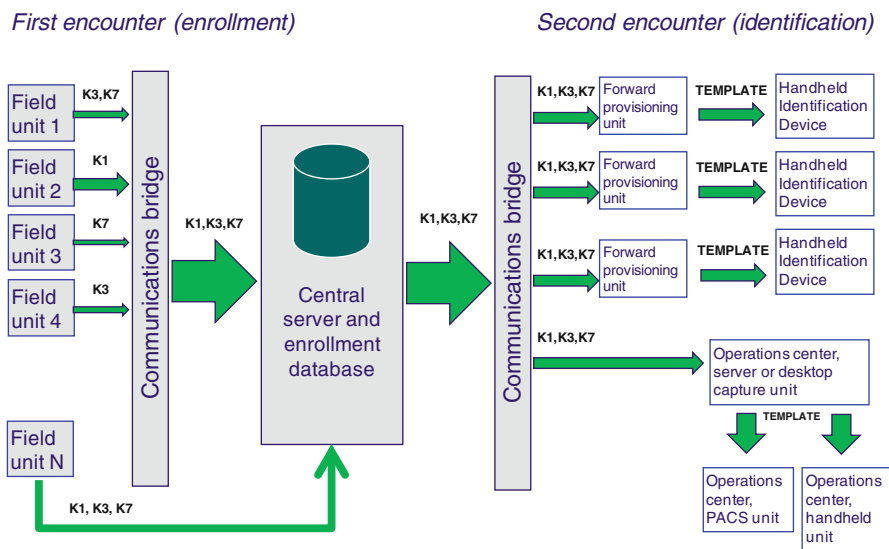


Fig. 3.2 Example of a network-centric application where standardized iris data is stored centrally and in the field. The left-to-right data flow represents the temporal sequence of enrollment followed by identification

- Central matching facility: Compressed standard samples are submitted to a central dataset. These are the first-encounter enrollment samples. Subsequently, compressed samples are transmitted to the central facility and are matched against the enrollments. Compression is implied by operational network bandwidth constraints. This scenario is typical in open-universe one-to-many applications such as visa fraud detection and watchlists. In this case, compression is applied to both authentication and enrollment samples.

3.5 Standard Storage Formats

Standardized iris images (i.e., iris records) are not iris templates. Rather, they are specialized interoperable images designed for the efficient storage and transmission of iris data. Templates contain proprietary “black box” feature representations that are typically specific to only a particular provider’s matching algorithm. As such, their content is nonstandard, non-interoperable, and not suitable for cross-agency exchange of iris data. The role of standardized images is depicted in Fig. 3.3. Iris segmentation precedes the second stage of processing in which features are extracted from iris records to form a template. The last stage, recognition, involves matching of templates to produce comparison scores.

IREX I evaluated three potential storage formats for iris images. Each is depicted in Fig. 3.4. Conversion of a raw iris image into any of the record types requires localization of iris features. Preparation of a KIND 3 record requires detection of the iris center and limbus boundary, which are then used to crop and center the iris. A KIND 7 record undergoes similar processing with the additional step of detecting the iris-eyelid and iris-sclera boundaries. The boundaries are used to apply a pixel masking operation to the regions outside of the iris. The purpose of the masking operation is to make it easier to compress the image to smaller storage sizes. For either format, the part of the image containing the iris features is left unchanged. A KIND 16 record stores the iris image in a polar format. Preparation of this record type involves defining concentric circles inside the pupil and outside the iris followed by a rectilinear-to-polar mapping. The polar format stores a fixed number of circumferential pixels at each radial distance, which effectively stores the inner part of the iris at a higher sampling rate (see Fig. 3.1b).




Fig. 3.3 Role of standardized imagery in a typical iris recognition system

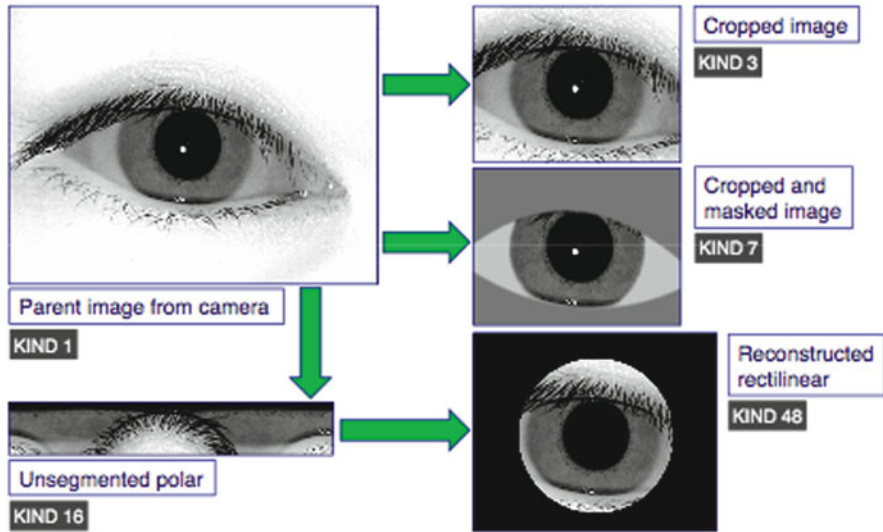


Fig. 3.4 Iris record formats tested in IREX I

3.6 Compression Formats

Compression is applied to images to reduce their storage size and transmission times. Compression algorithms are either lossless or lossy. Lossless compression maintains perfect fidelity to the original data. Lossy compression, on the other hand, discards some of the information, usually to achieve higher compression. When lossy compression is applied to an iris image, it alters the pixel values within the iris region. At low compression, the changes are inconsequential, but when compression is applied with sufficient severity, visible artifacts appear that can lead to recognition errors.

The ISO/IEC 19794-6 standard includes support for two methods of lossy image compression. The first, JPEG [19], is a popular and widely supported lossy storage format that uses a variety of methods to compress the image. Although optimized for digital photography, JPEG has proven useful for storing biometric samples for a range of modalities [9, 10, 16]. The second, JPEG 2000 [4], was developed to supersede JPEG and uses newly designed wavelet-based technology. Although less widely supported than its predecessor, JPEG 2000 has been generally shown to achieve superior compression [7, 14]. Figure 3.5 illustrates the types of artifacts that are introduced when iris images are heavily compressed. JPEG compression artifacts manifest as “macroblocking” and “mosquito noise,” while JPEG 2000 compression artifacts manifest as blurring and ringing around areas of sharp contrast.

Certain properties of an iris image determine how well it compresses to a particular storage size. Generally speaking, images with a lot of high gradient




Fig. 3.5 A single-iris record compressed to different file sizes. From *left to right*, the image is uncompressed and compressed to 4, 3, and 2 K. JPEG 2000 is applied to the *top row* images, and JPEG to the *bottom*. Beneath each image is a horizontal scan line of pixel intensities across the center of the image

information are more difficult to compress since they tend to lack the redundancy that compression algorithms exploit. Some of these high gradient features are irrelevant for the purposes of matching (e.g., eyelashes and skin texture) and can be discarded. The KIND 7 format was specifically formulated to improve the efficiency of compression by replacing the sclera, eyelids, and eyelashes with a uniform color, since regions of uniform color contribute minimally to the encoding space. A disadvantage of the KIND 7 format is that it requires additional processing and localization of a greater number of features around the iris.

3.7 Effect of Compression on Matching Accuracy

Several studies have investigated the effect of image compression on iris recognition accuracy. Ives et al. [13] applied JPEG 2000 compression to iris images and found that a compression ratio of 20 to 1 was possible without a discernible effect on impostor scores and only a minimal effect on genuine scores when using the Masek algorithm. Rakshit et al. [18] also applied JPEG 2000 compression to iris images and concluded that compressing the images to 0.5 bits-per-pixel actually improves ROC performance through noise reduction. This result could not be confirmed in IREX I. The authors also concluded that compressing the images using JPEG 2000 is generally preferable to downsampling to lower resolutions. Daugman [7] tested both JPEG and JPEG 2000 compression and found that, with the appropriate adjustments and compression parameters, samples could be reduced to as little as 2,000 bytes with little impact on accuracy. The conclusions from these studies are

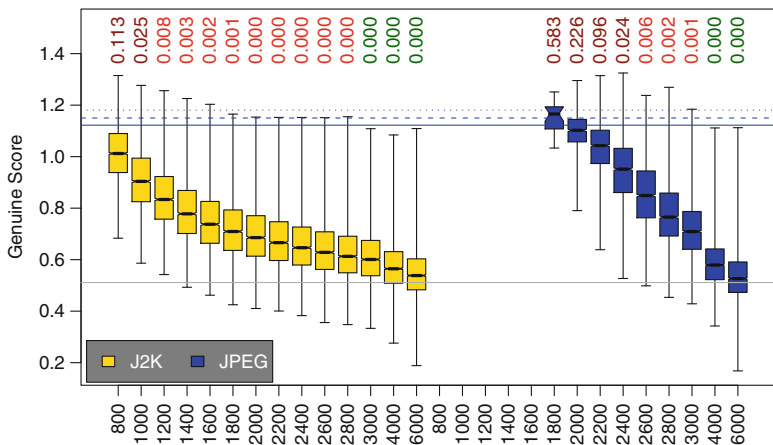


Fig. 3.6 Distribution of genuine scores for a matching algorithm when KIND 7 enrollment records were compressed to specific file sizes (specified along the vertical axis). The three blue lines correspond from the top to FMRs of $10^{(-2,-3,-4)}$, while the gray line refers to the median score for uncompressed images

based on data collected from one or two matching algorithms run over relatively small datasets. IREX I sought to verify some of the key results of these studies over much larger datasets and with matching algorithms from several providers.

IREX I explored the effect of compression on iris impostor comparisons. The evaluation found that the impostor distribution is generally invariant to JPEG 2000 compression for most algorithms. This is a desirable property since iris recognition systems can be easily calibrated to operate at known false match rates, even in the presence of varying amounts of compression. Indeed, a strength of the well-known IrisCode matching algorithm is that it has been shown to produce a stable and predictable impostor distribution under a variety of circumstances. For some algorithms, JPEG compression caused an increase in impostor scores. Although this typically only occurred at sample sizes below 2,000 bytes, it contraindicates the use of JPEG compression.

When the impostor distribution is stable, variations in matching accuracy are reflected in the genuine distribution. Results of the IREX evaluation demonstrate that compressed iris images are more likely to cause false non-matches. Figure 3.6 shows how progressively greater amounts of compression adversely affect the impostor distribution for one algorithm submitted to IREX I. Only enrollment images were compressed (as in the identity credential scenario). Typically, matching accuracy begins to degrade around 6,000 bytes and occurs more rapidly for JPEG compression. At higher compression ratios, the distribution becomes more long tailed, indicating the median genuine score alone does not fully capture possible changes in matching accuracy.

3.8 Compression Guidelines

Iris images can be compressed to as little as 2 kB, making them suitable for storage on ISO/IEC 7816-integrated circuit smart card identification tokens. However, the high amount of compression leads to a reduction in matching accuracy. IREX I provides guidance for the effective usage of standard iris storage formats. Application-specific recommendations on compression and format based on the storage limitations of the system are presented in Fig. 3.7. The evaluation also makes the following general recommendations and observations:

- JPEG 2000 is generally superior at compressing iris images than JPEG. When matching accuracy is paramount and storage space is limited to 20 kB or less, JPEG 2000 should be used.
- The cropped-and-masked KIND 7 image format should be retained and advanced as the primary format for the exchange of compact iris images smaller than 3 kB. At larger sizes or lower compression ratios, the KIND 3 format should be used since it is more easily and safely instantiated.
- Recognition error rates associated with the KIND 16 unsegmented polar format are much larger than those attained with the rectilinear formats. For this reason, IREX I recommended that the polar format be rejected from inclusion in the ISO/IEC 19794-6 standard.
- The KIND 7 format is best at preserving the iris texture when the image must be compressed to a small target record size. However, if storage space and network bandwidth are not limiting factors, the additional masking step required to convert an iris image to a KIND 7 record is unnecessary and can introduce problems if the eyelids are not correctly localized.



Fig. 3.7 Application-specific recommendations on compression and format from the IREX I evaluation

3.9 Summary

Biometric data standards define interchange formats for the storage and transmission of biometric data. Well-structured and widely implemented biometric data standards serve to facilitate the exchange of biometric data across different agencies and organizations and help promote fair competition in the marketplace. Both INCITS 379 and ISO/IEC 19794-6 define two rectilinear formats and one polar format for iris images. In 2009, NIST initiated the IREX I evaluation to develop and test compact iris storage formats. Such formats are important for applications where network bandwidth is limited or when iris records must be stored on smart card chips that have limited storage capacities. The evaluation determined that iris images can be compressed to as little as 2 kB, making them suitable for storage on ISO/IEC 7816-integrated circuit smart card identification tokens. If storage space is extremely limited (<3 kB per image), iris images should be stored as KIND 7 records using JPEG 2000 compression. When storage constraints are more relaxed, KIND 3 records are preferred since they are easier and safer to instantiate.

References

1. Working Group 3. ISO/IEC 19794 Biometric Data Interchange Formats (2005). http://www.iso.org/iso/home/store/catalogue_tc/catalogue_detail.htm?csnumber=50862
2. Working Group 3. ISO/IEC CD 19794-6 Information Technology – Biometric Data Interchange Formats – Part 6: Iris Image Data. International Standard. JTC1 :: SC37 (2006). http://www.iso.org/iso/home/store/catalogue_tc/catalogue_detail.htm?csnumber=50862
3. Working Group 5. ISO/IEC 10918-1 Digital compression and coding of continuous-tone still images: requirements and guidelines. International Standard. JTC1 :: SC29 (1994). <http://www.iso.org/>
4. Working Group 5. ISO/IEC 15444-1 JPEG 20000 image coding system: core coding system. International Standard. JTC1 :: SC29 (2004). <http://www.iso.org/>
5. Working Group 5. ISO/IEC 15948 Portable Network Graphics (PNG): functional specification. International Standard. JTC1 :: SC24 (2004). <http://www.iso.org/>
6. Abhyankar, A., Schuckers, S.: Iris quality assessment and bi-orthogonal wavelet based encoding for recognition. *Pattern Recognit.* **42**(9), 1878–1894 (2009)
7. Daugman, J.: Biometric decision landscapes. Technical Report TR482, University of Cambridge Computer Laboratory (2000)
8. Daugman, J., Downing, C.: Effect of severe image compression on iris recognition performance. Technical report, University of Cambridge, Computer Laboratory (2007)
9. Finger Image-Based Data Interchange Format. ANSI-INCITS 381-2004. Washington. <http://www.iso.org/> (2004)
10. Griffin, P.: Face recognition format for data interchange. Committee Draft ISO/IEC 19794-5 SC37 M1/03-0494 SC37 Document 342. Identix Corporate Research Center (2003)
11. Grother, P., Tabassi, E., Quinn, G.W., Salamon, W.: Performance of iris recognition algorithms on standard images. Technical report, National Institute of Standards and Technology (2009)
12. ISO/IEC 7816-15. International Standard. JTC1 :: SC17. <http://www.iso.org/> (2004)
13. Ives, R.W., Broussard, R.P., Kennell, L.R., Soldan, D.: Effects of image compression on iris recognition system performance. *J. Electronic Imaging* **17**(1), (2008)

14. Mascher-Kampfer, A., Stögner, H., Uhl, A.: Comparison of compression algorithms' impact on fingerprint and face recognition accuracy. In: Proceedings of the SPIE, vol. 6508, pp. 650810.1–650810.12. San Jose, CA (2007)
15. Matey, J., Naroditsky, O., Hanna, K., Kolczynski, R., Lolacono, D., Mangru, S., Tinker, M., Zappie, T., Zhao, W.: Iris on the move: acquisition of images for iris recognition in less constrained environments. In: Proceedings of the IEEE, vol. 94 (2006)
16. Onyshczak, R., Youssef, A.: Chapter 19: Fingerprint image compression and the wavelet scalar quantization specification. In: Ratha, N., Bolle, R. (eds.) Automatic Fingerprint Recognition Systems. Springer Berlin Heidelberg (2004)
17. Proença, H., Alex, L.A.: Iris recognition: An analysis of the aliasing problem in the iris normalization stage. In: International Conference on Computational Intelligence and Security, Guangzhou, vol. 2, pp. 1771–1774 (2006)
18. Rakshit, S., Monro, D.M.: An evaluation of image sampling and compression for human iris recognition. *IEEE Trans. Inf. Forensics Secur.* **2**(3–2), 605–612 (2007)
19. Wallace, G.K.: The jpeg still picture compression standard. *Commun. ACM* **34**(4), 30–44 (1991)
20. Wei, Z., Tan, T., Sun, Z., Cui, J.: Robust and fast assessment of iris image quality. In: Advances in Biometrics. Lecture Notes in Computer Science vol. 3832, pp. 464–471. Springer Berlin Heidelberg (2005)
21. Wilson, C., Grother, P., Chandramouli, R.: Biometric data Specification for personal Identity verification. Technical Report NIST Special Publication 800-76-1, National Institute of Standards and Technology. http://csrc.nist.gov/publications/nistpubs/800-76-1/SP800-76-1_012407.pdf (2007)

Chapter 4

Iris Quality Metrics for Adaptive Authentication

Natalia A. Schmid, Jinyu Zuo, Francesco Nicolo, and Harry Wechsler

Abstract Iris sample quality has a number of important applications. It can be used at a variety of processing levels in iris recognition systems, for example, at the acquisition stage, at image enhancement stage, or at matching and fusion stage. Metrics designed to evaluate iris sample quality are used as figures of merit to quantify degradations in iris images due to environmental conditions, unconstrained presentation of individuals or due to postprocessing that can reduce iris information in the data. This chapter presents a short summary of quality factors traditionally used in iris recognition systems. It further introduces new metrics that can be used to evaluate iris image quality. The performance of the individual quality measures is analyzed, and their adaptive inclusion into iris recognition systems is demonstrated. Three methods to improve the performance of biometric matchers based on vectors of quality measures are described. For all the three methods, the reported experimental results show significant performance improvement when applied to iris biometrics. This confirms that the newly proposed quality measures are informative in the sense that their involvement results in improved iris recognition performance.

4.1 Introduction

Assessment of iris image quality is one of important research venues recently identified in the field of iris biometrics [2, 13, 14]. Iris quality is evaluated by means of quality metrics. The main role of these metrics is to quantify, at the stage of data acquisition or at a later processing stage, what information an iris image contains

N.A. Schmid (✉) • J. Zuo • F. Nicolo
West Virginia University, Morgantown, WV, USA
e-mail: Natalia.Schmid@mail.wvu.edu; jzuo@mix.wvu.edu; fnicolo@mix.wvu.edu

H. Wechsler
George Mason University, Fairfax, VA, USA
e-mail: wechsler@cs.gmu.edu

to discriminate the iris class, which it represents, from all other iris classes in a database. In many cases, the quality metrics are used to decide if the image should be discarded or should be further processed by iris recognition system.

Quality of iris images is determined by many factors that can be broadly divided into two groups: (1) environmental and camera effects and (2) unconstrained presentation of a subject. As an example, images can be of low quality due to insufficient lighting, defocus blur, off-angle presentation, and heavy occlusion. These factors affect iris segmentation and later encoding and matching.

Previous works on iris image quality can be placed into two categories: local and global analysis. Zhu et al. [16] evaluated quality by analyzing the coefficients of particular areas of iris texture by employing discrete wavelet decomposition. Chen et al. [3] classified iris quality by measuring the energy of concentric iris bands obtained from 2-D wavelets. Ma et al. [11] analyzed the Fourier spectra of local iris regions to characterize out-of-focus and motion blur and occlusions. Zhang and Salganicoff examined the sharpness of the region between the pupil and the iris. Daugman [4] and Kang and Park [8] characterized quality by quantifying the energy of high spatial frequencies over the entire image region. Belcher and Du [1] proposed a clarity measure by comparing the sharpness loss within various iris image regions against the blurred version of the same regions. The major feature of these approaches is that the evaluation of iris image quality is reduced to the estimation of a single [3, 4, 8] or a pair of factors [11], such as out-of-focus blur, motion blur, and occlusion. A more comprehensive set of quality factors was introduced by Kalka et al. [6, 7], where seven factors that can affect iris image quality were identified: out-of-focus and motion blur, occlusion, specular reflection, illumination, off-angle, and pixel count. Recently, National Institute of Standards and Technology (NIST) identified 16 iris image properties that influence the recognition accuracy in Iris Exchange (IREX) II Iris Quality Calibration and Evaluation (IQCE) [15].

The components of a vector of quality measures, however, rarely carry equal weight in terms of their relationship to the performance of the matcher. In practical applications (such as US Visit program), it is required to keep a single biometric quality measure in order to decide if biometric samples are suitable for further processing and matching. Research questions should thus be concerned with (1) what quality measures to use, (2) how to combine multiple quality measures into a single quality index without losing the information that the vector of quality measures contains, and (3) how to use this vector to improve performance of biometric systems.

Most of the quality-based matchers described in the literature involve biometric sample quality at the matching stage by concatenating matching scores due to the original matcher and quality measures. These matchers are known as Q-stack classifiers [9, 12]. In spite of the fundamental theory presented in these works in support of Q-stack classifiers, the improvement of performance is marginal, if at all (see [5, 9]). More noticeable improvements are reported for Q-stack classifiers operating on multiple algorithms or multiple matchers [5, 12].

This chapter suggests several methods on the use of biometric sample quality to improve the performance of a single matcher. It targets two main applications for quality measures: (1) to improve performance of a matcher by predicting its quality of sample (QS) index or confidence in score (CS) and using them to decide if the underlying biometric sample should be retained or discarded and (2) to design a nonlinear matcher that treats a vector of quality measures as a set of weak features.

The remainder of this chapter is organized as follows. Section 4.2 introduces ten new quality factors for a segmented iris template. Section 4.3 evaluates performance of metrics summarized in Sect. 4.2. Section 4.4 describes three proposed methods to boost the performance using the estimated quality factors. Section 4.5 concludes this chapter.

4.2 Quality Evaluation

A traditional iris recognition system processes, segments, and encodes iris images sequentially. Simultaneously, local iris quality factors are evaluated. These quality factors can be later used to enhance performance of an iris recognition system alone or of a multimodal system with iris being one of the modalities. In the following subsections, new individual iris quality factors are introduced, and procedures to evaluate them are described. The factors are segmentation scores, interlacing, illumination, lighting, occlusion, pixel count, dilation, off-angle, and blur. The subsections below present a summary of the development, while detailed description is provided in [18–20].

4.2.1 Iris Segmentation Scores

Since most of local iris quality measures are applied to segmented iris images, the metrics evaluating the precision of the segmentation should be given a higher priority compared to other factors. Two segmentation scores Q_{p_seg} and Q_{i_seg} introduced in [17] can be used as two distinct quality metrics related to the segmentation itself. These metrics analyze the gradient values along the pupil and limbic boundaries. Larger value of the metrics indicates more precise segmentation.

To evaluate the cumulative intensity gradient along the estimated pupil and iris boundaries, the iris image (Fig. 4.1a) is unwrapped to the template (Fig. 4.1b) and its corresponding noise mask (Fig. 4.1c). First, a band covering the estimated pupil boundary (Fig. 4.1d) is analyzed. For each horizontal pixel in unwrapped image, we evaluate the intensity gradient along the vertical direction. If a gradient value for one of vertical pixels exceeds a specified threshold, we will say that at this location (horizontal location) the boundary is detected (Fig. 4.1e). We further count the number of horizontal locations with the gradient exceeding the threshold and divide by the total number of horizontal pixels in the unwrapped iris. A similar




Fig. 4.1 The steps performed by the automatic algorithm for evaluation of precision of iris segmentation: (a) the original image, (b) the unwrapped template, (c) the unwrapped mask, (d) the extended portion of unwrapped iris on the pupil side, (e) the result of the edge detection on the pupil side, (f) the extended portion of unwrapped iris on the sclera side, and (g) the result of the edge detection on the sclera side

evaluation is performed for the iris boundary (Fig. 4.1f). Since in majority of cases the limbic boundary is not sufficiently sharp, a vertically compressed version is used in place of the original unwrapped image to detect the edge (Fig. 4.1g). Denote by Q_{p_seg} and Q_{i_seg} the percentage of horizontal pixels with the gradient exceeding the prespecified threshold where Q_{p_seg} is for the pupil boundary and Q_{i_seg} is for the limbic boundary. However, occluded parts will not be taken into account. The larger the values of these parameters, the better are the estimates of the boundaries.

4.2.2 Interlacing

Poor interlacing is a disturbing artifact. Interestingly, a poorly interlaced image may result in a high focus score in spite of strong defocus of either even or odd lines. These artifacts should be detected. If there is a large difference between odd and even lines, then there must be a clear motion related interlacing effect. The image

may be either discarded or divided to two sub-images: odd rows and even rows. The difference between odd rows and even rows Inter_1 can be calculated as

$$\text{Inter}_1 = \frac{\sum_{i=1:2:m-1} 2 \sum_{j=1}^n |I_{i,j} - I_{i+1,j}|}{m \times n} \quad (4.1)$$

for an image I with m rows and n columns. The function must be normalized by subtracting Inter_2 calculated using odd or even rows only

$$\text{Inter}_2 = \frac{\sum_{i=1}^{m-2} \sum_{j=1}^n |I_{i,j} - I_{i+2,j}|}{(m-2) \times n} \quad (4.2)$$

resulting in

$$Q_{\text{interlacing}} = \text{Inter}_1 - \text{Inter}_2. \quad (4.3)$$

Note that the high values of $Q_{\text{interlacing}}$ indicate poor interlacing.

4.2.3 Illumination

The contrast of the image is mainly determined by the level and strength of the illumination. The illumination level is the mean intensity value of the iris area:

$$Q_{\text{illumination}} = \frac{\sum_{\text{unaffected iris area}} I_{i,j}}{\sum_{\text{unaffected iris area}} 1}. \quad (4.4)$$

To get a more precise estimation of this factor, only unaffected (by occlusion or specular reflections) area is considered. The value that $Q_{\text{illumination}}$ can take ranges from 0 to 255. This factor can be affected by the color of the iris. Large values of the measure indicate high illumination value.

4.2.4 Lighting

Sided or uneven illumination of the iris often results in performance degradation. Illumination pattern can be treated as a low-frequency signal that distorts encoded iris images. The variance of the mean intensity evaluated over small blocks is proposed as a measure of the uneven illumination. The calculation of the lighting factor is similar to the procedure described in [7] but without normalization to $[0, 1]$. Note that bad lighting condition is characterized by a large value of the metric.

4.2.5 Occlusion

This attribute measures the fraction of the iris area occluded by other objects such as eyelids, eyelashes, and specular reflections. The proposed metric evaluates the percentage of the unoccluded area in the unwrapped template. Denote by M the binary occlusion mask of the unwrapped iris template. The value “1” at location (i, j) in the mask indicates that it is a point of occlusion. Then the occlusion metric is

$$Q_{\text{occlusion}} = \frac{\sum_{\{(i,j): M_{i,j}=0\}} 1}{\sum_{\{(i,j): M_{i,j} \geq 0\}} 1}. \quad (4.5)$$

The usage of the percentage can reduce the correlation between this quality factor and the resolution factor. This quality factor is similar to the pixel count factor in [7]. Large values of the metric indicate smaller occlusions.

4.2.6 Pixel Count

To distinct it from the occlusion factor, pixel count finds the total iris area including its occluded area:

$$Q_{\text{pixel count}} = \sum_{\text{iris area}} 1. \quad (4.6)$$

Large values of the metric correspond to high pixel counts.

4.2.7 Dilation

The dilation factor measures the degree of the pupil dilation. The value of the dilation factor is calculated by taking the ratio of $Q_{\text{pixel count}}$ to the total iris and pupil area:

$$Q_{\text{dilation}} = \frac{\sum_{\text{iris area}} 1}{\sum_{\text{iris area}} 1 + \sum_{\text{pupil area}} 1}. \quad (4.7)$$

The metric Q_{dilation} takes values between 0 and 1. This factor also affects pixel count. Note that small pupil dilations are characterized by large values of the metric.

4.2.8 Off-Angle

This factor measures the relative orientation of the iris with respect to the camera. Assuming that the frontal view iris has a circular shape, the off-angle view becomes an ellipse. The off-angle quality factor is a ratio of the two main axes of the ellipse fitted into the iris boundary:

$$Q_{\text{off angle}} = \frac{b}{a}, \quad (4.8)$$

where b is the minor axis and a is the major axis of an ellipse. These values are obtained after the iris has been segmented. Note that the large values of the metric indicate that the image is close to frontal view.

4.2.9 Blur

Both motion and defocus blurs are treated simultaneously. This method uses spectral components of an iris image and involves a number of preprocessing steps.

First, the area of interest (iris region) is segmented, and an ellipse with the center at (x_i, y_i) and the major axes a and b is fitted into the iris region. We then expand the image around the area of interest by 250% along the major axis a and choose the iris center (x_i, y_i) as the center of the expanded image.

A small median filter is then applied to denoise the image. To compensate the difference in resolution, the area of interest is normalized to the size 151×151 . This selection is due to resolution requirement imposed by standards: 120 or more pixels across the iris. A 2-D FFT transform is applied to normalized image in order to extract the frequency information denoted here by P :

$$P = \log_{10} |\text{FFT}(I_{\text{crop}})|, \quad (4.9)$$

where I_{crop} is the cropped iris area after normalization.

The power distribution of P is analyzed, and its central area is used to calculate a decision threshold. Here the threshold is empirically selected. We select the average power of a centered 13 pixel diamond-shaped area as the value of the threshold γ :

$$\gamma = \frac{\sum_{\{(i,j): (i-76)^2+(j-76)^2 < 16\}} P_{i,j}}{13 \times 1.5}. \quad (4.10)$$

Then the number of pixels with the power exceeding the threshold value is counted. This number is scaled by a coefficient involving the dilation information. The final expression for the Q_{blur} becomes

$$Q_{\text{blur}} = (1 + Q_{\text{dilation}}^6) \sum_{P_{i,j} > \gamma} 1. \quad (4.11)$$




Fig. 4.2 Example of blur estimation: (a) the cropped area, (b) the power spectrum, and (c) the results of thresholding

Steps illustrating the evaluation of the blur quality score are shown in Fig. 4.2. This specific image has the blur score equal to 5,953.6.

Note that larger values of the metric correspond to a smaller amount of blur.

In the following two sections, the application of the designed quality measures is two fold: (1) we show that good quality images can be separated from poor quality images by involving the proposed quality measures, and (2) we use quality measures to design quality-based adaptive iris authentication systems.

4.3 Performance of Quality Metrics

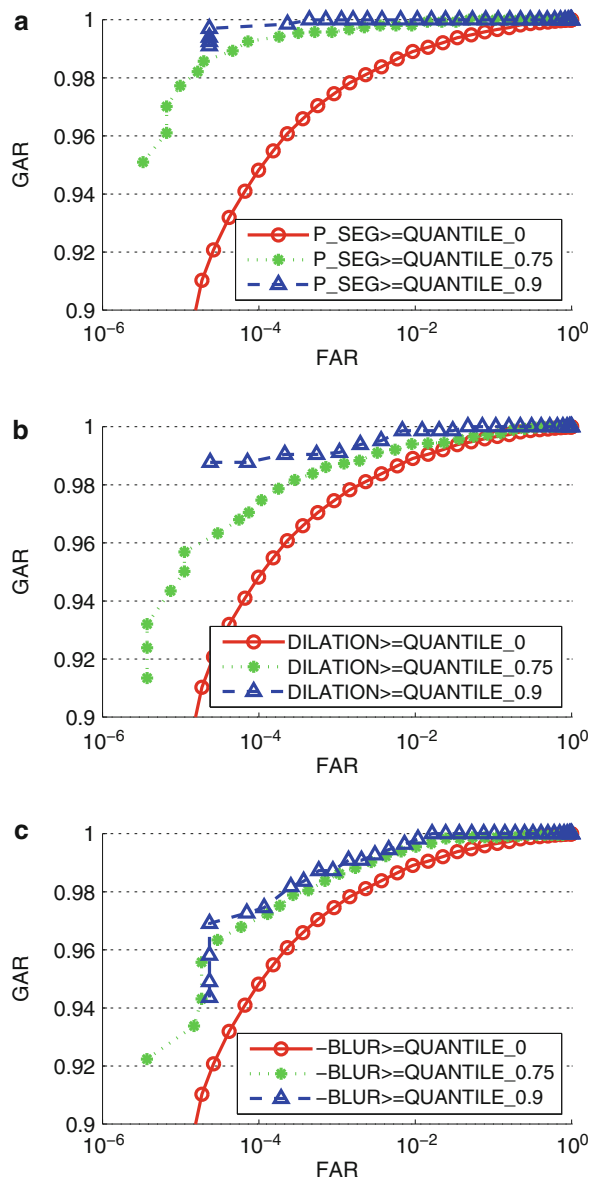
All experiments are performed using ICE 2005 dataset [10]. The enhancement, encoding, and matching procedures follow Daugman's implementation.

We perform a number of experiments. For each individual factor, the ICE 2005 dataset is used to form three subsets of images. The first subset was composed of the entire ICE dataset. To form the second and the third subsets, we involve the distribution of values of a selected quality factor. The second set includes all images with the value of selected quality factor exceeding 0.75th quantile. The third set is composed of all images with the value of selected quality factors exceeding 0.9th quantile.

The panels in Figs. 4.3 and 4.4 each displays three receiver operating characteristic (ROC) sets obtained using data in subsets 1, 2, and 3. Note that all results can be placed into those based on a relative quality score (in our case, it is the difference of two quality values for two distinct images) and those based on an absolute measure. Examples of relative measures include interlacing, illumination, pixel count, and off-angle (Fig. 4.4). The other measures are used as absolute.

From Figs. 4.3 and 4.4 regardless of the type of the measure, the difference between ROCs formed from the three subsets of ICE 2005 dataset are quite noticeable. This indicates that each individual factor proposed in this work does influence recognition performance of a Gabor filter-based system.

Fig. 4.3 ROC curves for ICE2005 dataset (a) selecting images using pupil segmentation score, (b) selecting images using dilation measure, and (c) selecting images using minus blur measure



4.4 Quality-Based Adaptive Authentication

This section suggests another use of quality factors. Three different methods to enhance performance of biometric systems due to quality factors are presented. In the first two methods, the quality factors are combined to produce a single

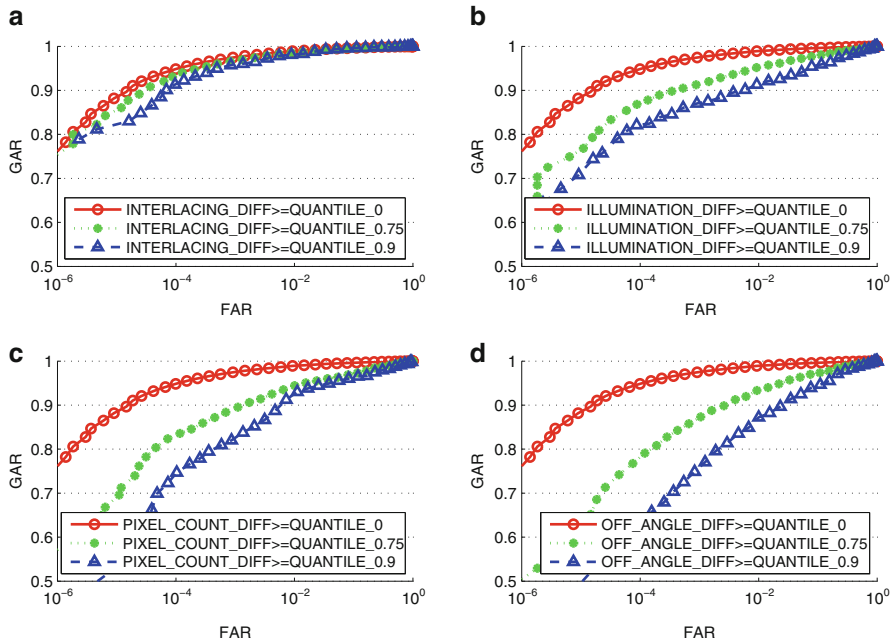


Fig. 4.4 ROC curves for ICE2005 dataset (a) selecting matching scores using interlacing measure, (b) selecting matching scores using illumination measure, (c) selecting matching scores using pixel count measure, and (d) selecting images using off-angle measure

metric either characterizing the overall quality of an iris image or characterizing the confidence level of a matching score. The third method treats quality factors as additional weak features that improve performance of the matcher. In each of the three methods, the functional relationship between vectors of quality measures and the predicted (estimated) measures is not known. The problem of estimating this relationship is stated as a multivariate regression problem:

$$Y = f(X_1, \dots, X_K), \quad (4.12)$$

where f is a multivariate adaptive mapping, variable Y is the output variable, and X_1, \dots, X_K is a vector of K input variables. To estimate $f(\cdot)$, we involve a set of labeled training data. The multivariate adaptive mapping $f(\cdot)$ can be implemented using a variety of multivariate functions and systems. The results reported below are obtained using a feedforward neural network (FFNN). Once the function $f(\cdot)$ is estimated, it is used to predict the output variable from unlabeled input data.

4.4.1 Quality of Sample

Here the vector of quality factors introduced in Sect. 4.2 is used to generate a single quality index by fusing the entries of the vector. Since the index has to be indicative of the performance of iris matcher, we propose to use d-prime index as the combined quality index. This simple single value performance measure requires only a small amount of labeled data per iris class to estimate it. Denote by QS_A the overall quality index of sample A . Then the expression for QS_A is given by

$$QS_A = \frac{|m(\text{Imp. Scores})_A - m(\text{Gen. Scores})_A|}{\sqrt{\text{var}(\text{Imp. Scores})_A + \text{var}(\text{Gen. Scores})_A}}, \quad (4.13)$$

where $m(\cdot)_A$ and $\text{var}(\cdot)_A$ are the sample mean and sample variance of genuine and imposter scores formed by involving the sample A .

Using the labeled training set, QS is estimated for every biometric sample. This requires that a set of genuine matching scores and a set of imposter matching scores involving the same biometric sample be formed. For unlabeled biometric samples, this task becomes almost impossible. However, having quality vectors associated with each biometric sample makes it possible to *predict* the QS of unlabeled data. The QS of an unlabeled sample can be obtained as the output parameter of a nonlinear multivariate adaptive mapping applied to a vector of quality measures from the same sample. Let $\mathbf{Q}_A = [Q_{A,1}, \dots, Q_{A,K}]^T$ be a vector of K quality measures characterizing a biometric sample A . The superscript T indicates the transpose operation. Let $f_{QS}(\cdot)$ be a nonlinear multivariate adaptive mapping that maps a vector of quality measure \mathbf{Q}_A into the quality index QS_A . Let $\hat{f}_{QS}(\cdot)$ be its estimated version. Then QS_A is predicted as $QS_A = \hat{f}_{QS}(\mathbf{Q}_A)$. The predicted value of QS can be then used to decide if the underlying biometric sample should be retained or discarded to improve the performance of the original matcher.

4.4.2 Confidence in Scores

The second method evaluates the confidence level assigned to matching scores associated with a pair of biometric samples. The confidence in genuine and imposter scores (CS) is defined as follows:

$$CS_G = \begin{cases} 0, & HD_G < Q(HD_G)_x, \\ -\frac{HD_G - Q(HD_G)_x}{Q(HD_G)_y - Q(HD_G)_x}, & \text{otherwise,} \end{cases} \quad (4.14)$$

$$CS_I = \begin{cases} 0, & HD_I > Q(HD_I)_{1-x}, \\ -\frac{Q(HD_I)_{1-x} - HD_I}{Q(HD_I)_{1-x} - Q(HD_I)_{1-y}}, & \text{otherwise,} \end{cases} \quad (4.15)$$




Fig. 4.5 Illustration of the confidence in scores (CS)

where $Q(\text{HD}_G)_x$ and $Q(\text{HD}_I)_y$ are the quantile points at the quantile x and y for genuine and imposter scores, respectively. The levels of the quantiles were optimized empirically. The values resulting in significantly improved verification performance are $x = 0.7$ and $y = 0.9$. Figure 4.5 illustrates genuine and imposter distributions typical for Hamming distances (HDs).

The CS of an unlabeled biometric sample is predicted using a nonlinear adaptive mapping and a vector of quality measures (input parameters). The nonlinear adaptive mapping is trained using a set of labeled data in the form of vectors of quality measures and the corresponding CS values obtained using Eqs. (4.14) and (4.15). At the testing stage, the CS value is predicted based on a vector of quality measures only. This information is used to keep or discard the corresponding matching score. The procedure of predicting the CS of matching scores between two biometric samples A and B is $\text{CS}_{AB} = \hat{f}_{\text{CS}}(\mathbf{Q}_A, \mathbf{Q}_B)$, where $\hat{f}_{\text{CS}}(\cdot)$ is a multivariate adaptive mapping (FFNN in our case) estimated using training data.

4.4.3 Quality Sample and Template Features

The third method suggests to treat quality vectors as weak features that can be combined with biometric template features. Let \mathbf{Q}_A and \mathbf{C}_A be a vector of K quality measures and a template vector associated with a biometric sample A . Then the extended template, denote it by \mathbf{F}_A , of the sample A is the vector $\mathbf{F}_A = [\mathbf{Q}_A^T, \mathbf{C}_A^T]^T$.

When two biometric samples A and B are compared, the distance (or similarity) between \mathbf{C}_A and \mathbf{C}_B will be saved as MS_{AB} , while the quality vectors \mathbf{Q}_A and \mathbf{Q}_B will be treated as extra dimensions that may improve performance of the original matcher, provided these dimensions contain sufficient discriminative information.

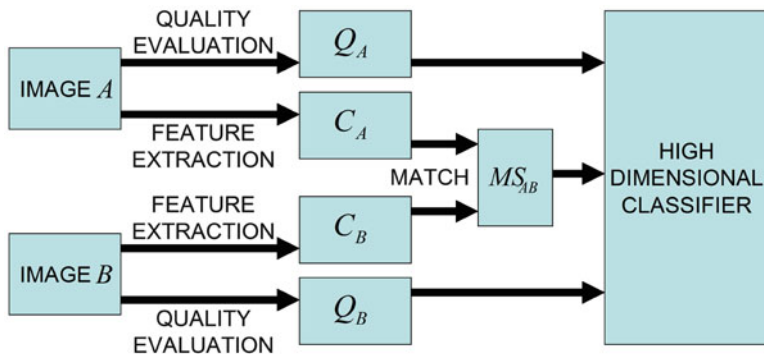


Fig. 4.6 Combining quality vectors and templates

A block diagram of the proposed quality enhanced matcher is shown in Fig. 4.6. The input to the high-dimensional classifier is a vector $[\mathbf{Q}_A^T, \mathbf{Q}_B^T, \text{MS}_{AB}]^T$. The output of the classifier is a decision made by the classifier. The decision is a binary-valued variable corresponding to $\{\text{genuine}, \text{imposter}\}$. The high-dimensional classifier is implemented using a nonlinear adaptive mapping. Denote $\eta_{A,B}$ as the output variable predicted using vectors of the quality measures and the matching score of the biometric samples A and B . Then the prediction procedure is described as $\eta_{A,B} = \hat{f}_{\text{QST}}(\mathbf{Q}_A^T, \mathbf{Q}_B^T, \text{MS}_{AB})$, where \hat{f}_{QST} is an estimated version of f_{QST} .

4.4.4 Experimental Results

All experiments were performed using ICE 2005 dataset and processing steps described in Sect. 4.3. We generated 26,867 genuine matching scores and 4,331,761 imposter matching scores from 2,953 iris biometric samples.

4.4.4.1 Neural Network

The nonlinear mapping is implemented using feedforward neural network (FFNN). Training data are assigned labels according to the functional use of the mapping and a set of input and output parameters. The final design is achieved by trading off the complexity and the performance of the network with two hidden layers. For the iris experiments, the first hidden layer of the FFNN is composed of 16 neurons, while the second layer is composed of 2 neurons. The training data are divided randomly in two subsets: a learning subset composed of 60% of training data and a validation subset made of remaining 40% of data. The training process stops when the mean square error drops below 10^{-4} . The experimental results described below are obtained using codes from the Neural Network Toolbox in MATLABTM.

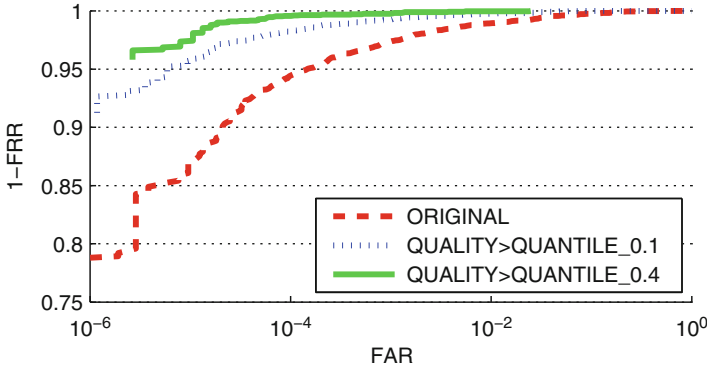


Fig. 4.7 Performance improvement achieved by selecting only images with high QS values

4.4.4.2 QS Evaluation

The performance of the QS method is evaluated by randomly selecting 1,500 iris images from the ICE 2005 dataset to form the training set. The remaining 1,453 iris images are used to form the testing set. The QS of unlabeled images from the testing set is predicted based on the quality vector (quality factors 1–10) and by using an FFNN trained on labeled samples. During performance evaluation, unlabeled images with the value of predicted QS above a preset quantile are retained. Figure 4.7 displays three receiver operating characteristic (ROC) curves parameterized by 0, 10, and 40% quantile levels. The ROC curve marked as “original” is parameterized by zero quantile level, which means that no poor quality biometric samples were discarded. Note that by discarding only 10% of iris images with the low predicted QS index, a considerable performance improvement can be achieved. Figure 4.8 displays a box plot of the equal error rate (EER) as a function of the quantile used to select iris samples with high QS value. It is a summary of ten independent trials, where training and testing data are sampled at random. It can be observed that regardless of the composition of training and testing data, removing iris images characterized by low predicted QS improves matching performance of the original matcher. The higher the value of QS is, the better the performance is.

4.4.4.3 CS Evaluation

To assess the performance of the CS method, 20,000 genuine matching scores, 200,000 matching imposter scores, and the vectors of quality measures associated with iris images were used to train an FFNN. The remaining data were used for testing. The matcher was designed to be symmetric with respect to quality vectors, that is, if \mathbf{Q}_A and \mathbf{Q}_B are two vectors of quality measures associated with iris image A and B , training included both the pair $(\mathbf{Q}_A, \mathbf{Q}_B)$ and the pair $(\mathbf{Q}_B, \mathbf{Q}_A)$ and the associated matching score. The testing experiment is similar to the experiment of

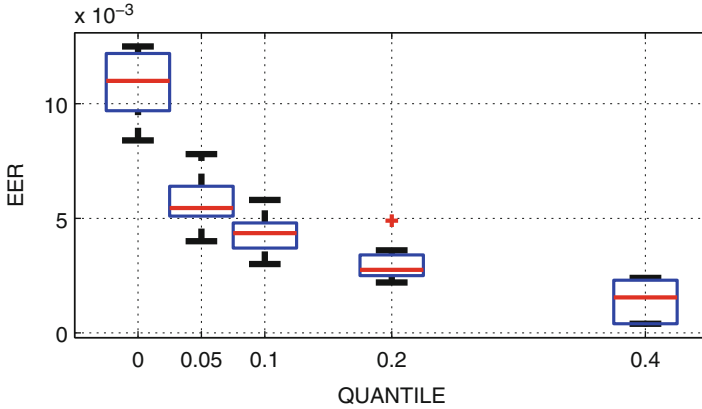


Fig. 4.8 Performance improvement by selecting only images with a quality value larger than a certain quantile

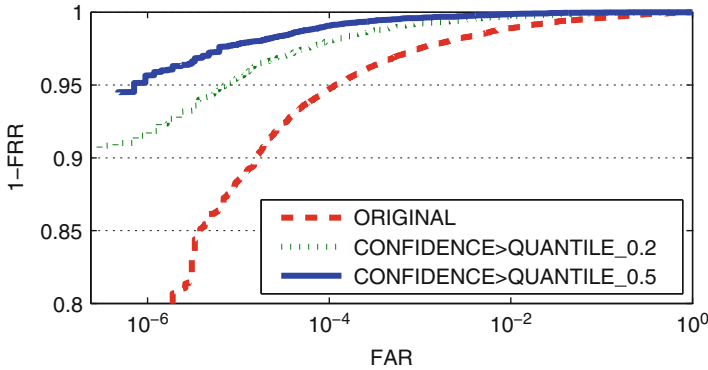


Fig. 4.9 Performance improvement achieved by selecting only matching scores with high CS values

the previous subsection with the difference that pairs of quality vectors are used to predict the CS values.

Figure 4.9 shows three ROC curves: the original curve, the curve formed from iris data with the predicted CS values exceeding 20% quantile, and the curve formed from iris data with the predicted CS values exceeding 50% quantile. Performance improves when low-confidence matching scores are discarded. Figure 4.10 summarizes the results of ten trials. Again, training set is formed by randomly sampling iris images from a larger set. The trends and results are consistent.

4.4.4.4 QST Evaluation

The matcher is now an FFNN trained and tested as follows. During training, the label “1” is assigned to all genuine vectors on the input, and the label “-1” is

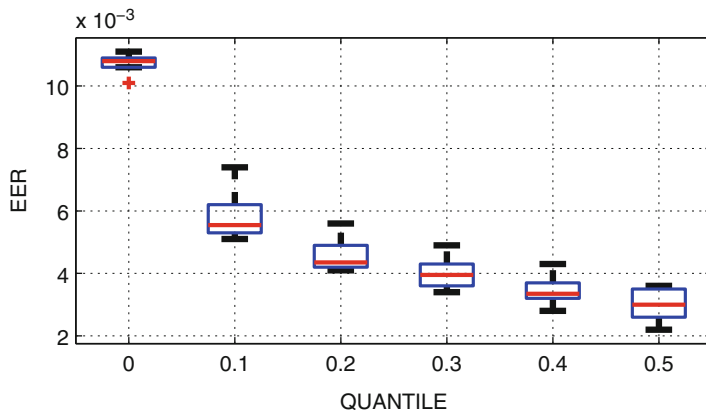


Fig. 4.10 Performance improvement by selecting only matching scores with a confidence level higher than a certain quantile

Table 4.1 Performance with/without quality factors

	Original EER	New EER	Original FRR at 0.001 FAR	New FRR at 0.001 FAR
Trial 1	0.0107	0.0065	0.0251	0.0197
Trial 2	0.0110	0.0076	0.0246	0.0186
Trial 3	0.0109	0.0072	0.0251	0.0192
Trial 4	0.0109	0.0068	0.0249	0.0193
Trial 5	0.0106	0.0086	0.0244	0.0229

assigned to all imposter vectors on the input to the neural network. During testing, the output label is predicted based on the input vector of quality metrics and the original matching score. The output label in this case is a real number. The high-dimensional classifier makes decision in favor of genuine class if the output label is close to 1. It decides in favor of imposter, if the output label is closer to -1 . When the decision threshold varies, the performance of the high-dimensional classifier is characterized by the ROC curve.

To assess the performance of the QST method, 20,000 genuine scores and 200,000 imposter scores and associated quality vectors were involved in training. The remaining vector-triplets were used for testing. The success of the QST method depends on the data selected for training and testing. A single trial out of set of 20 trials resulted in a perfect separation of genuine and imposter matching scores. In the other cases, the % improvement was between 20 and 35%. The results of five first trials are shown in Table 4.1. In columns 2 and 3, it displays the values of EER without/with quality factors. In columns 3 and 4, the table displays the values of false reject rate (FRR) evaluated at 0.001 false accept rate (FAR).

4.5 Conclusions

This chapter introduced a number of new absolute and relative (global and local) quality factors for iris images. The performance of the proposed measures was evaluated by analyzing the relationship between the quality of iris images and verification performance of the system (in terms of ROC curves). These relationships indicate that proposed quality measures have good selectivity in terms of recognition performance.

Using quality factors as a means to characterize the overall iris image quality, the confidence in matching score or to treat them as weak features is an alternative application of the quality factors. Three methods for matching iris biometrics using quality metrics are presented. The methods are adaptive and use nonlinear mappings for making predictions on quality measures and corresponding verification scores. The reported experimental results illustrate the importance of predictive and selective integration of quality measures for biometric authentication and show significant advantages compared to existing methods.

References

1. Belcher, C., Du, Y.: Information distance based selective feature clarity measure for iris recognition. In: Proceedings of the SPIE Symposium on Defense and Security. Conference on Human Identification Technology IV, Orlando, FL, vol. 6494, pp. 64,940 E1–E12 (2007)
2. Bowyer, K.W., Hollingsworth, K., Flynn, P.J.: Image understanding for iris biometrics: a survey. *J. Comput. Vis. Image Underst.* **110**, 281–307 (2008)
3. Chen, Y., Dass, S.C., Jain, A.K.: Localized iris image quality using 2-D wavelets. In: Proceedings of International Conference on Biometrics (ICB'06), Hong Kong, pp. 373–381 (2006)
4. Daugman, J.: Biometric personal identification system based on iris analysis. US Patent 5,291,560, March 1994
5. Fronthaler, H., Kollreider, K., Bigun, J., Fierrez-Aguilar, J., Alonso-Fernandez, F., Ortega-Garcia, J., Gonzalez-Rodriguez, J.: Fingerprint image-quality estimation and its application to multialgorithm verification. *IEEE Trans. IFS* **3**(2), 331–338 (2008)
6. Kalka, N.D.: Image quality assessment for iris biometric. Master's thesis, West Virginia University, Morgantown, WV (2005)
7. Kalka, N.D., Zuo, J., Schmid, N.A., Cukic, B.: Image quality assessment for iris biometric. In: SPIE 6202: Biometric Technology for Human Identification III, vol. 6202, Morgantown, WV, pp. 6202:D1–D11 (2006)
8. Kang, B.J., Park, K.R.: A study on iris image restoration. In: International Conference on Audio- and Video-Based Biometric Person Authentication, Rye Brook, NY, pp. 31–40 (2005)
9. Kryszczuk, K., Drygajlo, A.: Improving classification with class-independent quality measures: *Q-stack* in face verification. In: Proceedings of ICB 2007, Korea University, Seoul, pp. 1124–1133 (2007)
10. Liu, X., Bowyer, K.W., Flynn, P.J.: Experiments with an improved iris segmentation algorithm. In: Fourth IEEE Workshop on Automatic Identification Technologies (AutoID), Buffalo, NY, pp. 118–123 (2005)
11. Ma, L., Tan, T., Wang, Y., Zhang, D.: Personal identification based on iris texture analysis. *Pattern Recognit.* **25**(12), 1519–1533 (2003)

12. Nandakumar, K., Chen, Y., Dass, S.C., Jain, A.K.: Likelihood ratio-based biometric score fusion. *IEEE Trans. PAMI* **30**(2), 342–347 (2008)
13. National Institute of Standards and Technology (NIST): Biometric Quality Workshop (2006). http://www.nist.gov/itl/iad/ig/bio_quality_wkshopi.cfm
14. National Institute of Standards and Technology (NIST): Biometric Quality Workshop II (2007). http://www.nist.gov/itl/iad/ig/bio_quality_wkshopii.cfm
15. National Institute of Standards and Technology (NIST): Iris Exchange (IREX) II: Iris Quality Calibration and Evaluation (IQCE) (2011). <http://www.nist.gov/itl/iad/ig/irexii.cfm>
16. Zhu, X., Liu, Y., Ming, X., Cui, Q.: A quality evaluation method of iris images sequence based on wavelet coefficients in “region of interest”. In: Proceedings of the Fourth International Conference on Computer and Information Technology (CIT’04), pp. 24–27. IEEE Computer Society, Washington, DC (2004)
17. Zuo, J., Schmid, N.A.: An automatic algorithm for evaluating the precision of iris segmentation. In: International Conference on Biometrics: Theory, Applications, and Systems (BTAS’08), Washington, DC (2008)
18. Zuo, J., Schmid, N.A.: Global and local quality measures for NIR iris video. In: Proceedings of the 2009 IEEE CVPR Workshops, Miami, FL, pp. 1–6 (2009)
19. Zuo, J., Schmid, N.A.: On a methodology for robust segmentation of nonideal iris images. *Systems Man Cybern. Pt. B* **40**(3), 703–718 (2010)
20. Zuo, J., Nicolo, F., Schmid, N.A., Wechsler, H.: Adaptive biometric authentication using nonlinear mappings on quality measures and verification scores. In: 17th IEEE International Conference on Image Processing (ICIP’10), Hong Kong, pp. 4077–4080 (2010)

Chapter 5

Quality and Demographic Investigation of ICE 2006

P. Jonathon Phillips and Patrick J. Flynn

Abstract Four major experimental evaluations of iris recognition technology have been conducted in recent years. These four experimental evaluations have many differences in design and implementation, but in the areas where the specifications of the studies intersect, the results are consistent. In the ICE 2006 evaluation, participants were allowed to submit quality measures for iris images. This chapter investigates the properties of the quality measures submitted as part of the ICE 2006 evaluation.

5.1 Introduction

There have been four major experimental evaluations of iris recognition technology in recent years: the ITIRT evaluation conducted by the International Biometric Group, the Iris 2006 evaluation conducted by Authenti-Corp, and the Iris Challenge Evaluation (ICE) 2006 and Iris Exchange (IREX) conducted by the National Institute of Standards and Technology. These experimental evaluations employed different vendor technologies and experimental specifications but yield consistent results in the areas where the specifications intersect. Thus, these studies jointly established an overall performance level for iris recognition [1, 3, 5, 6, 8]. Overall

P.J. Phillips (✉)
National Institute of Standards and Technology, Gaithersburg, MD, USA
e-mail: jonathon@nist.gov

P.J. Flynn
University of Notre Dame, IN, USA
e-mail: flynn@nd.edu

performance on a dataset does not give a complete characterization of performance. For real-world applications, scientific knowledge, and technology advancement, it is necessary to have a finer-grained characterization of performance.

Formally speaking, a covariate is a variable that is potentially predictive of an outcome. The focus of this study is the effect of quality measurements (covariates) on iris recognition system performance (the outcome). In this context, covariates that can affect performance fall into two categories. The first category is subject covariates. These are properties of an individual or of an iris. Examples are sex, race, eye color, the presence/absence and type of corrective lenses worn during acquisition, and the eye (left or right). The second category of covariates contains image covariates. These are properties of the image of the iris and can vary from image to image of the same iris. Examples of these covariates include iris size (e.g., diameter in pixels), iris dilation (e.g., radius of pupil in pixels), degree of occlusion by eyelids or eyelashes, and focus quality.

Quality metrics have a special place in iris recognition. Iris sensors employ embedded algorithms for quality assessment that are used to discard images and (if necessary) trigger reacquisition. The precise nature of the quality metrics used is unknown. This quality metric is often used to prescreen iris images and trigger reacquisition in the event that quality is judged to be poor. One can identify and casually define several potential degradations to the subjective quality of an iris image, such as defocus, motion blur, strong occlusion, large pupil size, off-axis eye pose (i.e., nonzero gaze angle relative to the optical axis), the presence of corrective lenses and lens reflections, contact lens registration marks or brands, refractive artifacts from corrective lenses, spoofing attempts using cosmetic lenses, surgical consequences such as deformed irises, pathological conditions such as cataracts, iris nevi (sometimes presenting as eye carcinoma), aniridia (lack of irises), or albinism (melanin deficiency). Not all of these conditions have been studied at length in the literature; few are amenable to rigorous analysis today due to the paucity of data presenting the conditions involved. Although one can subjectively judge iris image quality and perhaps implement software that parallels a human judgment, it does not follow that iris recognition system performance on such data will degrade in tandem with the subjective judgment.

In the ICE 2006, participants were allowed to submit quality measures. For the three iris algorithms reported in Phillips et al. [8], we investigate the properties of their quality submissions. To look at the question of “are the three quality measures the same thing?” we looked at the correlation among the three quality measures. We investigated the effect of demographics on the quality measures. Finally, we investigated the ability of the quality measures to predict performance. For numerous applications, this is the acid test of the effectiveness of quality measures.

5.2 ICE 2006: Methodology

5.2.1 Data

The ICE 2006 data was collected at the University of Notre Dame as part of a larger multiple biometric data collection effort [8]. The images were acquired between January 2004 and May 2005. The iris images were collected with a LG EOU 2200 that allowed for the collection of images with a broader range of quality than the sensor would normally acquire. This includes iris images that did not pass the quality control software embedded in the LG EOU 2200. The LG EOU 2200 is a complete acquisition system and has automatic image quality control checks.

The image quality software embedded in the LG EOU 2200 is one of the numerous iris quality measures. Prior to the start of the data collection, an arrangement was made to minimize the effect of the LG EOU 2200 quality screening software on the data collection. By agreement between University of Notre Dame and Iridian, a modified version of the acquisition software was provided. The modified software allowed all images from the sensor to be saved under certain conditions, as explained below.

These iris images are 640 (width) by 480 (height) pixels in resolution. For most “good” iris images, the diameter of the iris in the image exceeds 200 pixels. The images are stored with 8 bits of intensity, but every third intensity level is unused. This is the result of a contrast stretching automatically applied within the LG EOU 2200 system. The iris images were digitized from NTSC video, and the iris images may have interlace artifacts due to motion of the subjects.

In our acquisitions, subjects were seated in front of the system. The system provided recorded voice prompts to aid the subject in positioning their eye at the appropriate distance from the sensor. The system took images in “shots” of three. Each image was illuminated by one of three near-infrared light-emitting diodes (LEDs).

For a given subject at a given iris acquisition session, two “shots” of three images were taken for each eye, for a total of 12 images; see Fig. 5.1 for an example set of images from an acquisition session. Figure 5.2 shows examples of “lower” quality iris images. The system provided a feedback sound when an acceptable shot of images was taken. An acceptable shot had one or more images that passed the LG EOU 2200s built-in quality checks, but all three images were saved. If none of the three images passed the built-in quality checks, then none of the three images were saved. At least one-third of the iris images did pass the Iridian quality control checks, and up to two-thirds did not pass. A manual quality control step was performed at Notre Dame to remove images in which, for example, the eye was not visible at all due to the subject having turned their head.




Fig. 5.1 Example of iris images from an acquisition session

5.2.2 Protocol

The ICE 2006 evaluated algorithms. Participants had to deliver executable versions of their algorithms to NIST. The executables had to be received by NIST by June 15, 2006. The ICE 2006 was open to academia, industry, and research laboratories. Participants could submit multiple algorithms.

The format for submissions was binary executables that could be run independently on the test server. All submitted executables had to run using a specified set of command line arguments. The command line arguments included an experiment parameter file, files that contained the sets of biometric samples to be matched, and name of the output similarity file.

The ICE 2006 had an optional iris image quality task. For the quality task, executables gave a quality score for each iris image. The quality score had to be



Fig. 5.2 Examples of “lower” quality iris images in the ICE 2006 dataset

Table 5.1 The list of recognition and quality algorithms covered in this chapter. Column headings identify each participant group. The organization that submitted an algorithm is listed in the group column. The abbreviations used in this chapter are presented in the iris recognition algorithm and iris quality algorithm columns

Group	Iris recognition algorithm	Iris quality algorithm
U. of Cambridge	CAM-2	CAM-2-Q
Iritech	IRTCH-2	IRTCH-2-Q
Sagem-Iridian	SI-2	SI-2-Q

an integer in the range between 0 and 100, with 100 being the highest quality. A quality score is a number that rates an image’s utility to a recognition system and should be predictive of performance [4].

The test system hardware for the ICE 2006 was a Dell PowerEdge 850 server with a single Intel Pentium 4 3.6 GHz 660 processor and 2 MB of 800 MHz cache. All systems had 4 GB of 533 MHz DDR2 RAM. At no time did the test system have access to the Internet. The ICE 2006 allowed executables that would run under Windows Server 2003 (standard edition) and Linux Fedora Core 3 operating systems.

Table 5.1 lists the ICE 2006 algorithms whose results are presented in this chapter. These are the same algorithms reported in Phillips et al. [8]. All three of these submissions included a quality measure. All three matching algorithms and their quality algorithms were fully automatic.

5.3 Overview of the ICE 2006 Algorithm Performance Results

In this section, we summarize the ICE 2006 results that were published in Phillips et al. [8].

The images for the ICE 2006 were collected over three academic semesters: Spring 2004, Fall 2004, and Spring 2005. In computing performance, all similarity scores are cross semesters; that is, iris images taken in the same semester were not compared. There were 30 partitions for the left eye and 30 partitions for the right eye. For each algorithm, the false reject rate (FRR) at a false accept rate (FAR) of 0.001 was computed independently for each partition. The performance for each eye at a FAR of 0.001 was characterized by 30 FRRs which were summarized by a boxplot.

Performance for the ICE 2006 benchmark is presented in Fig. 5.3 for algorithms from three groups: Sagem-Iridian (SI-2), Iritech (IRTCH-2) and Cambridge (CAM-2). The interquartile range for all three algorithms overlaps, with the largest amount of overlap between Iritech (IRTCH-2) and Cambridge (CAM-2).




Fig. 5.3 Summary of the performance of the ICE 2006. Results are presented for three groups: Cambridge (Cam-2), Iritech (Irtch-2) and Sagem-Iridian (SI-2). Performance is broken out by right and left eyes. The false reject rate (FRR) at a false accept rate (FAR) of 0.001 is reported. Performance is reported for 29,056 right and 30,502 left iris images from 240 subjects with 30 partitions for each eye (This figure is from Phillips et al. [8])

5.4 Quality Scores for ICE 2006

The ICE 2006 protocol allowed for algorithms to report a quality measure for each iris image. Because the input to the quality module was an iris image, by their nature, the quality metrics are image-based. Quality scores were required to be reported as an integer between 0 and 100, with 100 being the highest quality. Figure 5.4 shows the histograms of quality scores for Sagem-Iridian (SI-2-Q), Iritech (IRTCH-2-Q), and Cambridge (CAM-2-Q). The quality score distribution for Sagem-Iridian is bimodal with modes at 0 and 100.

5.4.1 Quality Score Correlation Study

Intuitively, one would expect that quality scores measure generic iris image properties, such as amount of occlusion, in-focus measure, or amount of motion blur. While high correlation between quality measures might not be expected, some correlation is to be expected. This is especially true if a quality metric is intended to generalize to multiple algorithm approaches. Our first experiment explored the correlation among the quality scores returned by Sagem-Iridian (SI-2-Q), Iritech (IRTCH-2-Q), and Cambridge (CAM-2-Q). Scatterplots for comparing the quality measures are given in Fig. 5.5. Each point in a scatterplot is an iris image with coordinates being the quality score for two participants. In each scatterplot, there are 59,558 points, one for each iris image in the ICE 2006 dataset.

The corresponding Pearson and Spearman correlation coefficients calculated for pairs of participant quality scores are given in Table 5.2. These coefficients reflect little correlation between quality scores for the different participants.

5.4.2 Correlation Between Right and Left Irises

One of the open questions in biometrics is whether or not there exist correlations between the left and right irises of a person. Analysis performed on the ICE 2005 results showed that for submitted quality scores, there was a correlation between the left and right irises of a person [9]. This analysis is repeated for the quality scores for the ICE 2006.

The first step computes the average quality score for each person for both the left and right iris images. The average quality score for all the left irises for person i is denoted by $\rho_L(i)$, and the corresponding average for the right iris images is $\rho_R(i)$. Figure 5.6 shows scatterplots for $\rho_L(i)$ versus $\rho_R(i)$ for each of the three quality measures Sagem-Iridian (SI-2-Q), Iritech (IRTCH-2-Q), and Cambridge (CAM-2-Q). Each point in the scatterplot is a subject, and the coordinates of the point are $(\rho_L(i), \rho_R(i))$. Quality scores from left and right irises from 240 subjects were

Fig. 5.4 Histogram of the quality scores for (a) Sagem-Iridian (SI-2-Q), (b) Cambridge (CAM-2-Q), and (c) Iritech (IRTCH-2-Q)

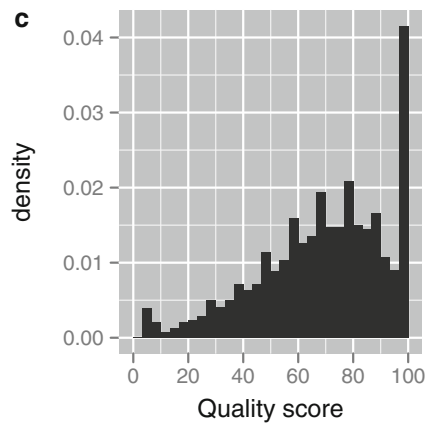




Fig. 5.5 Scatterplot and correlation between quality scores for the quality submissions to the ICE 2006. (a) Shows the correlation between CAM-2-Q and SI-2-Q, (b) shows the correlation between IRTCH-2-Q and SI-2-Q, and (c) shows the correlation between IRTCH-2-Q and CAM-2-Q




Table 5.2 Correlation between quality scores. Correlation coefficients are given for Pearson’s r and Spearman’s ρ

Quality scores	Pearson’s r	Spearman’s ρ
SI-2-Q vs. CAM-2-Q	0.122	0.131
SI-2-Q vs. IRTCH-2-Q	0.349	0.348
CAM-2-Q vs. IRTCH-2-Q	0.120	0.108

Table 5.3 Correlation coefficients between mean quality scores for the left and right irises. Correlation coefficients are given for Pearson’s r and Spearman’s ρ

Quality score	Pearson’s r	Spearman’s ρ
SI-2	0.362	0.347
CAM-2	0.896	0.896
IRTCH-2	0.790	0.781

included in this analysis. Included in Fig. 5.6 is the regression line for $\rho_L(i)$ versus $\rho_R(i)$. Table 5.3 lists the Pearson’s r and Spearman’s ρ correlation coefficients for mean subject quality scores for the left and right irises, $\rho_L(i)$ versus $\rho_R(i)$.

The results show that the average left and right quality measure are highly correlated for Cambridge and Iritech. The lower amount of correlation between the left and right irises for Sagem-Iridian SI-2-Q module may be due to the bimodal nature of the module’s quality score distribution, as shown in Fig. 5.4.

5.4.3 Demographics and Quality Scores

In face recognition, the sex, race, age, and other subject covariates effect performance [2, 7]. It has been suggested that eye color and race effect the performance of iris recognition systems.

The two largest demographic groups in the ICE 2006 data-set are Caucasians and East Asians. The eye color for the East Asians was uniformly labeled dark. However, for the Caucasian, subject had a full range of eye colors including green, blue, hazel, and brown. The Caucasian eye colors were grouped into light and dark to ensure that there were sufficient number of subjects in each category. The light eyes consisted of green and blue eyes, and the dark colors consisted of hazel and brown eyes. Since race and eye color are not independent factors, for our analysis on the effect of race and eye color, we combine these factors into one factor with three levels: Caucasian with light irises, Caucasian with dark irises, and East Asians.

Since the quality scores from the three quality algorithms have different distributions, for the analysis in this section, the quality scores were normalized. The first step in the normalization process was to compute the empirical cumulative distribution function \hat{F} . The normalized quality scores are then computed by

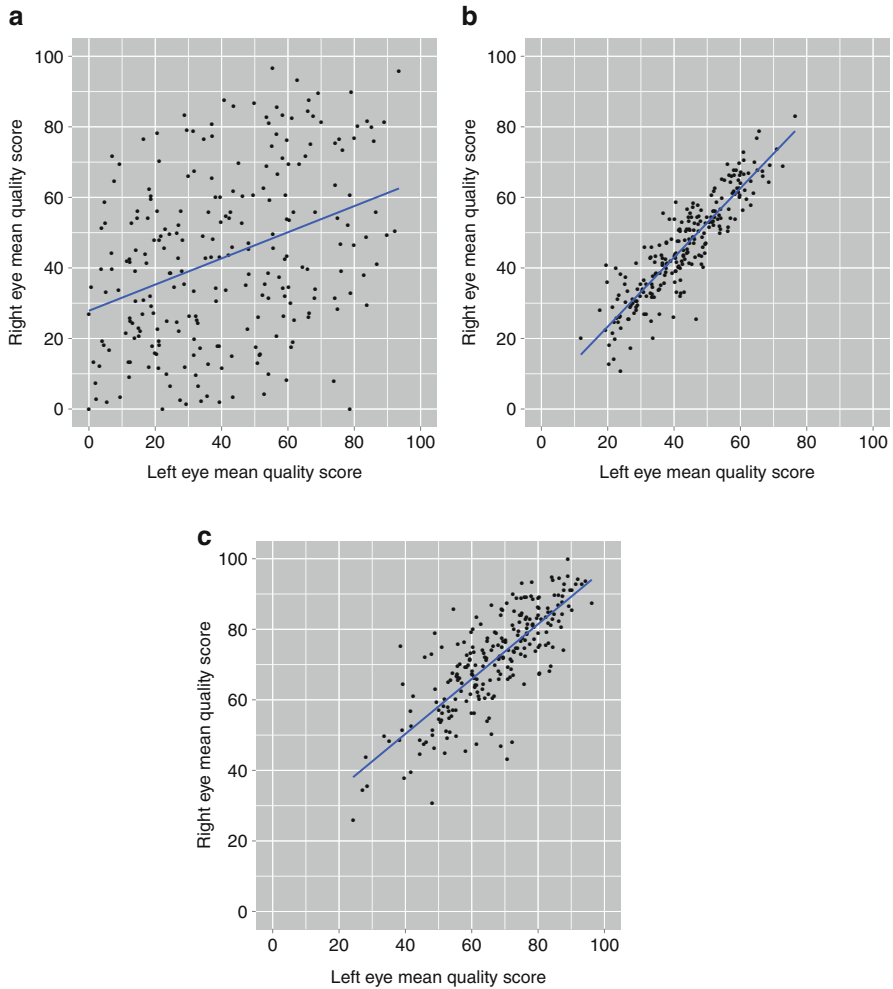



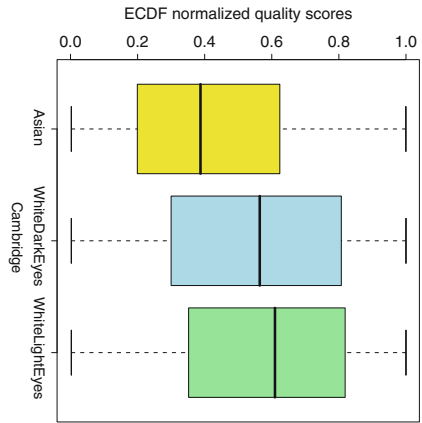

Fig. 5.6 Scatterplot and correlation between subjects’ mean quality score for the left and right eyes. The scatterplots are for quality modules (a) SI-2-Q, (b) CAM-2-Q, and (c) IRTCH-2-Q

$q' = \hat{F}^{-1}(q)$ for the original quality score q . The normalized quality scores q' are between 0.0 and 1.0 and have an approximately uniform distribution. (Note: Because the distribution of the Sagem-Iridian quality scores has peaks at 0 and 100, the normalized quality scores do not have a uniform distribution.)

Figure 5.7 shows boxplots of the normalized quality for broken out by East Asian, Caucasian with dark irises, and Caucasian with light irises. Results are reported for Sagem-Iridian (SI-2-Q), Iritech (IRTCH-2-Q), and Cambridge (CAM-2-Q).

The plots in Fig. 5.7 clearly show that the quality measures are sensitive to race and eye color. For Sagem-Iridian and Iritech, the median quality scores are ordered

Fig. 5.7 Boxplots of quality score broken out by East Asian, Caucasians with dark iris color, and Caucasians with light iris color. Results are reported for (a) the Sagem-Iridian (SI-2-Q), (b) Iritech (IRTCH-2-Q), and (c) Cambridge (CAM-2-Q) quality modules



by East Asian, Caucasians with dark colored irises, and then Caucasians with light colored irises. For the Cambridge quality module, the ordering of the median of the quality scores is the opposite. Figure 5.7 shows that the quality modules respond qualitatively different to race and eye color.

We looked at the effect of the sex of the subject on the distribution of quality scores. For the ICE 2006 submissions, the results of our analysis showed that the sex of the subject did not affect the distribution of the quality scores.

5.5 Quality Scores and Performance

Quality scores are computed independently for each iris image. The underlying assumption is that quality scores are predictive of performance. If two iris images have “high-”quality scores, then when they are compared, the matching algorithm is likely to give the correct answer. The obvious use of quality scores is to exclude “low-”quality iris images from the matching step. This is a sensible operational step for enrollment (population of the gallery set), since additional images are relatively easy to acquire from a cooperative subject at enrollment. Some scenarios may also allow multiple images to be acquired on-line, then selecting the highest quality image for enrollment. Our analysis models the above scenario where the quality controls are at the points of enrollment and acquisition. In this model, when two iris images are compared, both iris images independently meet the quality standard for the system. In the analysis in this section, we systematically vary the quality threshold and assess its impact on performance.

One of the goals in quality research is the development of quality measures that are effective for multiple algorithms. This property is sometimes referred to as universality of quality metrics. To investigate universality, we explore the effectiveness of each quality measures against each of the three matching algorithms.

5.5.1 *Measuring Performance of Quality Measures*

It is well established in biometrics that performance of an algorithm needs to report both type I and type II errors, with FRR and FAR as the basis for our analysis. In this section, we extend the traditional FRR and FAR trade-off to include quality measures, which follows [4]. Computation of FRR and FAR is based on similarity scores that results from comparing two iris images. The first step takes quality scores for an iris image and creates a quality score for a pair of iris images. The second step plots the trade-off in performance among FRR, FAR, and quality scores.

The first step is to compute image quality scores for each image in both the target and query sets. The quality for an iris image q is denoted by $r(q)$.

The next step is to define a quality measure for a pair of iris images q and t . Computing a quality measure $r(t, q)$ for a pair of iris images makes it possible

to associate each similarity score $s(i, j)$ with a quality score. For the analysis in this chapter, the quality measure $r(t, q) = \min(r(t), r(q))$. The selection of the min function models the situation where the target and query images are acquired separately, and both images must have a minimum quality score to be saved. In computing $r(q, t)$, other possible functions include average and weighted average. The selection of the appropriate function depends on the scenario being modeled.

The last step is to extend the characterization of verification performance to include quality scores. Whereas the traditional characterization of verification performance is the traded-off between FRR and FAR, the extended version shows the trade-off among FRR, FAR, and quality rate (QR). For the quality performance analysis in this chapter, we need a decision threshold and the quality range. In the quality analysis, a decision threshold λ serves the same purpose as in traditional signal detection analysis. Image quality is characterized by a quality rate which is defined as

$$QR(\tau_l, \tau_u) = \frac{|\tau_l \leq r(t), r(q) \leq \tau_u|}{|\mathcal{T}| + |\mathcal{Q}|}.$$

The QR gives the fraction of iris images that fall within the quality range. For each image quality threshold τ , FRR and FAR are computed from the similarity scores that have a quality score above τ . Formally $FRR(\tau_l, \tau_u, \lambda)$ and $FAR(\tau_l, \tau_u, \lambda)$ are

$$FRR(\tau_l, \tau_u, \lambda) = \frac{|s(t, q) \leq \lambda, \text{ where } (t, q) \text{ is a match and } \tau_l \leq r(t), r(q) \leq \tau_u|}{|(t, q) \text{ that are matches and } \tau_l \leq r(t), r(q) \leq \tau_u|},$$

and

$$FAR(\tau_l, \tau_u, \lambda) = \frac{|s(t, q) \geq \lambda, \text{ where } (t, q) \text{ is a nonmatch and } \tau_l \leq r(t), r(q) \leq \tau_u|}{|(t, q) \text{ that are nonmatches and } \tau_l \leq r(t), r(q) \leq \tau_u|}.$$

For each experiment in this chapter, the same base decision threshold λ_0 was selected for each algorithm. The decision threshold λ_0 corresponds to a FAR of 0.001 computed the entire ICE 2006 experiment as described in Sect. 5.3. Given this threshold λ_0 , the quality performance statistics $QR(\tau_l, \tau_u)$, $FRR(\tau_l, \tau_u, \lambda_0)$, and $FAR(\tau_l, \tau_u, \lambda_0)$ are computed for different values of the quality thresholds τ_l and τ_u . For the analysis in this chapter, the selection criteria fall into two categories. In the first, the upper threshold τ_u is set to the maximum quality score ($\tau_u = 100$), and τ_l is varied. This selection criteria computes performance when the quality scores are above a threshold set by τ_l . This method measures the trade-off among QR, FRR, and FAR as the quality of iris images increases. In the second category, the lower threshold τ_l is set to the minimum quality score ($\tau_l = 0$), and τ_u is varied. This selection criteria computes performance when the quality scores are below the threshold set by τ_u . This method measures the trade-off among QR, FRR, and FAR on lower quality images.

5.5.2 *Effect of Quality Scores on Performance*

The underlying premise for quality measures is that they are predictive of performance. In addition, a desirable property of quality measures is that they are universal. A quality measure is universal if it is predictive of performance for a class of algorithms.

In this experiment, we test the predictive property of the quality measures in the ICE 2006 and their universality. To test predictive property of the ICE 2006 quality measures, performance is computed as the minimum threshold of quality scores is increased. For this analysis, the upper threshold τ_u is set to the maximum value of 100. Then QR, FRR, and FAR are computed for 20 increasing values of the lower threshold. The values for τ_l are 0, 5, 10, ..., 95. The case when $\tau_l = 0$ corresponds to no pruning by quality. The lower thresholds of 5, 10, ..., 95 correspond to requiring increasing levels of quality of the target and query images.

The analysis in this chapter presents results for three matching algorithms. Each of these groups also submitted a quality algorithm. To test universality, we measured the predictive power of each quality measure against all three matching algorithms.

To show the trade-off among QR, FRR, and FAR, we plot QR versus FRR and QR versus FAR. In plots of QR versus FRR, the horizontal axis is $QR(\tau_l, 100)$, and the vertical axis is $FRR(\lambda_0, \tau_l, 100)$. When $\tau_l = 0$, $QR(\tau_l, 100) = 1$ and $FRR(\lambda_0, \tau_l, 100) = FRR(\lambda_0)$, which is the FRR at a FAR of 0.001. In plots of QR versus FAR, the horizontal axis is $QR(\tau_l, 100)$, and the vertical axis is $FAR(\lambda_0, \tau_l, 100)$. When $\tau_l = 0$, $FAR(\lambda_0, \tau_l, 100) = 0.001$.

If this premise is true, then as lower quality biometric samples are removed, performance should improve. It is desirable that quality scores are predictive of performance for multiple algorithms.

Figures 5.8–5.10 present the performance triplet values by matching algorithm, by triplet element, by quality scorer, and by quality pruning threshold. Figure 5.8 contains results for matcher SI-2, Fig. 5.9 contains results for matcher IRTCH-2, and Fig. 5.10 contains results for matcher CAM-2. In each figure, the top plot contains FRR curves for each of the three quality scorers SI-2-Q, IRTCH-2-Q, and CAM-2-Q, (plotted in different colors), and the bottom plot shows FAR curves. The points in each plot indicate the components of the operating point triplet for quality pruning thresholds proceeding from 0.0 at the left edge to 1.0 at the right edge. The horizontal axis measures the percentage of iris pairs discarded by the quality pruning threshold for each curve.

In Fig. 5.8, the red curve in the bottom plot shows the values of FAR for the SI-2 matcher as iris samples are discarded using quality module CAM-2-Q. With no quality-based pruning at the left, the FAR = 0.001. As more images are pruned by the quality module CAM-2-Q, the FAR increase under the fraction pruned is about 0.9. For the other two quality modules, FAR decreases with quality-based pruning.

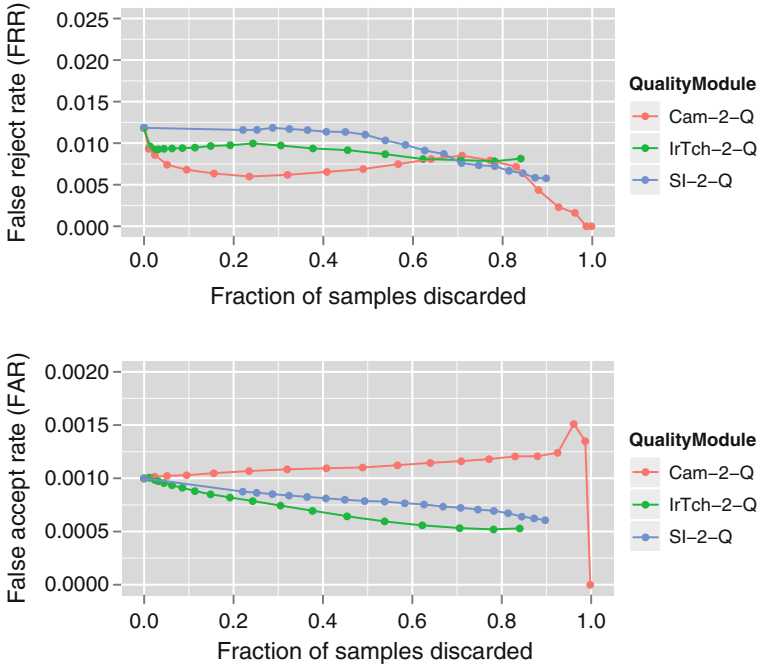


Fig. 5.8 Effect of all three quality measures on matcher SI-2

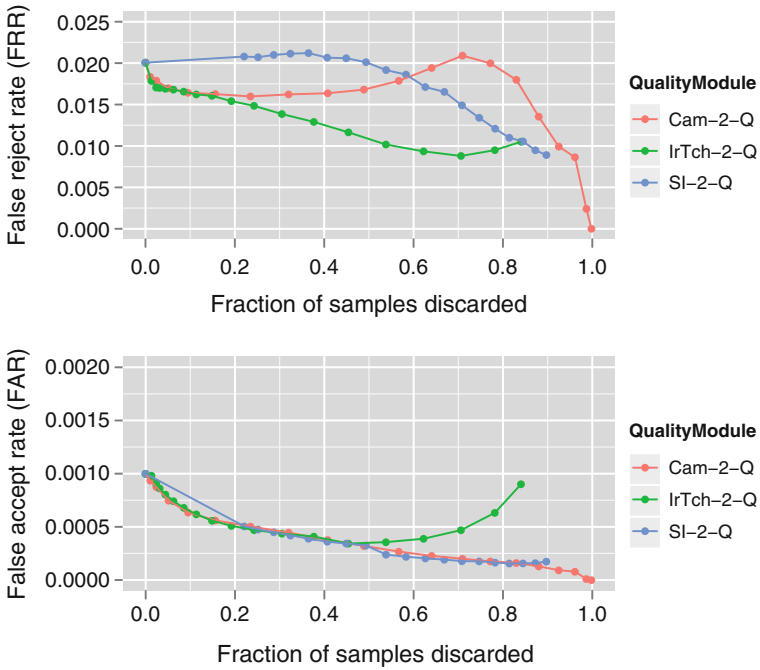


Fig. 5.9 Effect of all three quality measures on matcher IRTCH-2

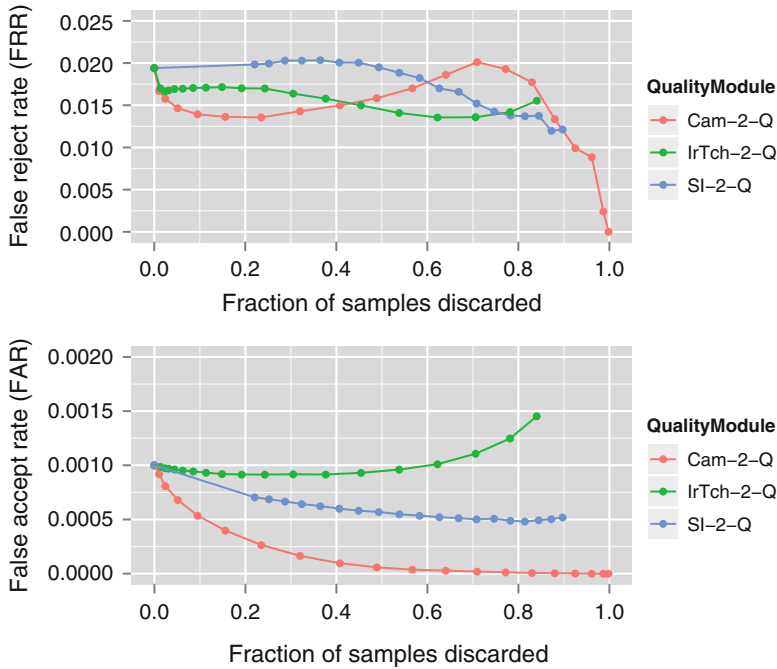


Fig. 5.10 Effect of all three quality measures on matcher CAM-2

5.6 Discussions and Conclusions

The core experiments in this analysis looked at the universality of the ICE 2006 quality measures. The first study looked at the correlation among the three quality measures. Spearman’s correlation coefficient ranged from 0.108 to 0.348. This showed minimal correlation among the quality measures.

One study looked at the effect of eye color and race on the quality measures. For the SI-2-Q and IRTECH-2-Q measures, the median quality measures were ordered Asian, Caucasian with light colored eyes, and Caucasian with dark colored eyes (from highest to lowest). For the CAM-2-Q measure, the median quality measures had the reverse order (Caucasian with dark colored eyes, Caucasian with light colored eyes, and Asian). This showed that the quality measures were sensitive to the combination of eye color and race. In addition, the sensitivity varied by quality measures.

In the final experiment, we looked at the ability of the quality measures to predict performance. The results showed that none of the three quality measures could universally predict performance.

These three core experiments suggest that the three quality modules are measuring different properties of an iris. Also, the three matchers respond differently to these properties.

Since the ICE 2006, conducted, research in quality measures for iris images has continued. The IREX II Iris Quality Calibration and Evaluation (IQCE) will measure progress in the development of iris quality algorithms.¹

Acknowledgements We acknowledge the support of Department of Homeland Security's Science and Technology Department and Transportation Security Administration (TSA), the Director of National Intelligence's Information Technology Innovation Center, the Federal Bureau of Investigation (FBI), the National Institute of Justice, and the Technical Support Working Group (TSWG). This work was undertaken during PJF's sabbatical leave at NIST. The support of NIST is gratefully acknowledged. Biometrics research at the University of Notre Dame is supported by the National Science Foundation under grant CNS01-30839, by the US Department of Justice, and by Unisys Corp. The identification of any commercial product or trade name does not imply endorsement or recommendation by the National Institute of Standards and Technology or the University of Notre Dame.

References

1. Authenti-Corp: Iris recognition study 2006 (IRIS06). Tech. Rep. version 0.40, Authenti-Corp (2007). <http://www.authenti-corp.com/iris06/report/>
2. Beveridge, J.R., Givens, G.H., Phillips, P.J., Draper, B.A.: Factors that influence algorithm performance in the face recognition grand challenge. *Comput. Vis. Image Underst.* **113**, 750–762 (2009)
3. Grother, P.J., Tabassi, E., Quinn, G.W., Salamon, W.: IREX I performance of iris recognition algorithms on standard images. Tech. Rep. NISTIR 7629, National Institute of Standards and Technology (2009)
4. Grother, P., Tabassi, E.: Performance of biometric quality measures. *IEEE Trans. PAMI* **29**, 531–543 (2007)
5. International Biometric Group: Independent testing of iris recognition technology. Tech. rep., International Biometric Group (2005). <http://www.ibgweb.com/reports/public/ITIRT.html>
6. Newton, E.M., Phillips, P.J.: Meta-analysis of third-party evaluations of iris recognition. *IEEE Trans. SMC-A* **39**(1), 4–11 (2009)
7. Phillips, P.J., Grother, P.J., Micheals, R.J., Blackburn, D.M., Tabassi, E., Bone, J.M.: Face recognition vendor test 2002: evaluation report. Tech. Rep. NISTIR 6965, National Institute of Standards and Technology (2003). <http://www.frvt.org>
8. Phillips, P.J., Scruggs, W.T., O'Toole, A.J., Flynn, P.J., Bowyer, K.W., Schott, C.L., Sharpe, M.: FRVT 2006 and ICE 2006 large-scale results. *IEEE Trans. PAMI* **32**(5), 831–846 (2010)
9. Phillips, P., Bowyer, K.W., Flynn, P.J., Liu, X., Scruggs, W.T.: The Iris challenge evaluation 2005. In: *Second IEEE International Conference on Biometrics: Theory, Applications, and Systems*, Washington, DC (2008)

¹For details on IREX II IQCE, see <http://iris.nist.gov>.

Chapter 6

Iris Recognition with Taylor Expansion Features

Algirdas Bastys, Justas Kranauskas, and Volker Krüger

Abstract The random distribution of features in an iris image texture allows to perform iris-based personal authentication with high confidence. In this chapter, we describe three iris representations. The first one is a phase-based iris texture representation which is based on a binarised multi-scale Taylor expansion. The second one describes the iris by using the most significant local extremum points of the first two Taylor expansion coefficients. The third method is a combination of the first two representations. For all methods, we provide efficient similarity measures which are robust to moderate iris segmentation inaccuracies. Using three public iris datasets, we show (a) the compact template size of the first two representations and (b) their effectiveness: the first two representations alone already perform well, but in combination, they outperform significantly state-of-art iris recognition approaches.

6.1 Introduction

Iris images have a complex texture of visible particles of different sizes. The particles come from freckles, furrows, stripes, coronas, occluding eyelashes, etc. The distribution of these particles on the iris in terms of their locations, shapes, and sizes is random and differs from eye to eye which is the reason why iris images are such an effective biometric for verification and identification.

Generally, there are three different types of approaches to capture the individual distribution of particles on the iris: phase-based techniques [7–10, 20], techniques based on zero-crossings [4, 21, 27, 28], and more intuitive iris texture descriptors

A. Bastys (✉) • J. Kranauskas
Vilnius University, Vilnius, Lithuania
e-mail: algirdas.bastys@mif.vu.lt

V. Krüger
Aalborg University, Aalborg, Denmark
e-mail: vok@cvmi.aau.dk

[15, 23, 29–32]. Despite seemingly differences in phase, zero-crossing, and texture-based iris descriptors, the approaches are all closely related. For example, all three iris descriptions preserve information about the original image texture with a large likelihood.

In this chapter, we present results of using a multi-scale Taylor expansion of the iris pattern [2] from which we are computing local extrema and phase information.

Phase-based iris recognition algorithms are the most popular ones and show good speed and verification performances (see [7–10, 20]). Our phase-based iris representation uses features that are computed from a binarised multi-scale Taylor expansion. Our phase information is an iris code in which 0 codes a negative and 1 codes a positive sign of the expansions [2].

Local extrema of wavelet, Gabor, or other expansions for iris recognition were used already earlier [1, 16, 17]. But, e.g. unlike the 1D local extrema as used in [17], we use 2D local extrema points of the multi-scale expansion. Such an approach reduces the size of the iris template which is important for efficient processing in large-scale applications, and it eliminates correlations of the local extrema points along the radial direction which improves reliability. Furthermore, the first two coefficients of the multi-scale Taylor expansion have a transparent interpretation, and they can be computed very efficiently. Our local extrema are computed from approximations of the first- and second-order Taylor coefficients of the multi-scale Taylor expansion [2, 3], averaged over different scales. The most significant local extrema of the Taylor expansions define the most significant iris texture variations.

As we will show in our experiments, the 2D local extrema and the phase-based approaches alone already reach very good results, but, interestingly, their combination reaches one of the highest recognition rate on the databases Casia 2.0, ICE-1, and MBGC-31.


For all methods, we employ efficient and elastic iris template matching schemes [13, 26] that allow to compensate for slight segmentation and localisation errors of the iris features and that assure efficient processing.

This chapter is structured as follows: We will briefly discuss the automatic segmentation used for preprocessing of the database images in Sect. 6.2. In Sects. 6.3, 6.3.2 and 6.3.3, we discuss the binary and the local extrema-based iris representations. In Sect. 6.3.4, we discuss the elastic matching scheme for the phase-based approach and the efficient matching scheme for the local extrema-based approach. The joint distance measure is a combination of the phase-based and the local extrema-based distance measures, and it is discussed in Sect. 6.4. The thorough evaluation for all three approaches is provided in Sect. 6.4. We conclude with final remarks in Sect. 6.5.

6.2 Segmentation

The iris segmentation process consists of mainly three steps [2, 3]: (1) localising the inner and outer boundaries of the iris, (2) the optional detection of occluding

Fig. 6.1 Approximation of the *inner* and *outer* iris boundaries by convex contours (The iris image is taken from the Casia V3 interval database. Figure reproduced with permission from [2])



upper and lower eyelids, and (3) the detection and removal of reflections on the cornea or the eyeglasses. For iris segmentation, we use standard state-of-the-art image processing techniques (see also [11] for a general overview on eye detection).

Figure 6.1 illustrates a typical segmentation result, the contours are approximated by short Fourier series.

6.3 Iris Features from Multi-scale Taylor Expansion

In this section, we discuss the two Taylor-based approaches in detail. For both approaches, we use the first two coefficients of the Taylor expansion that models the iris at different radii and scales. We compute the binary features by thresholding and binarising the first two Taylor coefficients. For the local extrema approach, we use the local extrema points of the first two Taylor coefficients on different scales.

In the following, we give a brief overview of the Taylor expansion as far as we need it for understanding the later discussions. In Sect. 6.3.2, we discuss the binary features, and Sect. 6.3.3 goes into the details of the local extrema-based features.

6.3.1 Brief Introduction to the Taylor Expansion

Consider the one-dimensional iris pattern at a specific pseudo-polar radius $r = \text{const}$ as an analytic one-dimensional signal $u = u(x)$ where x then denotes the pseudo-polar angle. Then, we are able to express the iris pattern using a classic Taylor series expansion centred on any fixed point $x = x_i$. The zero coefficient of the expansion, given by $u(x_i)$, is strongly illumination dependent and will not use it. The first- and second-order coefficients are given as $\frac{u'(x_i)}{1!}$ and $\frac{u''(x_i)}{2!}$. The first derivative $u'(x)$ measures the transition rate of the signal around point x , i.e. a big

positive value of $u'(x)$ reflects a fast increasing transition of the signal around point x and a big negative $u'(x)$ reflects a fast decay around the same point. The sign of the second derivative $u''(x)$ provides information about the type of concavity of a graph $\{x, u(x)\}$ around x , i.e. if $u''(x) > 0$, then $\{x, u(x)\}$ is concave upward and if $u''(x) < 0$, then the graph $\{x, u(x)\}$ is concave downward around point x .

At a local negative minimum or positive maximum of the first derivative, the second derivative will have a zero-crossing. Note that the derivatives $u'(x)$ and $u''(x)$ can be estimated only numerically, and therefore, results of the approximations significantly depend on the scale at which variations of the signal are estimated. As pointed out above, we neglect the 0-th order Taylor coefficients of $u(x_i)$ because of their illumination dependency. The first two terms of the Taylor series expansion define the local linear and the quadratic behaviour of u around the point of the expansion. Extrema of the first derivative indicate points where the signal u has its greatest local asymmetry. Similarly, extrema of second derivative indicate points of great local symmetry of u . It is interesting to note that the real part of the familiar Gabor response at a point x is similar to the second-order Taylor coefficient centred at the point x , and imaginary part the Gabor response is alike to the first-order Taylor coefficient.

Let us give a formal definition. The well-known Taylor series expansion

$$u(x) = \sum_{n=0}^{\infty} \frac{u^{(n)}(a)}{n!} (x-a)^n \quad (6.1)$$

assures us that a sufficiently regular function u can be reconstructed from its Taylor coefficients $\frac{u^{(n)}(a)}{n!}$. In theory, Taylor coefficients can be derived from the $u(x)$ values known at any narrow surrounding of the centre point $x = a$. Let us consider some details of the estimation at different scales of the first two derivatives of the iris texture. The estimation technique of the derivatives at different scales is similar to the wavelet multi-resolution analysis. Let us consider some even smooth density $p = p(x) \geq 0$ that has its maximum at $x = 0$ and a normalised L_1 energy

$$\int_{-\infty}^{\infty} p(x) dx = 1. \quad (6.2)$$

Let us fix some scale $\sigma > 0$ and consider the convolutions

$$(u' * D_{\sigma} p)(x) = 1/\sigma \int_{-\infty}^{\infty} p(\xi/\sigma) u'(\xi - x) d\xi \quad (6.3)$$

and

$$(u'' * D_{\sigma} p)(x) = 1/\sigma \int_{-\infty}^{\infty} p(\xi/\sigma) u''(\xi - x) d\xi, \quad (6.4)$$

where D_{σ} is the dilation operator for a given scale σ . We can observe that for a larger-scale σ the blurring of the signals $u'(x)$ and $u''(x)$ increases. If σ vanishes, then $(u' * D_{\sigma} p)(x)$ and $(u'' * D_{\sigma} p)(x)$ converge to $u'(x)$ and $u''(x)$, respectively.

The density $p = p(x)$ defines the rule for averaging the derivatives. A function meeting the requirements of p is $\text{sech}(x) = 2/(e^x + e^{-x})$: It has the required symmetry, tends exponentially to zero at infinity and its L_1 norm equals to π . Therefore, we use the following density function:

$$p(x) = \frac{\text{sech}(x)}{\pi}. \quad (6.5)$$

The iris texture along the angular direction can be described as a discrete periodic sequence $u_i = u(i)$ of some length NX . Defining $u(x)$ through linear interpolation $u(x) = u_i + (x - i)(u_{i+1} - u_i)$, $x \in [i, i + 1)$, we have

$$u''(x) = \sum_{i=-\infty}^{\infty} (u_{i+1} - 2u_i + u_{i-1})\delta(x - i) \quad (6.6)$$

and

$$\begin{aligned} (u'' * D_{\sigma}p)(x) &= \frac{1}{\sigma} \sum_{i=-\infty}^{\infty} (u_{i+1} - 2u_i + u_{i-1})p\left(\frac{x+i}{\sigma}\right) \\ &= \sum_{i=0}^{NX-1} \Delta u_i p_{\sigma}^{NX}(x+i) \\ &= \sum_{i=0}^{NX-1} u_i \Delta p_{\sigma}^{NX}(x+i). \end{aligned} \quad (6.7)$$

Here, $\delta(x)$ is Dirac's delta function,

$$\Delta u_i = u(i+1) - 2u(i) + u(i-1), \quad (6.8)$$

and

$$p_{\sigma}^{NX}(x) = \frac{\sum_{i=-\infty}^{\infty} p\left(\frac{x+i*NX}{\sigma}\right)}{\sigma}. \quad (6.9)$$

Similarly, for the first derivative, we have

$$\begin{aligned} (u' * D_{\sigma}p)(j) &= - \sum_{i=0}^{NX-1} \Delta u_i q^{NX} \sigma(j+i), \\ q_{\sigma}^{NX}(x) &= \sum_{i=-\infty}^{\infty} 2 \arctan\left(\frac{x+i*NX}{\sigma}\right). \end{aligned} \quad (6.10)$$

Due to the symmetry of the density function (6.5) Eq. (6.7) defines a filtration of the 1D data $\{u_i\}$ by an even-symmetric filter and (6.10) defines an odd-symmetric filter. In wavelet terminology, the even and odd filters have two and one vanishing moments, respectively. The proposed filters have more freedom in choosing the

Table 6.1 Derivatives and averaging parameters used in derivation of multi-scale Taylor expansion and local features

Scale and filter type index s	Blurred derivative	$\sigma_x = \sigma$ of horizontal direction	$\sigma_y = \sigma$ of vertical direction
0	u_x	$\frac{6}{8}$	$\frac{1}{2}$
1	u_x	$\frac{9}{8}$	$\frac{1}{2}$
2	u_x	$\frac{15}{8}$	$\frac{1}{2}$
3	u_{xx}	$\frac{9}{8}$	$\frac{1}{2}$
4	u_{xx}	$\frac{15}{8}$	$\frac{1}{2}$
5	u_{xx}	$\frac{24}{8}$	$\frac{1}{2}$
6	u_y	$\frac{1}{2}$	$\frac{9}{8}$
7	u_{yy}	$\frac{1}{2}$	$\frac{15}{8}$

basic function $p^\sigma = p^\sigma(x)$ than the dyadic wavelet filtration, and they do not have dyadic restrictions for the scale σ . To regularise the filtration results that approximate second- (even filters) or first (odd filters)-order derivatives at different scales σ , we apply smoothing in the orthogonal direction by using the same density $p^\sigma = p^\sigma(y)$ but with a smaller scale value σ' .

Table 6.1 summarises the scales, type of symmetries, and directions that we have used to calculate the differences. The differences calculated along the x and y axis correspond to horizontal- and vertical-type filters. Derivatives in x direction were estimated at three different scales, while derivatives in y direction were estimated only at one scale. The number of derivatives along angular and radial directions was chosen experimentally. It seems that derivatives along radial (y) direction are less useful since sometimes they enhance the edges from undetected eyelids or pupil/iris boundary. These segmentation errors lead to increased inter-class similarity.

6.3.2 Binary Features as Local Iris Features

Following Daugman's traditions, we define the iris binary features simply by binarising the blurred first- and second-order Taylor coefficients u_x , u_{xx} , u_y , and u_{yy} . To be precise, consider the points (i, j) belonging to the iris region bounded by the two convex iris contours. If after blurring with scales σ_x and σ_y , we have $u_x(i, j) > 0$ at a point (i, j) then we set the binary feature $(u_x)_{i,j} = 1$. If we have for the same conditions $u_x(i, j) < 0$, then the binary feature is set to $(u_x)_{i,j} = 0$. Finally, if the point (i, j) does not belong to the iris region bounded by two convex contours or if $u_x(i, j) = 0$, then the binary feature $(u_x)_{i,j}$ is chosen randomly with equal probability.

Differently from binary features that are used to describe the phase of Gabor responses, our binary features have a transparent interpretation. For example, if u_x has the value $(u_x)_{i,j} = 1$ when blurred at scales σ_x and σ_y , it means that the iris




Fig. 6.2 240861 and 243395 iris images from the NIST ICE2005 dataset (*above*). Unwrapped 240861 and 243395 iris images (*below*) (Figure reproduced with permission from [2])

texture at scales σ_x and σ_y increases along angular direction around (i, j) point. If, on the other hand, $(u_{yy})_{i,j} = 0$, one can state that the iris texture is concave along radial direction around (i, j) point. A random definition of binary features occurs at points that correspond to regions of iris texture that is occluded by eyelids or that are located in areas with constant grey values. For example, in specular reflection regions, one can have $u(x, y) \equiv 255$ that leads to $u_x(x, y) \equiv 0$ and $u_y(x, y) \equiv 0$. The random definition of binary features at problematic iris regions reduces iris inter-class similarity variation and avoids iris similarity scores normalisation problem that arises when binary features corresponding to eyelids and eyelashes are ignored in similarity evaluation.

Since we are using eight different continuous two-dimensional multi-scale Taylor expansion coefficients (see Table 6.1), we are obtaining eight binary maps after binarisation. Figure 6.2, top, illustrates 240861.tiff and 243395.tiff iris images of NIST ICE2005 iris dataset and their unwrapped versions (Fig. 6.2, bottom). These two images constitute a genuine pair. Figure 6.3 shows the eight binary maps of the multi-scale Taylor features of the two unwrapped iris images from Fig. 6.2. The white and black pixels indicate locations of positive and negative averaged first- and second-order Taylor coefficients, respectively. Combining at any fixed location (i, j) eight bits of binary features, we form a byte-type feature

$$\mathbf{f}[i][j] = \sum_{k=0}^7 \{\text{feature bit of } k\text{th binary map at position } (i, j)\} 2^k. \quad (6.11)$$

Here, indexes i and j correspond to the unwrapped angular and radial positions. After some experimentation, we found that an iris template size of 256×32 gives the best compromise between speed and accuracy for the iris template matching. We will assume this iris template size ($256 \times 32 = 8,192$ (byte)) in the description of iris template matching technique in Sect. 6.3.4.

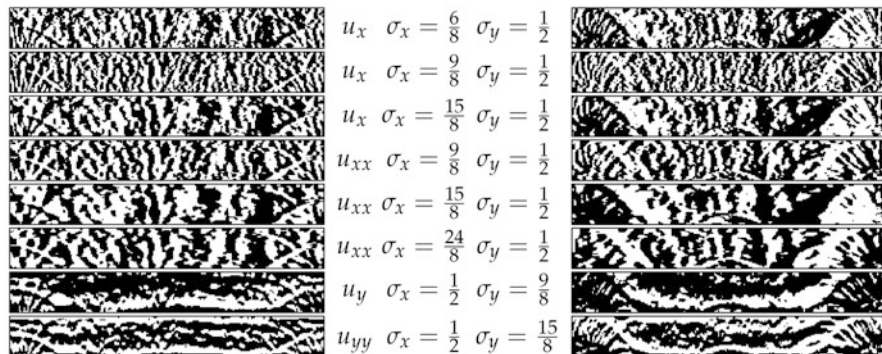



Fig. 6.3 Eight binary maps of multi-scale Taylor expansion of unwrapped 240861 and 243395 iris images. *White pixels* correspond to binary 1 and *black pixels* to binary 0 (Figure reproduced with permission from [2])

Fig. 6.4 The iris is divided into the different sectors 0...15 (Figure reproduced with permission from [2])



6.3.3 Local Extrema as Iris Features

It is the coefficients with largest squared values in the Taylor expansion and in the wavelet transform that give the main contribution for the approximation of the original signal. Thus, we use local extrema (positive local maxima and negative local minima) of blurred derivatives as local descriptors of the iris texture. To assure a better stability, an extremum is used as a feature only if (a) the local extrema at a particular scale σ_0 exceed the same derivative at the same position but with larger scales $\sigma > \sigma_0$ and (b) if its magnitude is large enough. However, to avoid thresholding the magnitude at this stage, we chose the best extrema points in the following way:

1. Subdivide the iris texture along the polar angle direction into 16 equal sectors.
2. Enumerate sectors clockwise according to Fig. 6.4.

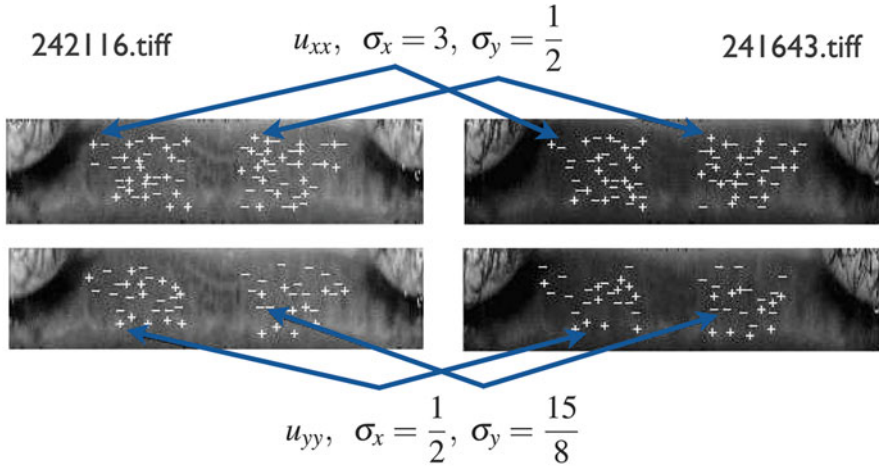


Fig. 6.5 Unwrapped 242116 and 241643 iris images from the NIST ICE2005 dataset. “+” and “-” mark local maxima and minima of the averaged u_{xx} and u_{yy} derivatives (u_{xx} – above, u_{yy} – below). Arrows points to the examples of extrema points that have correspondences in both images (Figure reproduced with permission from [2])

3. Ignore sectors 0, 1, 2, 7, 7, 13, 14, and 15 because the iris texture in these sectors is likely to be occluded by the upper and low lids.
4. Choose no more than K biggest local extrema from each scale and each of the remaining sector for some pre-defined K . For the reminder of the paper, we will refer to the remaining sectors as *active sectors*.

Figure 6.5 illustrates local extrema points that were found in the active sectors for the 242116.tiff and 241643.tiff iris images from the NIST ICE2005 iris dataset. These two images constitute genuine pairs. The two presented derivatives u_{xx} and u_{yy} are calculated using the averaging parameters defined by the sixth ($s = 5$) and eighth ($s = 7$) row of the Table 6.1. The “+” and “-” mark local maxima and minima points of the averaged Taylor expansion. The presented iris texture has no clearly expressed freckles, furrows, stripes, or coronas in the active sectors. However, there are blurred dark and bright blobs where local maxima and minima indicate the blobs of the size that resonates with the averaging scale σ .

6.3.4 Elastic Similarity Metric for Comparison of Binary Feature Maps

Let $\mathbf{f}^A[i][j]$ and $\mathbf{f}^B[i][j]$, $(i, j) \in \{256 \times 32\}$ be two sets of binary features formed from the multi-scale Taylor expansions of iris images A and B . The traditional Daugman approach of matching two binary iris templates consists of calculating the

Hamming distance for different angular shifts and the final distance is the minimum of all calculated distances. To be more consistent with our local extrema-based iris matching technique (see [2]), we will construct a similarity metric instead of a distance measure. A similarity measure of a pair of iris images can be defined simply by the maximal number of matched bits of their binary features maps where the maximisation is performed over a fixed number of angular shifts, i.e.

$$\text{similarity}(A, B) = \max_{\alpha} \sum_{i,j} \#_0(\mathbf{f}^A[i][j] \hat{\ } \mathbf{f}^B[i + \alpha][j]). \quad (6.12)$$

Here, $\hat{\ }$ denotes bit-wise XOR operation of 2 bytes and $\#_0(\mathbf{b})$ is number of 0 bits in byte \mathbf{b} . However, such a Daugman type similarity measure (6.12) underestimates the similarity of genuine pairs that have slight segmentation inaccuracies or nonlinearly deformed texture. To make the iris similarity measure more robust for such deformations and to have the possibility to control the eyelids and eyelashes influences, we again subdivided the unwrapped iris into sectors (as done above in Sect. 6.3.3) and Fig. 6.4.

Let $k \in \{3, 4, 5, 6, 9, 10, 11, 12\}$ be a sector number and α and β two shifts in the angular and radial direction. The similarity of two shifted k sectors is computed as

$$\text{score}_k(A, B; \alpha, \beta) = \sum_{i=k*16}^{(k+1)*16-1} \sum_{j=0}^{31} \#_0(\mathbf{f}^A[i][j] \hat{\ } \mathbf{f}^B[i + \alpha][j + \beta]). \quad (6.13)$$

Because of the angular periodicity of the iris texture and the size of 256×32 (see Sect. 6.3.2), we assume 256 periodicity of $\mathbf{f}^B[i + \alpha][j + \beta]$ in the first variable i . Similarly, in order to properly define $\mathbf{f}^B[i + \alpha][j + \beta]$ for at all $[j + \beta]$, we assumed 32 periodicity in the second variable, but this assumption does not have a natural interpretation. Maximising Eq. (6.13) over the angular shifts α compensates possible iris rotations; maximising over the radial shifts β allows to reduce possible influences of slight segmentation inaccuracies.

By allowing α_k and β_k to maximise the accumulative similarity scores of the two k sectors independently for different sectors k , we are able to compensate for nonlinear deformations of the iris texture. Unfortunately, such degree of freedom undesirably increases also impostor pairs similarity variation. To reach a compromise between the increase of genuine pair average similarity and a rise of variation of impostor pairs similarity, we define the elastic iris similarity measure as follows:

Definition 6.1. Let A and B be two iris images templates formed from binary maps of the Taylor multi-scale expansion features. Then, their elastic similarity is

$$\widetilde{\text{score}}^*(A, B) = \max_{\{\alpha_k, \beta_k\}_{k=3}^{12} \in T} \sum_{k=3}^{12} \text{score}_k(A, B; \alpha_k, \beta_k), \quad (6.14)$$

$$T = \left\{ \{ \alpha_k, \beta_k \}_{k=3}^{12} : |\alpha_k - \alpha_{k-1}| \leq 1, |\beta_k - \beta_{k-1}| \leq 1, k = 4, \dots, 12; \right. \\ \left. |\alpha_k| \leq KX, |\beta_k| \leq KY, k = 3, 4, \dots, 12 \right\}. \quad (6.15)$$

Here, KX and KY are integer parameters defining $2KX + 1$ rotations and $2KY + 1$ radial shifts under which the similarity scores of shifted and rotated iris templates are evaluated.

The concrete choice of KX and KY values can be adjusted to iris database and magnitude of segmentation inaccuracies. In our experimentations, the average used values are $KX = 10$ and $KY = 2$. The set T defines the admissible shifts $\{ \alpha_k, \beta_k \}_{k=3}^{12}$ under which the similarity scores are maximised. The $|\alpha_k| \leq KX$ and $|\beta_k| \leq KY$ conditions restrict the maximal admissible shifts. The $|\alpha_k - \alpha_{k-1}| \leq 1$ and $|\beta_k - \beta_{k-1}| \leq 1$ inequalities ensure the continuity of admissible deformations in the iris matching routine. The sum $\sum_{k=3}^{12} \text{score}_k(A, B; \alpha_k, \beta_k)$ includes the similarity of the sectors with indices $k = 7$ and $k = 8$ that we proposed to avoid because the iris texture in these sectors is likely to be occluded by lower lid. Therefore in the calculation of (6.14), we put $\text{score}_7(A, B; \alpha_7, \beta_7) = \text{score}_8(A, B; \alpha_8, \beta_8) = 0$. However, values of $\alpha_7, \beta_7, \alpha_8$, and β_8 influence the accumulative elastic similarity scores through the set T that defines all the admissible shifts.

Direct maximisation of the elastic similarity scores over the set T would require the evaluation of $4KXKY9^9$ accumulated similarity scores. In the following, we provide a dynamic programming approach for calculating (6.14) that reduces the maximisation complexity to $4KXKY9$ arithmetic operations. In the reminder of this section, we will skip, for better readability, the explicit reference to the iris images A and B in the notation. For example instead of $\text{score}_k(A, B; \alpha, \beta)$, we will write $\text{score}_k(\alpha, \beta)$.

Proposition 6.1. *Let the initial values of $(2KX + 1) \times (2KY + 1)$ of accumulative scores be set to*

$$\widetilde{\text{sum}}_3(\alpha, \beta) = \text{score}_3(\alpha, \beta), \quad |\alpha| \leq KX, |\beta| \leq KY. \quad (6.16)$$

For $k = 4, 5, \dots, 12$, the accumulative scores are calculated using the following dynamic programming rule:

$$\widetilde{\text{sum}}_k(\alpha, \beta) = \text{score}_k(\alpha, \beta) + \max_{\substack{i, j = -1, 0, 1: \\ |\alpha+i| \leq KX \\ |\beta+j| \leq KY}} \widetilde{\text{sum}}_{k-1}(\alpha+i, \beta+j). \quad (6.17)$$

Then, the maximum of $\widetilde{\text{sum}}_{12}$ gives the final elastic similarity scores (6.14), i.e.

$$\widetilde{\text{score}}^*(A, B) = \max_{\substack{|\alpha| \leq KX \\ |\beta| \leq KY}} \widetilde{\text{sum}}_{12}(\alpha, \beta). \quad (6.18)$$

Proof. Let

$$\max_{\substack{|\alpha| \leq KX \\ |\beta| \leq KY}} \widetilde{\text{sum}}_{12}(\alpha, \beta) = \widetilde{\text{sum}}_{12}(\alpha_{12}^*, \beta_{12}^*).$$

Then, using repeatedly (6.17), we have

$$\widetilde{\text{sum}}_k(\alpha_k^*, \beta_k^*) = \text{score}_k(\alpha_k^*, \beta_k^*) + \widetilde{\text{sum}}_{k-1}(\alpha_{k-1}^*, \beta_{k-1}^*), \quad k = 12, 11, \dots, 4,$$

where for all $k = 12, 11, \dots, 3$ the following inequalities hold true:

$$|\alpha_k^*| \leq KX, |\beta_k^*| \leq KY,$$

and for $k > 3$

$$|\alpha_k^* - \alpha_{k-1}^*| \leq 1, |\beta_k^* - \beta_{k-1}^*| \leq 1.$$

Therefore, we have a set of shifts $\{\alpha_k^*, \beta_k^*\}_{k=3}^{12} \in T$ and

$$\begin{aligned} \widetilde{\text{sum}}_{12}(\alpha_{12}^*, \beta_{12}^*) &= \sum_{k=3}^{12} \text{score}_k(\alpha_k^*, \beta_k^*) \\ &\leq \max_{\{\alpha_k, \beta_k\}_{k=3}^{12} \in T} \sum_{k=3}^{12} \text{score}_k(\alpha_k, \beta_k) = \widetilde{\text{score}}^*(A, B). \end{aligned} \quad (6.19)$$

On the other hand, let a set $\{\alpha_k^*, \beta_k^*\}_{k=3}^{12}$ result in a maximum of Eq. (6.14), i.e. the set is admissible (belongs to T) and

$$\max_{\{\alpha_k, \beta_k\}_{k=3}^{12} \in T} \sum_{k=3}^{12} \text{score}_k(\alpha_k, \beta_k) = \sum_{k=3}^{12} \text{score}_k(\alpha_k^*, \beta_k^*).$$

We will show that for this set we have

$$\widetilde{\text{sum}}_3(\alpha_3^*, \beta_3^*) = \text{score}_k(\alpha_3^*, \beta_3^*) \quad (6.20)$$

and

$$\widetilde{\text{sum}}_k(\alpha_k^*, \beta_k^*) = \text{score}_k(\alpha_k^*, \beta_k^*) + \widetilde{\text{sum}}_{k-1}(\alpha_{k-1}^*, \beta_{k-1}^*), \quad k = 4, 5, \dots, 12. \quad (6.21)$$

The equality (6.20) holds true simply by the definition in (6.16). Let us suppose that for some k Eq. (6.21) fails. Then, for this k using (6.17), we will have

$$\widetilde{\text{sum}}_{k-1}(\alpha_{k-1}^*, \beta_{k-1}^*) < \widetilde{\text{sum}}_{k-1}(\alpha_k^* + i, \beta_k^* + j),$$

where $|i| \leq 1, |j| \leq 1, |\alpha_k^* + i| = |\bar{\alpha}_{k-1}| \leq KX$, and $|\beta_k^* + j| = |\bar{\beta}_{k-1}| \leq KY$. The last inequality gives that there exists a shift set $\{\bar{\alpha}_3, \bar{\beta}_3, \dots, \bar{\alpha}_{k-1}, \bar{\beta}_{k-1}, \alpha_k^*,$

Table 6.2 An example of similarity scores score_k for sections $k = 6, 7, 8,$ and 9 of ICE 240861 and 243395 iris images with horizontal and vertical translations: $\alpha = -2, -1, 0, 1, 2, \beta = -1, 0, 1.$ sum_k give accumulated similarity scores and $\widetilde{\text{sum}}_k$ corresponding elastic sums

Section score						Accumulative score						Elastic score					
$\text{score}_6(\alpha, \beta)$						$\text{sum}_6(\alpha, \beta)$						$\widetilde{\text{sum}}_6(\alpha, \beta)$					
$\frac{\alpha}{\beta}$	-2	-1	0	1	2	$\frac{\alpha}{\beta}$	-2	-1	0	1	2	$\frac{\alpha}{\beta}$	-2	-1	0	1	2
-1	38	38	45	57	73	-1	38	38	45	57	73	-1	38	38	45	57	73
0	38	40	48	62	75	0	38	40	48	62	75	0	38	40	48	62	75
1	40	45	53	64	69	1	40	45	53	64	69	1	40	45	53	64	69
$\text{score}_7(\alpha, \beta)$						$\text{sum}_7(\alpha, \beta)$						$\widetilde{\text{sum}}_7(\alpha, \beta)$					
$\frac{\alpha}{\beta}$	-2	-1	0	1	2	$\frac{\alpha}{\beta}$	-2	-1	0	1	2	$\frac{\alpha}{\beta}$	-2	-1	0	1	2
-1	41	44	54	68	72	-1	79	82	99	125	145	-1	81	92	116	143	147
0	39	44	57	71	72	0	77	84	105	133	147	0	84	97	121	146	147
1	44	50	58	67	63	1	84	95	111	131	132	1	89	103	122	142	138
$\text{score}_8(\alpha, \beta)$						$\text{sum}_8(\alpha, \beta)$						$\widetilde{\text{sum}}_8(\alpha, \beta)$					
$\frac{\alpha}{\beta}$	-2	-1	0	1	2	$\frac{\alpha}{\beta}$	-2	-1	0	1	2	$\frac{\alpha}{\beta}$	-2	-1	0	1	2
-1	48	61	74	73	59	-1	127	143	173	198	204	-1	145	182	220	220	206
0	47	60	73	74	60	0	124	144	178	207	207	0	150	182	219	221	207
1	48	60	68	69	57	1	132	155	179	200	189	1	151	182	214	216	204
$\text{score}_9(\alpha, \beta)$						$\text{sum}_9(\alpha, \beta)$						$\widetilde{\text{sum}}_9(\alpha, \beta)$					
$\frac{\alpha}{\beta}$	-2	-1	0	1	2	$\frac{\alpha}{\beta}$	-2	-1	0	1	2	$\frac{\alpha}{\beta}$	-2	-1	0	1	2
-1	60	68	64	49	38	-1	187	211	237	247	242	-1	242	288	285	270	259
0	64	71	65	48	36	0	188	215	243	255	243	0	246	291	286	269	257
1	60	66	61	50	41	1	192	221	240	250	230	1	242	285	282	271	262

* $\beta_k, \dots, * \alpha_{12}, * \beta_{12}\} \in T$ such that

$$\begin{aligned}
 \sum_{i=3}^{12} \text{score}_i(*\alpha_i, *\beta_i) &= \sum_{i=3}^{k-1} \text{score}_i(*\alpha_i, *\beta_i) + \sum_{i=k}^{12} \text{score}_i(*\alpha_i, *\beta_i) \\
 &< \sum_{i=3}^{k-1} \text{score}_i(\bar{\alpha}_i, \bar{\beta}_i) + \sum_{i=k}^{12} \text{score}_i(*\alpha_i, *\beta_i), \quad (6.22)
 \end{aligned}$$

which contradicts to the assumption that the set $\{*\alpha_k, *\beta_k\}_{k=3}^{12}$ is a maximum of (6.14). The obtained contradiction proves Eq. (6.21). Therefore,

$$\begin{aligned}
 \widetilde{\text{sum}}_{12}(\alpha_{12}^*, \beta_{12}^*) &= \max_{|\alpha| \leq KX, |\beta| \leq KY} \widetilde{\text{sum}}_{12}(\alpha, \beta) \\
 &\geq \widetilde{\text{sum}}_{12}(*\alpha_{12}, *\beta_{12}) = \widetilde{\text{score}}^*(A, B). \quad (6.23)
 \end{aligned}$$

By combining (6.19) and (6.23) inequalities, we get (6.18).

Table 6.2 illustrates ordinary and elastic similarity scores. For short only similarity scores of sections 6–9 for ICE 240861 and 243395 iris images are presented. The scores are represented as percentage of alike bits of shifted sectors. For this example,

ordinary similarity scores are $\text{score}^*(A, B) = \max_{\alpha, \beta} \text{sum}_9(\alpha, \beta) = \text{sum}_9(1, 0) = 255 = 255/4$ (%) = 64 (%). The elastic similarity scores are $\widetilde{\text{score}}^*(A, B) = \max_{\alpha, \beta} \widetilde{\text{sum}}_9(\alpha, \beta) = \widetilde{\text{sum}}_9(-1, 0) = 291 = 291/4$ (%) = 73 (%). The admissible shifts optimal set is the following: $(\alpha_6^*, \beta_6^*) = (2, 0)$, $(\alpha_7^*, \beta_7^*) = (1, 0)$, $(\alpha_8^*, \beta_8^*) = (0, -1)$, $(\alpha_9^*, \beta_9^*) = (-1, 0)$. The ICE (240861,243395) is a genuine pair, and for this pair the elastic vs. ordinary measure gives a significant increase of similarity scores.

6.3.5 Similarity Metric for Comparison of Local Feature Maps

In this and the next subsections, we present a highly efficient similarity metric for the local extrema features.

Let A and B be sets of local features formed from our Taylor local extrema estimated at different scales. A particular feature $f_i = \{x_i, y_i, s_i, z_i\}$ has the following attributes:

- (x_i, y_i) – local extremum position defined by two integer numbers x_i and y_i .
- $s_i \in \{0, 1, \dots, 7\}$ – scale and filter-type index (see Table 6.1 for details).
- $z_i \in \{0, 1\}$ – extremum-type marker. 0 was used for local maximum points and 1 for local minimum points.

The similarity metric between two local feature sets is defined by a normalised sum of similarity scores of feature pairs. A particular feature pair (f_i^A, f_j^B) can gain non-zero similarity scores only if f_i^A and f_j^B correspond to the same filter and extremum type. In the following, we will call such features *congeneric*.

Definition 6.2. Two local features $f_i^A = \{x_i^A, y_i^A, s_i^A, z_i^A\}$ and $f_j^B = \{x_j^B, y_j^B, s_j^B, z_j^B\}$ are *congeneric* iff they (a) appear at the same scale, i.e. $s_i^A = s_j^B$, and (b) correspond to the same type of extremum, i.e. $z_i^A = z_j^B$.

Two congeneric features can gain a positive similarity only if their positions are close. The distance between two affine (x_i^A, y_i^A) and (x_j^B, y_j^B) points is estimated again by considering a possible alignment of the two irises by means of an angular rotation and radial translation. Angular rotation and radial translation correspond to some shift along x and y directions, respectively.¹ With fixed α and β shifts in x and y direction, we use the following score similarity expression:

$$s(f_i^A, f_j^B; \alpha, \beta) = \begin{cases} 0 & \text{if } s_i^A \neq s_j^B \text{ or } z_i^A \neq z_j^B; \\ |DX - |x_i^A - x_j^B - \alpha||_+ \times |DY - |y_i^A - y_j^B - \beta||_+ & \text{otherwise.} \end{cases} \quad (6.24)$$

Here, DX and DY are two positive integer parameters that define maximal acceptable distance along x and y directions, and

¹Shift along the x direction inherits angular periodicity.

$$|a|_+ = \begin{cases} a & \text{if } a > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (6.25)$$

For fixed angular and radial shifts α and β , the proposed expression allocates maximum similarity for the positions (x_i^A, y_i^A) and (x_j^B, y_j^B) that coincide after α and β shifts. The relative decay of the similarity is controlled by constants DX and DY (respectively in x and y directions). Greater DX and DY values mean that more intra-class variation of positions is tolerated. Note that the defined similarity measure is a piecewise bilinear function with respect to the angular and radial shifts. We will utilise this later for fast calculation of similarity scores.

The similarity scores between the feature set A and an α rotated and β radially shifted feature set B is defined as the sum of the similarity scores of local features pairs, i.e.

$$\text{score}(A, B; \alpha, \beta) = \sum_{i=1}^{NA} \sum_{j=1}^{NB} s(f_i^A, f_j^B; \alpha, \beta), \quad (6.26)$$

where NA and NB are given by the number of features in A and in B , respectively. The final similarity score between feature sets A and B is then defined by the following maximum:

$$\text{score}^*(A, B) = q(NA, NB) \max_{\substack{-KX \leq \alpha \leq KX \\ -KY \leq \beta \leq KY}} \text{score}(A, B; \alpha, \beta), \quad (6.27)$$

with KX and KY as defined in the previous section and $q(NA, NB) = \frac{1}{\min(NA, NB)}$ is a normalisation factor.

6.3.6 Efficient Calculation of Similarity Scores for Local Features Maps

The intuitive approach for calculating the similarity scores as defined by Eqs. (6.24)–(6.27) requires $\mathcal{O}((2KX + 1) \times (2KY + 1) \times NA \times NB)$ arithmetic operations. We will now provide a more efficient method for calculating the $(2KX + 1) \times (2KY + 1)$ score matrix elements $\text{score}(A, B; \alpha, \beta)$ with at most $\mathcal{O}((2KX + 1) \times (2KY + 1)) + \mathcal{O}(NA \times NB)$ arithmetic operations. The idea is based on the observation that $\text{score}(\alpha, \beta)$, as defined in (6.26) and (6.27), is piecewise bilinear with respect to α and β and that it consists of $|DX - |\alpha - a||_+ \times |DY - |\beta - b||_+$ terms. Thus, the numerical derivatives

$$\begin{aligned} s(\alpha, \beta)_{\alpha\alpha, \beta\beta} &= s(\alpha + 1, \beta)_{\beta\beta} - 2s(\alpha, \beta)_{\beta\beta} + s(\alpha - 1, \beta)_{\beta\beta}, \\ s(\alpha, \beta)_{\beta\beta} &= s(\alpha, \beta + 1) - 2s(\alpha, \beta) + s(\alpha, \beta - 1), \end{aligned} \quad (6.28)$$

are equal to zero at linear pieces. Here, and occasionally later, we use the short notation

$$s(\alpha, \beta) = \text{score}(A, B; \alpha, \beta). \quad (6.29)$$

The piecewise bilinearity of $\text{score}(A, B; \alpha, \beta)$ allows us to show the following Lemma:

Proposition 6.2. *Let $A = \{f_i^A = (x_i^A, y_i^A, s_i^A, z_i^A)\}_{i=1}^{NA}$ and $B = \{f_j^B = (x_j^B, y_j^B, s_j^B, z_j^B)\}_{j=1}^{NB}$ be local features sets, $s(\alpha, \beta) = \text{score}(A, B; \alpha, \beta)$ and*

$$\text{hits}(\alpha, \beta) = \#\{(i, j) : s_i^A = s_j^B, z_i^A = z_j^B, x_i^A - x_j^B = \alpha, y_i^A - y_j^B = \beta\} \quad (6.30)$$

is the counting function of pairs of congeneric features that can be superpositioned by the α and β shifts along x and y directions. Then, the following equation is satisfied:

$$s(\alpha, \beta)_{\alpha\alpha, \beta\beta} = \sum_{i=-1}^1 \sum_{j=-1}^1 \gamma_{i,j} \text{hits}(\alpha + iDX, \beta + jDY) \quad (6.31)$$

with

$$|\gamma_{i,j}|_{i,j=-1}^1 = \begin{vmatrix} 1 & -2 & 1 \\ -2 & 4 & -2 \\ 1 & -2 & 1 \end{vmatrix}. \quad (6.32)$$

Proof. Direct calculation of $\delta(\alpha, \beta) = (|DX - |\alpha - a||_+ |DY - |\beta - b||_+)_{\alpha\alpha, \beta\beta}$ for all integer $DX, DY, a = x_i^A - x_j^B$, and $b = y_i^A - y_j^B$, gives

$$\delta(\alpha, \beta) = \begin{cases} 4 & \text{if } \alpha = a, \beta = b, \\ -2 & \text{if } \alpha = a, \beta = b \pm DY \text{ or } \alpha = a \pm DX, \beta = b, \\ 1 & \text{if } \alpha = a \pm DX, \beta = b \pm DY, \\ 0 & \text{otherwise.} \end{cases} \quad (6.33)$$

Successive application of this formula to (6.27) and (6.26) yields (6.31).

Direct calculation of the counting matrix $\{\text{hits}(\alpha, \beta)\}$, with $\alpha = -KX \dots KX$ and $\beta = -KY \dots KY$, requires $\mathcal{O}(NA \times NB)$ arithmetic operations. However, using the pre-calculated tables

$$\text{index}^A(x, y; s, z) = \{i : s_i^A = s, z_i^A = z, x = x_i^A - \alpha, y = y_i^A - \beta, |\alpha| \leq KX, |\beta| \leq KY\}, \\ x = 0, 1, \dots, W - 1, y = 0, 1, \dots, H - 1,$$

the *hits* counting algorithm

$$\text{for}(j = 1; j \leq NB; ++j) \\ \text{for}(i \in \text{index}^A(x_j^B, y_j^B; s_j^B, z_j^B)) \\ ++ \text{hits}(x_i^A - x_j^B, y_i^A - y_j^B)$$

would require only $\mathcal{O}(NA \times NB \times KX \times KY / (W \times H))$ arithmetic operations. Here, W and H denote width and height of the unwrapped iris image, and it is assumed that local extrema points are uniformly distributed within the unwrapped rectangular area of the iris image. In our experimentations with real iris data, we had the following settings: NA and NB were in the order of 300, $KX = 10$, $KY = 2$, $W = 256$, and $H = 32$. With these, the $\mathcal{O}(NA \times NB \times KX \times KY / (W \times H))$ complexity is significantly lower than $\mathcal{O}(NA \times NB)$.

For identification, the pre-calculation cost of the index tables is not significant. However, pre-calculation may slow down the proposed *hits* counting algorithm in the verification case.

The fast computation of the *hits* function leads to fast calculation of the right hand side of the (6.31). To efficiently compute the similarity score matrix $\{s(\alpha, \beta)\}_{\alpha=-KX, \beta=-KY}^{KX, KY}$, we require a fast inversion of $s_{\alpha\alpha, \beta\beta}$. Let us consider first the inversion problem in the one-dimensional case with the second-order integer numerical derivative

$$s_{\alpha\alpha}(\alpha) = s(\alpha - 1) - 2s(\alpha) + s(\alpha + 1), \quad (6.34)$$

with the initial condition

$$s(\alpha) \equiv 0 \text{ for all } \alpha \leq 0. \quad (6.35)$$

We are looking for the values $s(\alpha)$. Numerical integration for $s(\alpha)$ gives

$$\begin{aligned} I s_{\alpha\alpha}(\alpha) &= \sum_{i < \alpha} s_{\alpha\alpha}(\alpha) = \sum_{i < \alpha} (s(i+1) - s(i) - (s(i) - s(i-1))) \\ &= s(\alpha) - s(\alpha - 1). \end{aligned} \quad (6.36)$$

Applying numerical integration and the initial conditions $s(\alpha) \equiv 0, \forall \alpha \leq 0$ for the second time, we get

$$\begin{aligned} I^2 s_{\alpha\alpha}(\alpha) &= \sum_{i \leq \alpha} I s_{\alpha\alpha}(\alpha) = \sum_{i \leq \alpha} (s(i) - s(i-1)) \\ &= \sum_{i \leq \alpha} s(i) - \sum_{i \leq \alpha-1} s(i) = s(\alpha). \end{aligned} \quad (6.37)$$

Notice that generally the double numerical integration is computationally unstable. However, under our assumptions where we are dealing with integer values $s(\alpha - 1) - 2s(\alpha) + s(\alpha + 1)$ and $s(\alpha) \equiv 0, \forall \alpha \leq 0$, the double numerical integration uses only integer arithmetics which is computationally stable.

In the two dimensional case, if $s(\alpha, \beta) = 0$ for all $\alpha \leq 0$ and $\beta \leq 0$, we have

$$\begin{aligned} I^{2,1} s_{\alpha\alpha, \beta\beta}(\alpha, \beta) &= \sum_{j < \beta} I^{2,0} s_{\alpha\alpha, \beta\beta}(\alpha, \beta) = \sum_{j < \beta} (s(\alpha, j+1) - 2s(\alpha, j) + s(\alpha, j-1)) \\ &= s(\alpha, \beta) - s(\alpha, \beta - 1) \end{aligned} \quad (6.38)$$

and

$$\begin{aligned} I^{2,2}s_{\alpha\alpha,\beta\beta}(\alpha,\beta) &= \sum_{j\leq\beta} I^{2,1}s_{\alpha\alpha,\beta\beta}(\alpha,\beta) = \sum_{j\leq\beta} (s(\alpha,\beta) - s(\alpha,\beta-1)) \\ &= s(\alpha,\beta). \end{aligned} \quad (6.39)$$

Thus, repeated summations allow to compute the inverse of $s_{\alpha\alpha,\beta\beta}$.

To summarise:

Proposition 6.3. *Let*

$$\begin{aligned} u(\alpha,\beta) &= 0 \text{ for all sufficient small } \alpha \text{ or } \beta, \\ I^{0,0}u(\alpha,\beta) &= u(\alpha,\beta), \\ I^{n,m}u(\alpha,\beta) &= \sum_{i\leq\alpha-n\%2} I^{n-1,m}u(i,\beta), n > 0, m \geq 0, \\ I^{n,m}u(\alpha,\beta) &= \sum_{j\leq\beta-m\%2} I^{n,m-1}u(\alpha,j), n \geq 0, m > 0, \end{aligned}$$

and coefficients $\{\gamma\}_{i,j=-1}^1$ are defined by (6.32). Then $\forall |\alpha| \leq KX, |\beta| \leq KY$,

$$\text{score}(A, B; \alpha, \beta) = \sum_{i=-1}^1 \sum_{j=-1}^1 \gamma_{i,j} I^{2,2} \text{hits}(\alpha + iDX, \beta + jDY). \quad (6.40)$$

6.4 Experimental Results

In this section, we document our experiments with the three approaches. For the experiments, we have used three publicly available iris databases in all experiments: the Chinese Academy of Sciences Casia 2.0 (device1) [6], the US National Institute of Standards and Technology (NIST) ‘‘Iris Challenge Evaluation’’, experiment 1 (right eye), (ICE-1) [22], and the NIST ‘‘Multiple Biometric Grand Challenge’’ (MBGC), Portal Challenge, experiment 3, left eye (MBGC-31) [25].

Initial iris segmentation was performed fully automatic for all iris images in the datasets ICE-1 and Casia 2.0 using the methods in Sect. 6.2. For an automatic segmentation of the MBGC-31 dataset, the methods in Sect. 6.2 were slightly enhanced. The MBGC Portal Challenge MBGC-31 dataset contains a gallery set with 571 high-resolution ($2,048 \times 2,048$) near-infrared (NIR) videos of faces that were acquired from a Sarnoff Iris on the Move (IoM) system [18]. The IoM system was designed to capture iris imagery as a person walks through the portal. We observed that our iris segmentation algorithm from Sect. 6.2 works slowly and unreliably for the high-resolution video data that contains almost a full face image. To improve the iris segmentation performance, we used an approach which detects reflections of portal

Table 6.3 Dependency of average size (in bytes) of one template and their standard deviations on iris databases and maximum allowable local extrema amount in one section (number K)

K	Casia 2.0	ICE-1	MBGC-31
1	190 ± 6	177 ± 7	191 ± 3
2	328 ± 13	288 ± 15	330 ± 9
3	452 ± 20	379 ± 24	457 ± 18
5	619 ± 34	498 ± 42	637 ± 43
7	679 ± 44	540 ± 51	710 ± 62

illumination system. The reflections form patterns of eight bright spots, which are usually situated on the iris or even the pupil. When reflection patterns of the left or right eye are detected, a rectangle, centred by the average position of reflections, is cropped. This cropped image contains the iris in the usual proportions so that it can be reliably handled by our original segmentation algorithm [3].

6.4.1 Size of Template

Different irises have each a distinct number of local sharp variations, so that the template size for different irises using our proposed features is variable. The average template size can be controlled by the parameter K , which defines the upper limit of the number of *congeneric* local features in one sector (see Fig. 6.4). In all our experiments, we used the eight sectors in $\{3, 4, 5, 6, 9, 10, 11, 12\}$.

Each sector in an unwrapped rectangular iris image (of width $W = 256$ and height $H = 32$) has 16×32 pixels. Congeneric features of one sector are written to the template in the following order: $\{N, x'_1, y_1, x'_2, y_2, \dots, x'_N, y_N\}$, where N is number of congeneric features, with (x'_n, y_n) , $n = 1, \dots, N$ being the positions of local features in a sector. We assumed that $K \leq 7$. Thus, $0 \leq n \leq 8$, $0 \leq x'_i \leq 16$, and $0 \leq y_i \leq 32$ can be represented in 3, 4, and 5 bits, respectively. Therefore, the maximal number of bytes in one template is at most $((bitsForN + (bitsForX'_i + bitsForY_i) * K) * (\# \text{ of used sectors}) * (\# \text{ of different congeneric classes}))/8 = (3 + (4 + 5) * 7) * 8 * 16)/8 = 1,056$. Table 6.3 shows the real average iris template sizes for the three different databases corresponding to maximal allowable local extrema number K per sector. The average template size in Casia 2.0, ICE-1, and MBGC-31 is 643 bytes which constitutes 61% of the possible maximal template size for maximum value $K = 7$. For $K = 5$ the average template size is comparable to the size of Daugman's iris code [7]. Furthermore, we can observe clusters of local extrema at the same (x'_i, y_i) locations but at different scales. These clusters can be used to compress the iris templates even more; however, a further analysis of such compressibility was not performed.

6.4.2 Verification Performance

We use detection error tradeoff (DET) curves [19] to document the quality of our approaches. The DET curves plot the false acceptance rate (FAR) vs. the false rejection rate (FRR) along the log-log axes for different similarity or distance threshold values. Our experiments show that the verification quality as expected increases with increasing maximum number K of congeneric features per sector. The results presented in this section were achieved with $K = 7$.

In the following, we present the results for:

1. Matching two local maxima feature sets A and B , $\widetilde{\text{score}}_{\text{local extr.}}^*(A, B)$.
2. For matching two binary feature sets A and B , $\widetilde{\text{score}}_{\text{phase}}^*(A, B)$.
3. After proper normalisation, these two elastic similarities are combined to define the joint feature:

$$\widetilde{\text{score}}_{\text{local extr.} + \text{phase}}^*(A, B) = \frac{\widetilde{\text{score}}_{\text{local extr.}}^*(A, B) + \widetilde{\text{score}}_{\text{phase}}^*(A, B)}{2}. \quad (6.41)$$

For Casia 2.0 (device 1) iris database, we have evaluated ${}_{20}C_2 \times 60 = 11,400$ genuine and ${}_{60}C_2 \times 20^2 = 708,000$ impostor similarity scores. It is recognised that Casia 2.0 iris database has lower quality than Casia 1.0 [20] images. The best result that can be found in the literature is $EER = 0.58\%$ [20]. In our experiments, we obtained considerably better results already for the proposed local extrema features of the multi-scale Taylor expansion coefficients: $EER = 0.13\%$, $\text{ZeroFAR} = 1.63\%$, and $\text{ZeroFRR} = 5.52\%$. The corresponding verification performance based solely on the elastic phase based metric is with $EER = 0.07\%$, $\text{ZeroFAR} = 2.1\%$, and $\text{ZeroFRR} = 0.49\%$ similar. The similarity measure for the joint features $\widetilde{\text{score}}_{\text{local extr.} + \text{phase}}^*$ gives even considerably better results: $EER = 0.026\%$, $\text{ZeroFAR} = 0.65\%$, and $\text{ZeroFRR} = 0.044\%$. In Fig. 6.6, the DET curves for Casia 2.0 iris database are shown. In this and all following DET curves, the results based on local maxima features are marked with filled circles, the results for the binarised phase-based features are marked by filled triangles, and the joint features are marked by + signs.

For the NIST ‘‘Iris Challenge Evaluation’’, experiment 1, (ICE-1) iris database local extrema-based verification algorithm relatively produced the worst results (See Fig. 6.7). The image database contains very difficult and corrupted examples, sometimes with big eyelid occlusion, off-angle, or with an iris partly outside the image frame. Poor focus of a part of the iris images especially degrades the quality of the verification results. Out of focus iris images lose all subtle details of the iris texture, which results in a significant decay of number of Taylor decomposition local extrema at tiny scales. For the ICE-1 iris database, we have evaluated 12,214 genuine and 1,002,386 impostor similarity scores. We obtained for the local maxima features $EER = 0.25\%$, $frr@far0.1 = 0.42\%$, $frr@far0.01 = 0.83\%$, the binarised phase features showed better results than the ones of the local extrema features: $EER = 0.15\%$, $frr@far0.1 = 0.20\%$, $frr@far0.01 = 0.45\%$.




Fig. 6.6 DET curves for Casia 2.0 (device 1) iris database (Figure reproduced with permission from [2])




Fig. 6.7 DET curves for ICE1 iris database (Figure reproduced with permission from [2])

This is slightly worse than the best reported results of a newer Daugman algorithm [10]: $EER = 0.11\%$, $frr@far0.1 = 0.12\%$, $frr@far0.01 = 0.30\%$. Daugman's algorithm concentrates on segmentation improvements, which tolerate off-angle and deformed iris inner and outer boundaries, which are present in the ICE-1 iris database. The iris features and matching routines in [10] still remain phase based. Using the $\widehat{score}_{local\ extr.+phase}^*$ similarity measure gives $EER = 0.063\%$, $frr@far0.1 = 0.04\%$, $frr@far0.01 = 0.18\%$, and considerably outperforms the ones of [10].

The design of MBGC problem was conditioned by a number of new iris recognition systems which should allow to perform human iris identification at a distance and on the move. The following iris data were collected for this challenge:

- Still iris images were collected with an LG2200 iris camera with intentionally degraded quality as described in [5].
- Iris video sequences were collected with the same LG2200 iris camera, digitised and transcoded to MPEG-4 format with high bit rate allowance, thus yielding near-lossless encoding.
- Very high-resolution near-infrared (NIR) video of faces were acquired from a Sarnoff Iris on the Move (IoM) system [18]. The IoM system was designed to capture iris imagery as a person walks through the portal. One frame of NIR video has $2,048 \times 2,048$ pixels, and average diameter of one iris is approximately 140 pixels.

Twelve experiments in five different categories were defined in this challenge. We chose experiment MBGC-3l in which NIR video data is compared with left iris still image. Following the decision of NIST, we skipped MBGC-3r experiment for right irises, since right irises were missing in 25 out of 571 videos.

The DET curve for MBGC-3l experiment was calculated following the MBGC protocol. The protocol allowed normalisation of similarity scores, and we used z -score normalisation for this purpose. Following the MBGC protocol, 27,671 genuine and 603,021 impostor similarity scores were evaluated. Because of pure iris resolution (only about 140 pixels in iris diameter) and different data acquisition conditions (video vs. still iris images), it is hard to expect very high verification results. We achieve with the local maxima features $EER = 2.27\%$, $ZeroFAR = 28\%$, $frr@far1 = 2.93\%$, and $frr@far0.1 = 6.4\%$. Only a small amount of work on iris recognition using the MBGC data has been published so far. Some preliminary results of mainly commercial systems were presented at a workshop [24]. In [12], Hollingsworth et al. proposed an improvement and speedup of iris recognition using a signal-level fusion of iris frames from video. Their paper presents some experimental results, where MBGC videos for both gallery and probe sets were used. In [14], methods to detect eyes in the MBGC portal videos and to measure the quality of the extracted eye images were presented. They achieved a false rejection rate of 43.90% at a false acceptance rate of 0.80%. In [33], Zhou et al. suggest to add some additional steps to the traditional iris system in order to improve the performance. After inclusion of four additional steps, they received $frr@far0.1 = 28.6\%$ and $frr@far0.01 = 42.7\%$ for the MBGC-3l dataset. For the




Fig. 6.8 DET curves for MBGC, Portal Challenge, experiment 3L, iris database (Figure reproduced with permission from [2])

joint similarity measure $\widetilde{\text{score}}_{\text{local extr. + phase}}^*$, we reach an error rate of $EER = 1.21\%$, $\text{ZeroFAR} = 12.6\%$, $frr@far1 = 1.27\%$, and $frr@far0.1 = 2.98\%$. Analysis of DET curves of the three different algorithms (see Fig. 6.8) indicates that for the MBGC iris dataset algorithm based on local extrema outperforms the phase-based algorithm. This can be caused by local extrema being more robust to different lightning conditions and low resolution of iris video images.

6.5 Summary and Conclusions

It is interesting to note that despite of the intuitive similarity to the Gabor features, the Taylor's coefficients have greater localisation in the space domain. The localisation is gathered at the expense of poor feature resolution in the frequency domain. Such proportion of localisation in space and frequency domains is natural for the iris where local objects as freckles, furrows, stripes, and coronas dominate the texture and where periodic texture patterns that are well localised in frequency domain are rare.

The local maxima and the binary features seem to be complementing each other so that their fusion gives a verification performance that is among the best verification results published so far in academic papers. Despite of the noticeable improvement of the iris verification quality due to the fusion of the similarity

scores, we should mention that our binarised phase matrix increases the template size by 2,048 bytes and significantly decreases the matching speed. The matching speed decreases because of the warped similarity calculation. Using only a 3.2 GHz Pentium 4 processor, we were able to compare $\approx 5,000$ iris templates per second. It is likely, however, that more accurate iris segmentation as done in [10] can eliminate the necessity to search for optimal shifts along radial direction.

References

1. Bae, K., Noh, S., Kim, J.: Iris feature extraction using independent component analysis. In: Proceedings of the International Conference on Audio- and Video-Based Biometric Person Authentication. Lecture Notes in Computer Science #2688, pp. 1059–1060. Springer-Verlag Berlin (2003)
2. Bastys, A., Kranauskas, J., Krüger, V.: Iris recognition by fusing different representations of multi-scale taylor expansion. *Comput. Vis. Image Underst.* **115**(6), 804–816 (2011)
3. Bastys, A., Kranauskas, J., Masiulis, R.: Iris recognition by local extremum points of multiscale taylor expansion. *Pattern Recognit.* **42**(9), 1869–1877 (2009)
4. Boles, W., Boashash, B.: A human identification technique using images of the iris and wavelet transform. *IEEE Trans. Signal Process.* **46**, 1185–1188 (1998)
5. Bowyer, K.W., et al.: Factors that degrade the match distribution in iris biometrics. *Identit. Inf. Soc.* **2**(3), 327–343 (2009) <http://www.springerlink.com/content/322732332v683301/>
6. Chinese Academy of Sciences: Institute of Automation Iris Database 1.0 (2003)
7. Daugman, J.: High confidence visual recognition of persons by a test of statistical independence. *IEEE Trans. Pattern Anal. Mach. Intell.* **15**(11), 1148–1161 (1993)
8. Daugman, J.: Statistical richness of visual phase information: update on recognizing persons by iris patterns. *Int. J. Comput. Vis.* **45**(1), 25–38 (2001)
9. Daugman, J.: Demodulation by complex valued wavelets for stochastic pattern recognition. *Int. J. Wavelets Multiresolution Inf. Process.* **1**(1), 1–17 (2003)
10. Daugman, J.: New methods in iris recognition. *IEEE Trans. Syst. Man Cybern.* **37**(5), 1167–1175 (2007)
11. Hansen, J., Witzner, D.W., Ji, Q.: In the eye of the beholder: a survey of models for eyes and gaze. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(3), 478–500 (2010)
12. Hollingsworth, K., et al.: Iris recognition using signal-level fusion of frames from video. *IEEE Trans. Inf. Forensics Secur.* **4**(4), 837–848 (2009)
13. Kruskal, J.B., Liberman, M.: The symmetric time-warping problem: from continuous to discrete. In: Sanko, D., Kruskal, J.B. (eds.) *Time Warps, String Edits, and Macromolecules: The Theory and Practice of Sequence Comparison*. Addison-Wesley, Reading (1983)
14. Lee, Y.Y., Phillips, P.J., Micheals, R.J.: An automated video-based system for iris recognition. In: Tistarelli, M., Nixon, M.S. (eds.) *Proceedings of the International Conference on Biometrics. Advances in Biometrics, Lecture Notes in Computer Science #5558* pp. 1160–1169. Springer Berlin Heidelberg (2009)
15. Lim, S., et al.: Efficient iris recognition through improvement of feature vector and classifier. *ETRI J.* **23**(2), 1–70 (2001)
16. Ma, L.: Person identification based on iris recognition. Ph.D. dissertation, Inst. Automation, Chinese Academy of Sciences (2003)
17. Ma, L., et al.: Efficient iris recognition by characterizing key local variations. *IEEE Trans. Image Process.* **13**(6), 739–750 (2004)
18. Matey, J.R., et al.: Iris on the move: acquisition of images for iris recognition in less constrained environments. *Proc. IEEE* **94**(11), 1936–1946 (2009)
19. Metz, C.E.: Basic principles of roc analysis. *Semin. Nucl. Med.* **8**, 283–298 (1978)

20. Miyazawa, K., et al.: A phase-based iris recognition algorithm. *Lect. J. Comput. Science. Advances in Biometrics, Lecture Notes in Computer Science #3832* pp. 356–365. Springer-Verlag Berlin Heidelberg (2005)
21. Monro, D.M., Rakshit, S., Zhang, D.: Dct-based iris recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(4), 586–595 (2007)
22. National Institute of Science and Technology (NIST), T.: Iris challenge evaluation (ICE) (2009) <http://www.nist.gov/itl/iad/ig/ice.cfm>
23. Park, C., et al.: Iris-based personal authentication using a normalized directional energy feature. *Audio- and Video-Based Biometric Person Authentication, Lecture Notes in Computer Science #2688* pp. 224–232. Springer-Verlag Berlin Heidelberg (2003)
24. Phillips, P.J.: Mbgc presentations and publications (2009) <http://www.nist.gov/itl/iad/ig/mbgc.cfm>
25. Phillips, P., et al.: Overview of the multiple biometrics grand challenge. In: Tistarelli, M., Nixon, M.S. (eds.) *Advances in Biometrics, Lecture Notes in Computer Science* vol. 5558, pp. 705–714. Springer, Berlin (2009)
26. Sahbi, H., Boujemaa, N.: Robust matching by dynamic space warping for accurate face recognition. In: *Proceedings of the International Conference on Image Processing*, vol. 1, pp. 1010–1013. IEEE, Piscataway (2001)
27. Sanchez-Avila, C., Sanchez-Reillo, R.: Iris-based biometric recognition using dyadic wavelet transform. *IEEE Aerosp. Electron. Syst. Mag.* **17**, 3–6 (2002)
28. Sanchez-Avila, C., Sanchez-Reillo, R.: Two different approaches for iris recognition using gabor filters and multiscale zero-crossing representation. *Pattern Recognit.* **38**, 231–240 (2005)
29. Sun, Z., Tan, T., Qiu, X.: Graph matching iris image blocks with local binary pattern. In: *Proceedings of International Conference on Biometrics. Advances in Biometrics, Lecture Notes in Computer Science #3832* pp. 366–372. Springer Berlin Heidelberg (2005)
30. Wang, Y., Zhu, Y., Tan, T.: Biometric personal identification based on iris patterns. *International Conference on Pattern Recognition, IEEE Publishing*, **2**, 801–804, Barcelona, Spain (2000)
31. Wildes, R., et al.: A machine-vision system for iris recognition. *Mach. Vis. Appl.* **9**, 1–8 (1996)
32. Yu, L., et al.: Multiscale wavelet texture based iris verification. In: *Proceedings of the International Conference on Wavelet Analysis and Its Applications (WAA)*, pp. 200–205. World Scientific, River Edge (2003)
33. Zhou, Z., Du, Y., Belcher, C.: Transforming traditional iris recognition systems to work in nonideal situations. *IEEE Trans. Ind. Electron.* **56**(8), 3203–3213 (2009)

Chapter 7

A Theoretical Model for Describing Iris Dynamics

Antwan D. Clark, Scott A. Kulp, Isom H. Herron, and A.A. Ross

Abstract We present a theoretical approach using what we know about tissue dynamics to explore the nonlinear dynamics of iris deformation. Current iris recognition algorithms assume a simple transformation to approximate the deformation of the iris tissue. Furthermore, current research work on iris deformation does not take into account the mechanical properties of the iris tissue nor the cause of deformation from the iris muscle activity. By looking at the tissue dynamics, we are able to gain a more comprehensive understanding of this deformation process. The results of this research work can potentially be leveraged into existing iris recognition systems.

7.1 Introduction

Biometrics is the science of establishing human identity based on the physical or behavioral attributes of an individual such as fingerprints, face, iris, voice, gait, and signature [9]. The deployment of biometric systems in several identity management and access control systems ranging from laptops to border security has demonstrated the potential of using biometrics for human recognition.


A.D. Clark (✉)
University of Maryland, Maryland, USA
e-mail: aclark.biometrics@gmail.com

S.A. Kulp
Rutgers University, Piscataway, NJ, USA
e-mail: sckulp@cs.rutgers.edu

I.H. Herron
Rensselaer Polytechnic Institute, Troy, NY, USA
e-mail: herroi@rpi.edu

A.A. Ross
West Virginia University, Morgantown, WV, USA
e-mail: arun.ross@mail.wvu.edu

Fig. 7.1 Iris segmentation using an integro-differential operator. *White lines* are used to denote the output of the circle-fitting algorithm



Iris recognition systems, in particular, are gaining interest because the iris's rich texture offers a strong biometric cue for recognizing individuals [19]. Located just behind the cornea and in front of the lens, the iris uses the dilator and sphincter muscles that govern pupil size to control the amount of light that enters the eye. Near-infrared (NIR) images of the iris's anterior surface exhibit complex patterns that computer systems can use to recognize individuals. Because NIR lighting can penetrate the iris's surface, it can reveal the intricate texture details that are present even in dark-colored irides. The iris's textural complexity and its variation across eyes have led scientists to postulate that the iris is unique across individuals. Further, the iris is the only internal organ readily visible from the outside. Thus, unlike fingerprints or palm prints, environmental effects cannot easily alter its pattern.

The history of iris recognition dates back to 1936 but was revolutionized in the mid-1990s when John Daugman [3] developed an algorithm to automate the process of iris recognition. Most iris recognition systems consist of five basic modules leading to a decision [19]:

1. The *acquisition* module obtains a 2D image of the eye using a monochromatic CCD camera sensitive to the NIR light spectrum.
2. The *segmentation* module localizes the iris's spatial extent in the eye image by isolating it from other structures in its vicinity, such as the sclera, pupil, eyelids, and eyelashes. Typically, a variant of the integro-differential operator is used for this purpose (Fig. 7.1).
3. The *normalization* module invokes a geometric normalization scheme to transform the segmented iris image from Cartesian coordinates to polar coordinates.
4. The *encoding* module uses a feature-extraction routine to produce a binary code.
5. The *matching* module determines how closely the produced code matches the encoded features stored in the database.

The role of the normalization module is to convert the localized iris from a near-circular entity to a rectangular entity. The process, often called iris unwrapping, has three advantages:




Fig. 7.2 Iris normalization. **(a)** The normalization routine converts the pixel coordinates in the annular region between the pupillary and limbus boundaries to pseudo-polar coordinates. This addresses the problem of variations in pupil size across multiple images. **(b)** Example of a normalized iris

1. It accounts for variations in pupil size due to changes in external illumination that might influence iris size.
2. It ensures that the irides of different individuals are mapped onto a common image domain in spite of the variations in pupil size across subjects.
3. It enables iris registration during the matching stage through a simple translation operation that can account for in-plane eye and head rotations.

The normalization module, denoted by $I(x(r, \theta), y(r, \theta)) \mapsto I(r, \theta)$, is given by the following transformation:

$$\begin{aligned} x(r, \theta) &= (1 - r)x_p(\theta) + rx_l(\theta) \\ y(r, \theta) &= (1 - r)y_p(\theta) + ry_l(\theta), \end{aligned} \quad (7.1)$$

where (x_p, y_p) and (x_l, y_l) correspond to the Cartesian coordinates of the pupil and limbus boundary, respectively. Figure 7.2 denotes the normalization process. After applying (7.1) to normalize the dilation effects, encoding is performed (typically using Gabor wavelets) in order to convert the information in the iris texture to a bit sequence used for comparison. Current commercial systems use Daugman's algorithms, which have demonstrated good matching accuracy on large datasets. However, there still exist challenges in iris recognition, including the effect of varying levels of pupil dilation. Recently, researchers showed that changes in pupil dilation have the potential to degrade iris recognition performance. Hence, there is a need to explore ways to counteract iris deformation caused by pupil dilation.

7.1.1 Related Work

The problem of pupil dilation and iris deformation has recently been explored and noted by researchers in biometrics as well as other areas. In their work, Ma et al. noticed that there were a number of false non-matches due to pupil dilation. Later, Thornton et al. [21] took pupil dilation into account in their work by using Bayesian estimation to recover the level of iris deformation. Their results illustrated that estimating the level of iris deformation leads to an improvement in iris recognition performance. In addition, Wei et al. [23] attempted to consider the effects of pupil dilation by modeling the nonlinear iris stretch as the sum of a linear stretch and Gaussian deviation term. Hollingsworth et al. [8] investigated this problem further by conducting experiments that demonstrated the impact of varying degrees of pupil dilation on iris recognition performance.

Pupil dilation has been explored in other fields including mathematical biology and, more recently, computer graphics. Previous research has shown that in addition to changes in light intensity, pupil motion is also dictated by levels of consciousness, focal length, heart rate, respiration, emotional factors, disease, and drug use. However, physiological models were created by taking variations of illumination into account since there is a large amount of data supporting this factor. Modeling pupil dynamics began with empirical studies based on varying ambient light intensity. One of the famous models is due to Moon and Spencer [13] who proposed the following model based on experimental data:

$$D = 4.9 - 3 \tanh(0.4(\log_{10}(L_b - 0.5))), \quad (7.2)$$

where D is the pupil diameter (expressed in mm) and L_b is the background luminance (expressed in blondels). Later, it was realized that it is necessary to explore the pupil dynamics in order to understand the changes in pupil size given the varying levels of illumination. Therefore, more experimental studies have been conducted to explore this effect rather than simply relying on changes in pupil size. Ellis [4] performed experiments to determine the average latency as well as the maximum constriction and dilation velocities of the pupil for a given light intensity. As a result, he found that the latency τ (expressed in ms) can be described as

$$\tau(L_{cd}) = 445.7 - 22.9L_{cd} + 76.2L_{cd}^2, \quad (7.3)$$

and the constriction V_c and dilation V_d velocities (expressed in mm/s), respectively, can be described as

$$V_c(L_{cd}) = 0.15 + 2.0L_{cd} - 0.17L_{cd}^2 \quad (7.4)$$

and

$$V_d(L_{cd}) = 0.16 + 0.72L_{cd} - 0.07L_{cd}^2. \quad (7.5)$$

In Eqs. (7.3)–(7.5), the variable L_{cd} represents the intensity of the light measured in candelas/m². Furthermore, Link and Stark [11] empirically determined the

pupillary latency, the time delay between the light pulse reaching the retina and the beginning of the pupillary reaction, which yields the following equation:

$$\tau(R, L_{fL}) = 253 - 14 \ln(L_{fL}) + 70R - 29R \ln(L_{fL}), \quad (7.6)$$

where τ is the latency in ms, L_{fL} is the luminance in foot-Lambert, and R is the frequency in Hz.

From the empirical models came the physiologically based models derived from the physiological and anatomical observations without the reliance on experimental data. Usui and Stark [22] developed a parametric model of the iris to describe the pupil characteristics in response to light while using probability distribution functions to describe the random fluctuations. Later, Stark [20] concluded that the pupillary response can be viewed as a negative feedback dynamical system. Though this expression of the pupil dynamics is attractive, it was Longtin et al. [12] who were the first to express the dynamics of the pupil light reflex as a nonlinear delay differential equation given by

$$\frac{dg}{dA} \frac{dA}{dt} + \alpha g(A) = \gamma \ln \left[\frac{\phi(t - \tau)}{\bar{\phi}} \right], \quad (7.7)$$

where A is the pupil area, ϕ is the retinal light flux, and $\bar{\phi}$ is the light level when there is no pupillary response. Furthermore, τ is the time delay due to retinal processing, which includes responses from the midbrain nuclei, and $\gamma > 0$ is an arbitrary constant. In the derivation of Eq. (7.7), $g(A)$ takes into account the inverse relationship between the pupil area and iris muscle activity. The direct relationship is given by

$$A = f(x), \quad (7.8)$$

where $f(x)$ represents the activity of the iris muscle. Longtin defines this activity in terms of the Hill function expressed as

$$f(x) = \frac{\beta_1 \theta^n}{\theta^n + x^n} + \beta_2, \quad (7.9)$$

where β_1 is the minimum pupil area, $\beta_1 + \beta_2$ is the maximum pupil area, and θ and n are positive constants. In his work, Longtin chose the Hill function because it best represents the fact that the pupil area is positive and has finite limits while accounting for the elastic properties of the iris muscle. In computer graphics, Pamplona [15] recently looked at Longtin's equation of the pupil light reflex in order to simulate the varying degrees of pupil dilation as well as iridal pattern deformation in animated human characters in order to enhance facial animation. In his work, Pamplona notes that though Longtin's model is well cited, there are ambiguities in estimating the various parameters ($\alpha, \beta_1, \beta_2, \gamma, \theta, n, \bar{\phi}$). The selection of these parameters are important to guarantee convergence as well as ensure that the simulations are as realistic as possible. By using the steady-state formulation

of Eq. (7.7), making the comparison with the Moon and Spencer model, as well as making the appropriate conversions to the dynamic case, yields

$$\frac{dM}{dD} \frac{dD}{dt} + 2.3026M(D) = 0.45 \ln \left[\frac{\phi(t - \tau)}{\bar{\phi}} \right], \quad (7.10)$$

where D is the pupil diameter and the function $M(D) = \tanh^{-1} \left(\frac{D-4.9}{3} \right)$.

Although pupil dynamics has been studied in different communities outside of biometrics, modeling iridal dynamics has received limited attention. It appears that Rohen [18] was the first researcher to study the collagen form of the iris. In his work, he proposed a model for the collagen fibers assuming that they are arranged in a series of parallel arcs, connecting the iris root to the pupil border in both clockwise and counterclockwise fashions. Later, Wyatt [24] used Rohen's formulation and created a mathematical expression for iris deformation as the sum of a linear stretch and nonlinear deviation in order to demonstrate the wear and tear of the collagen fibers in the iris region. Wei et al. [23] used pattern recognition approaches and concluded that the nonlinear deviation is Gaussian. Finally, Pamplona used an imaged-based model for iridal deformation by tracking points on the pupil border as well as various feature points throughout the iris region. Based on this, he calculated the distance from the pupil center, the distance of the points from the pupil border, and the percentage of the points in the iris region with respect to the pupil diameter. With these results, he concluded that approximating the iridal deformation as a linear function is good enough to generate a realistic simulation.

7.1.2 *Our Approach and Motivation*

One of our main motivations for this work comes from Hollingsworth et al. [8]. Their results show that there is a need to explore the possibility of correcting the various degrees of pupil dilation in order to improve iris recognition. Our other motivation comes from the observation that the impact of pupil dilation on the dynamics of the iris is not well studied in the literature. Though Thornton [21] and Wei [23] took dilation into account in their corrections, it appears that they did not consider the varying degrees of pupil dilation. Furthermore, both Thornton and Wei used pattern recognition approaches, which are limiting because they depend on a particular dataset and are not necessarily based on the physiology of the iris. We also found Pamplona's image-based model for iridal deformation to be incomplete since it cannot account for the nonlinearity in iris deformation. Having a realistic simulation might prove beneficial in computer graphics; however, the desire in biometrics and other biologically related fields is to be physically accurate as possible. Finally, Wyatt is also aware that his mathematical model is limiting because he stresses that his work is only a "meshwork skeleton" and does not form the basis for both linear and nonlinear iridal stretch, which comes from the variations

in the elastic properties of the iris as well as the iris musculature. Thus, the work of Wei et al. is also limiting in this regard since they apply the “meshwork skeleton” to determine that the nonlinear stretch is Gaussian.

Here, we present a novel approach to mathematically model iris deformation by using what we know about biomechanics – the application of continuum mechanics to biological medium [5, 6]. With biomechanics, we can predict the nonlinear dynamics of the iris while considering the varying degrees of pupil dilation. Such an approach removes the reliance on pattern recognition principles alone and incorporates the changes in the elastic properties of the iris while considering the muscle activity. After formulating our model, we will perform a mathematical analysis on it within the confines of the annular region of the iris. Our mathematical analysis is done in two parts: We first look at the entire orthotropic case. Next, we consider the assumption that the iris region is isotropic. Though we adopt the assumption by Lei et al. [10] that the iris region is orthotropic, exploring the isotropic case could prove beneficial because we do not have to be concerned about the variation of the material properties of the iris while gaining some insight into the orthotropic case.

7.2 Mathematical Formulation

We begin our mathematical formulation by viewing the iris region as a thin cylindrical shell, where the thickness in the z dimension is much smaller compared to that of the r and θ directions. Hence, we can look at the region in terms of a thin plate where the loads are applied uniformly over the z dimension. With this thin plate formulation, we can also make the assumption that the normal stress in the z direction, σ_z , as well as the shear stresses, τ_{rz} and $\tau_{\theta z}$, is negligible. This simplification reduces the dynamics of the iris region to the two-dimensional r - θ plane. Therefore, the Cauchy-Euler equations become [14]

$$\varepsilon_r = \frac{\partial u_r}{\partial r} - \frac{1}{2} \left\{ \left(\frac{\partial u_r}{\partial r} \right)^2 + \left(\frac{\partial u_\theta}{\partial r} \right)^2 \right\} \quad (7.11)$$

$$\varepsilon_\theta = \frac{1}{r} \frac{\partial u_\theta}{\partial \theta} + \frac{u_r}{r} - \frac{1}{2r^2} \left\{ \left(\frac{\partial u_\theta}{\partial \theta} \right)^2 + \left(\frac{\partial u_r}{\partial \theta} \right)^2 - 2 \frac{\partial u_r}{\partial \theta} u_\theta + 2 \frac{\partial u_\theta}{\partial \theta} u_r + u_r^2 + u_\theta^2 \right\} \quad (7.12)$$

$$\gamma_{r\theta} = \frac{1}{r} \frac{\partial u_r}{\partial \theta} + \frac{\partial u_\theta}{\partial r} - \frac{u_\theta}{r} - \frac{1}{r} \left\{ \left(\frac{\partial u_r}{\partial \theta} \right) \left(\frac{\partial u_r}{\partial r} \right) + \left(\frac{\partial u_\theta}{\partial r} \right) \left(\frac{\partial u_\theta}{\partial \theta} \right) + \frac{\partial u_\theta}{\partial r} u_r - \frac{\partial u_r}{\partial r} u_\theta \right\}, \quad (7.13)$$

where ε_r and ε_θ , in (7.11) and (7.12), are the normal strains and $\gamma_{r\theta}$, in (7.13), is the shear strain. Also, u_r and u_θ are the radial and azimuthal components of the displacement vector in the annular region. Consequently, the equilibrium conditions are expressed as [14]

$$\begin{aligned}\frac{\partial \sigma_r}{\partial r} + \frac{1}{r} \frac{\partial \tau_{r\theta}}{\partial \theta} + \frac{(\sigma_r - \sigma_\theta)}{r} &= 0 \\ \frac{\partial \tau_{r\theta}}{\partial r} + \frac{2}{r} \tau_{r\theta} + \frac{1}{r} \frac{\partial \sigma_\theta}{\partial \theta} &= 0,\end{aligned}\quad (7.14)$$

where σ_r and σ_θ are the normal stresses and $\tau_{r\theta}$ is the shear stress. Another observation is that the iris muscles (both the sphincter and dilator muscles) are equally distributed throughout the annular region. Furthermore, the pupillary response causes these muscles to produce an axisymmetric load throughout the iris region. Therefore, mathematically, we can also make the assumptions that $\frac{\partial}{\partial \theta}$ and u_θ are negligible, and consequently, it follows that the shear stress $\tau_{r\theta}$ is also negligible [2]. We next consider the fact that soft biological tissues, like the iris, experience finite (or large) strain. This affects the Cauchy relationships because the relationship between the strain components ε_r and ε_θ and the displacement are completely nonlinear. With these assumptions, the Cauchy-Euler equations become

$$\begin{aligned}\varepsilon_r &= u' - \frac{1}{2}(u')^2 \\ \varepsilon_\theta &= \frac{u}{r} - \frac{1}{2}\left(\frac{u}{r}\right)^2,\end{aligned}\quad (7.15)$$

and the equilibrium condition reduces to

$$\frac{d\sigma_r}{dr} + \frac{\sigma_r - \sigma_\theta}{r} = 0. \quad (7.16)$$

In Eq. (7.15), we have dropped the subscripts because our problem now simplifies to only examining variations in the radial displacement u_r . We also use the primes ($'$) to denote ordinary differentiation with respect to r because the displacement now only depends on the radial component (i.e., $u = u(r)$). Adopting Lei et al.'s assumption that the iris material is orthotropic, the relationship between the stress vector $\vec{\sigma} = \langle \sigma_r, \sigma_\theta \rangle$ and the strain vector $\vec{\varepsilon} = \langle \varepsilon_r, \varepsilon_\theta \rangle$ can be expressed in the following manner:

$$\begin{aligned}\varepsilon_r &= \frac{\sigma_r}{E_r} - \frac{\nu_{r\theta}}{E_\theta} \sigma_\theta \\ \varepsilon_\theta &= -\frac{\nu_{\theta r}}{E_r} \sigma_r + \frac{\sigma_\theta}{E_\theta},\end{aligned}\quad (7.17)$$

where E_r and E_θ are the Young's moduli (i.e., the material properties of the iris region), $\nu_{r\theta}$ is the Poisson's ratio in the azimuthal direction, and $\nu_{\theta r}$ is the Poisson's

ratio in the radial direction. It needs to be noted that for orthotropic formulation $v_{r\theta} \neq v_{\theta r}$; however, due to the symmetry with the stress and strain tensors, the following relationship must hold:

$$\frac{v_{\theta r}}{E_r} = \frac{v_{r\theta}}{E_\theta}. \quad (7.18)$$

Since the Poisson's ratio $v_{\theta r}$ is known [10], we use the relationship given from (7.18) to simplify Eq. (7.17) to be

$$\begin{aligned} \varepsilon_r &= \frac{\sigma_r}{E_r} - \frac{v_{\theta r}}{E_r} \sigma_\theta \\ \varepsilon_\theta &= -\frac{v_{\theta r}}{E_r} \sigma_r + \frac{\sigma_\theta}{E_\theta}. \end{aligned} \quad (7.19)$$

Solving for the stress terms σ_r and σ_θ results in

$$\begin{aligned} \sigma_r &= \frac{E_r}{(1 - \zeta v^2)} (\varepsilon_r + \zeta v \varepsilon_\theta) \\ \sigma_\theta &= \frac{E_\theta}{(1 - \zeta v^2)} (v \varepsilon_r + \varepsilon_\theta), \end{aligned} \quad (7.20)$$

where $v = v_{\theta r}$ and $\zeta = \frac{E_\theta}{E_r}$. Substituting Eq. (7.15) into Eqs. (7.17) and (7.16) yields the following differential equation:

$$u'' + \frac{u'}{r} - \frac{\zeta u}{r^2} - \frac{(1 - v\zeta)}{2r} (u')^2 - \frac{(v-1)\zeta}{2r} \left(\frac{u}{r}\right)^2 - \frac{1}{2} \frac{d}{dr} (u')^2 - \frac{v\zeta}{2} \frac{d}{dr} \left(\frac{u}{r}\right)^2 = 0, \quad (7.21)$$

where ζ is given by

$$\zeta = \frac{E_\theta}{E_r} \quad (7.22)$$

and E_r and E_θ are the radial and azimuthal Young's moduli, respectively. The boundary conditions in Eq. (7.21) on the annular domain $r \in (r_1, r_2)$ are

$$u(r_1) = \mu_1, \mu_1 > 0 \quad (7.23)$$

and

$$u(r_2) = 0. \quad (7.24)$$

Examining Eq. (7.21) shows that iris deformation is nonlinear. However, we will like to perform a theoretical analysis of the displacement $u(r)$ inside of the annular region $r \in (r_1, r_2)$ by examining (7.21) given the boundary conditions (7.23) and (7.24). Providing such a mathematical analysis results in the overall theoretical picture of the iridal dynamics.

7.2.1 Mathematical Analysis: Orthotropic Deformation

In order to begin our mathematical analysis of Eq. (7.21) with the boundary conditions (7.23) and (7.24), we first perform the following normalization:

$$\hat{u} = u/\mu_1, \quad \hat{r} = r/r_2. \quad (7.25)$$

Thus, Eq. (7.21) becomes

$$\hat{u}'' + \frac{1}{\hat{r}}\hat{u}' - \frac{\zeta}{\hat{r}^2}\hat{u} + \lambda \left[\frac{\nu\zeta - 1}{2\hat{r}}(\hat{u}')^2 + \frac{(1-\nu)\zeta}{2\hat{r}^3}\hat{u}^2 - \frac{1}{2} \frac{d}{d\hat{r}} \left\{ (\hat{u}')^2 + \nu\zeta \left(\frac{\hat{u}}{\hat{r}} \right)^2 \right\} \right] = 0, \quad (7.26)$$

and the normalized boundary conditions become

$$\hat{u}(\eta) = 1, \quad \hat{u}(1) = 0, \quad (7.27)$$

where

$$\lambda = \frac{\mu_1}{r_2}, \quad \eta = \frac{r_1}{r_2} \quad (7.28)$$

and the primes (\prime) denote the differentiation with respect to \hat{r} . It is important to note that the existence of the solution to Eq. (7.21) with boundary conditions defined in (7.23) and (7.24) is to be expected from the rational derivation of the equations in solid mechanics [1]. However, the question of uniqueness needs to be explored in more detail [16] so that we can apply a maximum principle for nonlinear differential operators. The idea is based on the maximum principle for linear operators of which the linearized problem ($\lambda = 0$) is an example. For nonlinear operators, we take advantage of the following theorem [16] to show uniqueness:

Theorem 7.1. *Let $u(r)$ be a solution of the boundary value problem*

$$u'' + H(r, u, u') = 0, \quad a < r < b, \quad (7.29)$$

$$\left. \begin{aligned} u(a) &= \gamma_1, \\ u(b) &= \gamma_2. \end{aligned} \right\} \quad (7.30)$$

Suppose that H , $\partial H/\partial u$, and $\partial H/\partial u'$ are continuous and that $\partial H/\partial u \leq 0$. Then, a solution to (7.29) which also satisfies (7.30) is unique.

In order to take advantage of Theorem 7.1, we expand Eq. (7.26) to obtain the following:

$$(1 - \lambda\hat{u}')\hat{u}'' + \frac{1}{\hat{r}}\hat{u}' - \frac{\zeta}{\hat{r}^2}\hat{u} + \lambda \left[\frac{\nu\zeta - 1}{2\hat{r}}(\hat{u}')^2 + \frac{(1+\nu)\zeta}{2\hat{r}^3}\hat{u}^2 - \nu\zeta \frac{\hat{u}'\hat{u}}{\hat{r}^2} \right] = 0, \quad (7.31)$$

which is in the same form as in (7.29) with H being defined as

$$H = \frac{1}{(1 - \lambda \hat{u}')} \left\{ \frac{1}{\hat{r}} \hat{u}' - \frac{\zeta}{\hat{r}^2} \hat{u} + \lambda \left[\frac{\nu \zeta - 1}{2\hat{r}} (\hat{u}')^2 + \frac{(1 + \nu) \zeta}{2\hat{r}^3} \hat{u}^2 - \nu \zeta \frac{\hat{u}' \hat{u}}{\hat{r}^2} \right] \right\} = 0 \quad (7.32)$$

with the boundary conditions defined in (7.27). Looking at the annular domain of the iris region $r \in (r_1, r_2)$, it is noted that as the pupil dilates, the iris tissue is compressed due to the iris musculature. Therefore, we can physically expect that $\hat{u}' \leq 0$, making the expression $1 - \lambda \hat{u}' > 0$. Furthermore, since the rest of the expression inside of the braces in Eq. (7.32) is in polynomial form, we can deduce that H and its derivatives are continuous.

Next, we compute $\partial H / \partial \hat{u}$ and get the following expression:

$$\frac{\partial H}{\partial \hat{u}} = \frac{\zeta}{(1 - \lambda \hat{u}')} \left\{ -\frac{1}{\hat{r}^2} + \lambda \left[(1 + \nu) \frac{\hat{u}}{\hat{r}^3} - \nu \frac{\hat{u}'}{\hat{r}^2} \right] \right\}. \quad (7.33)$$

Looking at Eq. (7.33), we see that since $\zeta / (1 - \lambda \hat{u}') > 0$, the condition $\partial H / \partial u \leq 0$ is satisfied when

$$\lambda \left[(1 + \nu) \frac{\hat{u}}{\hat{r}^3} - \nu \frac{\hat{u}'}{\hat{r}^2} \right] \leq \frac{1}{\hat{r}^2}. \quad (7.34)$$

Now, we have $\eta \leq \hat{r} < 1$, and since $\hat{u}' \leq 0$ and $\hat{u} \leq 1$, so condition (7.34) is satisfied as long as

$$\lambda \left[(1 + \nu) \frac{1}{\eta} + \nu |\hat{u}'| \right] \leq 1. \quad (7.35)$$

Though the conditions of the theorem are sufficient for uniqueness and are seemingly satisfied for the smaller values of λ , it does indicate the prominent role played by the nonlinearity. In Eq. (7.26), we perform a regular perturbation expansion to describe the behavior of the nonlinearity. For small values of λ , we can perform the perturbation expansion of \hat{u} as follows:

$$\hat{u}(\hat{r}; \lambda) = \hat{u}_0(\hat{r}) + \lambda \hat{u}_1(\hat{r}) + \lambda^2 \hat{u}_2(\hat{r}) + \dots \quad (7.36)$$

with simple boundary conditions $\hat{u}_0(\eta) = 1$, $\hat{u}_0(1) = 0$, and $\hat{u}_1(\eta) = \hat{u}_1(1) = \hat{u}_2(\eta) = \hat{u}_2(1) = \dots = 0$. Thus, \hat{u}_1 is given by

$$\hat{u}_1'' + \frac{1}{\hat{r}} \hat{u}_1' - \frac{\zeta}{\hat{r}^2} \hat{u}_1 = \frac{(1 - \nu \zeta)}{2\hat{r}} (\hat{u}_0')^2 - \frac{(1 + \nu) \zeta}{2\hat{r}^3} \hat{u}_0^2 + \nu \zeta \frac{\hat{u}_0' \hat{u}_0}{\hat{r}^2} := f_0(\hat{r}). \quad (7.37)$$

Each term can actually be represented as an integral such as

$$\hat{u}_n(\hat{r}) = \int_{\eta}^1 G(\hat{r}, \rho) f_{n-1}(\rho) d\rho, \quad n = 1, 2, \dots, \quad (7.38)$$

where $G(\hat{r}, \rho)$ is the Green's function satisfying

$$(\hat{r}G_{\hat{r}})_{\hat{r}} - \frac{\zeta}{\hat{r}}G = \delta(\hat{r} - \rho), \quad G(\eta, \rho) = G(1, \rho) = 0. \quad (7.39)$$

Therefore, in the term of order 1, we have the following expression:

$$\hat{u}_1(\hat{r}) = \int_{\eta}^1 G(\hat{r}, \rho) \left[\frac{(1 - \nu\zeta)}{2\rho} (\hat{u}'_0(\rho))^2 - \frac{(\nu + 1)\zeta}{2\rho^3} \hat{u}_0^2(\rho) + \nu\zeta \frac{\hat{u}_0(\rho)\hat{u}'_0(\rho)}{\rho^2} \right] d\rho. \quad (7.40)$$

Likewise, successive terms of (7.36) may be determined in this regular perturbation expansion.

7.2.2 Mathematical Analysis: Isotropic Deformation

The isotropic case, though different physically, is really only a special case analytically, by taking $\zeta = 1$ in Eq. (7.26). The uniqueness theorem therefore holds in the same way as before. The perturbation solution is a little simpler to write down. The appropriate Green's function in this case is explicitly

$$G(\hat{r}, \rho) = \begin{cases} [3(1 - \eta^2)]^{-1} \left(\hat{r} - \frac{\eta^2}{\hat{r}} \right) \left(\rho - \frac{1}{\rho} \right), & \hat{r} < \rho \\ [3(1 - \eta^2)]^{-1} \left(\rho - \frac{\eta^2}{\rho} \right) \left(\hat{r} - \frac{1}{\hat{r}} \right), & \hat{r} > \rho \end{cases}. \quad (7.41)$$

The solution of the perturbation expansion proceeds with

$$\hat{u}_0(\hat{r}) = \frac{\eta}{(\eta^2 - 1)} \left(\hat{r} - \frac{1}{\hat{r}} \right), \quad \hat{u}'_0(\hat{r}) = \frac{\eta}{\eta^2 - 1} \left(\frac{1}{\hat{r}^2} + 1 \right), \quad (7.42)$$

and the equation for \hat{u}_1 is the expression (7.40) with $\zeta = 1$.

7.3 Numerical Results and Simulation

The numerical simulation of the iris deformation for both the orthotropic and isotropic cases was done using the finite element method (FEM) [7, 17]. The simulation could be implemented via other numerical methods; however, we found that the FEM offers greater flexibility with regard to mesh generation. The FEM implementation consisted of converting Eq.(7.21) into a variational form and discretizing our annular domain into 100 vertices that are uniformly spaced from the inner radius r_1 to the outer radius r_2 (r_2 is set to 6 mm to represent the limbus boundary). Because the average pupil size at constriction is 3 mm in diameter, we

initialize the inner radius to 1.5 mm and solve numerically for $u(r)$. As we vary the inner radius, we continue to numerically solve for $u(r)$, noting the level of deformation of the iris tissue. Furthermore, due to the fact that the average pupil size at dilation is approximately 9.6 mm diameter (or 4.8 mm in radius), we use that as our stopping point. Each time we solve for $u(r)$, we compare that to the linear assumption. To visualize the deformation between our model and the linear approximation, we begin by adding concentric circles that are evenly spaced in the annular region. Next, once numerically solving for $u(r)$, we move the ring located at the starting radius r to the final position located at radius $u(r) + r$.

We adopt Pamplona's metrics [15] as follows: the distance from the interior point to the center of the annular region, the distance from the interior point to the pupil boundary, and the ratio of the distance between the tracked interior point to the pupil border and the local width of the iridal disk (otherwise known as the invariance). Doing so enables us to further analyze the nonlinear effects of the interior points of the annular region for both the orthotropic and isotropic cases.

7.3.1 Numerical Results: Orthotropic Deformation

For the numerical simulation of the orthotropic deformation, we adopt Lei et al.'s [10] formulation and assumptions for determining the material parameters E_r and E_θ in order to compute ζ for the various degrees of dilation. Because it is found that the Poisson's ratio for the iris $\nu \in (0.45, 0.5)$, we chose $\nu = 0.49$ for our simulation. Figures 7.3–7.5 show the results for each simulation where the inner radius was changed to 2, 3, and 4.5 mm to depict the various levels of iris deformation.

In Fig. 7.3 top left, we compare the solution $u(r)$ to Eq. (7.21) with the inner radius expanded to 2 mm against the previously assumed linear solution. The difference between the nonlinear solution given by (7.21) and the linear approximation is rather apparent in this figure. However, in the top right and bottom of Fig. 7.3, we see a minimum difference between the two methods in the effects of deformation when looking at the final positions $r + u(r)$. Furthermore, our model also shows that the compression of the iris tissue occurs in the innermost regions of the annular region, which is where most of the deformation occurs.

In Fig. 7.4, we see that when the inner radius is moved to 3 mm, there is the apparent difference between the solution $u(r)$ to (7.21) and the linear approximation. However, in this case, we now find that the differences between the linear and nonlinear solutions are more significant. Furthermore, the difference in the visualizations also become more discernible.

In Fig. 7.5, we see the results after the inner radius expands to 4.5 mm. For this case, while the iris boundary is now at 4.5 mm, elements originally at 2.5 mm move to approximately 4.4 mm behind the boundary. We conclude that this behavior is analogous to simulating folds in the iris tissue. It is important to note that these effects occur only in the radial direction due to the fact that we make the assumption that the dynamics in the z direction are negligible.

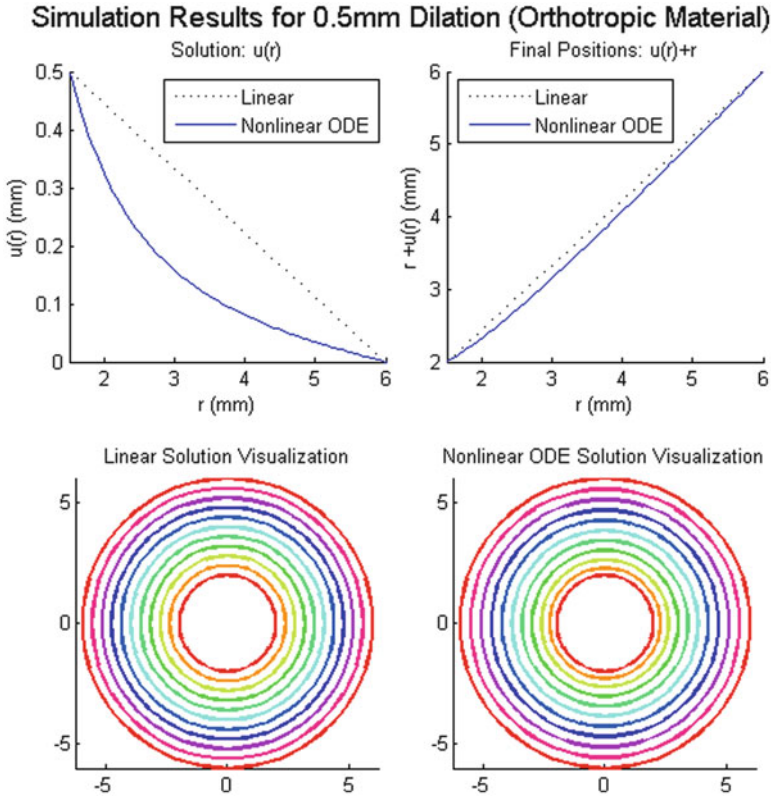


Fig. 7.3 Results in simulating a pupil with orthotropic material properties dilating from 1.5 to 2.0 mm. The solution $u(r)$ of Eq. (7.21) is nonlinear (*top left*); however, the difference in deformation between $u(r)$ and the linear approximation does not appear significant (*top right*). This is further shown in the visualizations (*bottom*)

We then ran the simulation for many more degrees of dilatation, tracked several locations on the iris through time, applied Pamplona's metrics, and plotted the results shown in Fig. 7.6. Each line in these graphs represents a single point on the iris that moves as the pupil dilates. The top-left graph represents the invariance, or the percentage that a given point lies within the annular region. So, for example, we see that a point which was originally 10% in the iris steadily moves inward as the pupil dilates, again suggesting that the inner regions are becoming compressed. In contrast, the point originally 90% into the iris barely moves at all as the pupil dilates, generating a straight line. We can see that as the pupil diameter exceeds 6.5 mm, the lines originating at 10–30% begin to intersect the 0 line, demonstrating the effects of folding. The top-right graph shows the final distance from a point to the pupil border after dilation. As expected, points closer to the outer edge of the iris are decreasing nearly linearly as pupil diameter increases, whereas the points closer to the pupil deform nonlinearly and again cross the horizontal axis. Finally, the bottom-center graph shows the distance from the center of the pupil to points

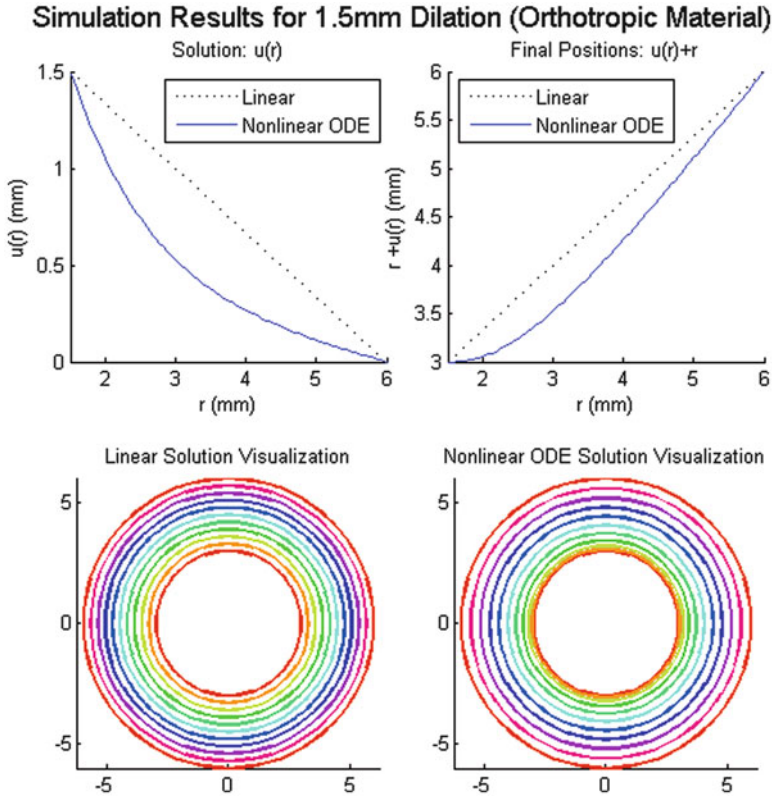


Fig. 7.4 Results in simulating a pupil with orthotropic material properties dilating from 1.5 to 3.0 mm. The solution $u(r)$, to (7.21), continues to remain nonlinear (top left); however, the differences between $u(r)$ and the linear approximation become significant (top right). This is further shown in the visualizations (bottom)

within the iris. We once again see that near the outer iris edge, points are moving very little compared to those near the pupil border.

7.3.2 Numerical Results: Isotropic Deformation

Our approach to simulate the isotropic deformation was the same as that for the orthotropic case except that we set $\zeta = 1$ to account for the material parameters being the same in all directions. The Poisson’s ratio ν was chosen to be 0.49 for our simulations for the same reason as that of the orthotropic case. Similar to the orthotropic case, Figs. 7.7–7.9 show the results for each simulation where the inner radius was changed to 2, 3, and 4.5 mm, depicting the various levels of deformation in the annular region.

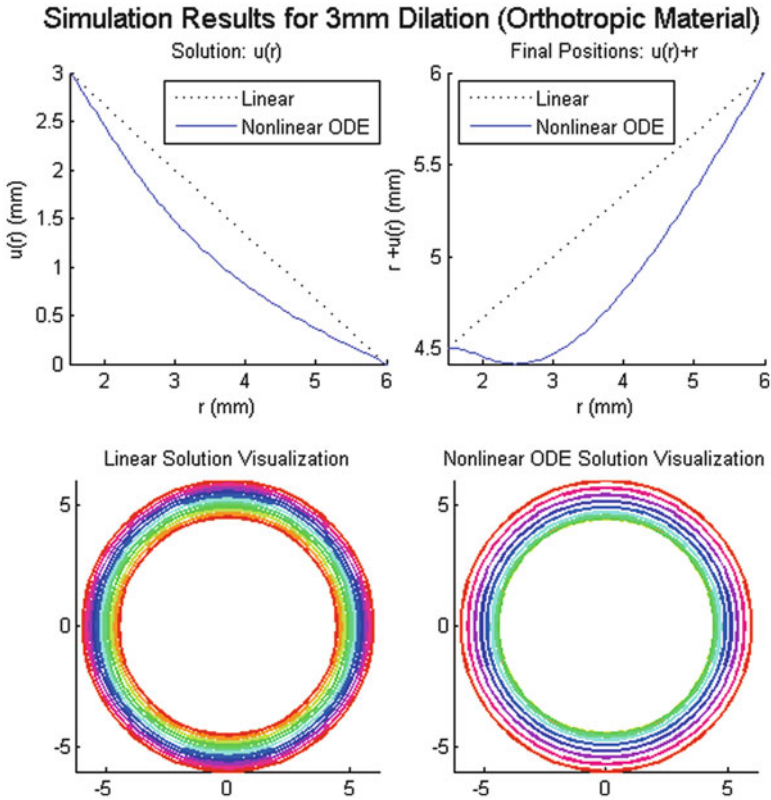


Fig. 7.5 Results in simulating a pupil with orthotropic material properties dilating from 1.5 to 4.5 mm. Note in the *top right*, the iris tissue moves backward, behind the pupil border, suggesting the need for three-dimension simulations

Analogous to the orthotropic case, Fig. 7.7 shows that for the case when the inner radius is expanded to 2 mm, the solution $u(r)$ is nonlinear compared to the linear approximation. However, just as in the orthotropic case, the distinctions between $u(r)$ and the linear approximation are not significant. Furthermore, compared to the orthotropic simulation, $u(r)$ does not appear to fall as rapidly at the beginning, and so we can conclude that the iris tissue is slightly less compressed in the inner regions.

We reach similar conclusions for the medium deformation case in Fig. 7.8. Again, we see that the nonlinear deformation is starting to be significantly different than the linear deformation.

Finally, in Fig. 7.9, we again see that the high degree of dilation causes the simulation to produce folding results. Just as in the orthotropic case, the reason for these effects is due to our assumption that the dynamics in the z direction are negligible.

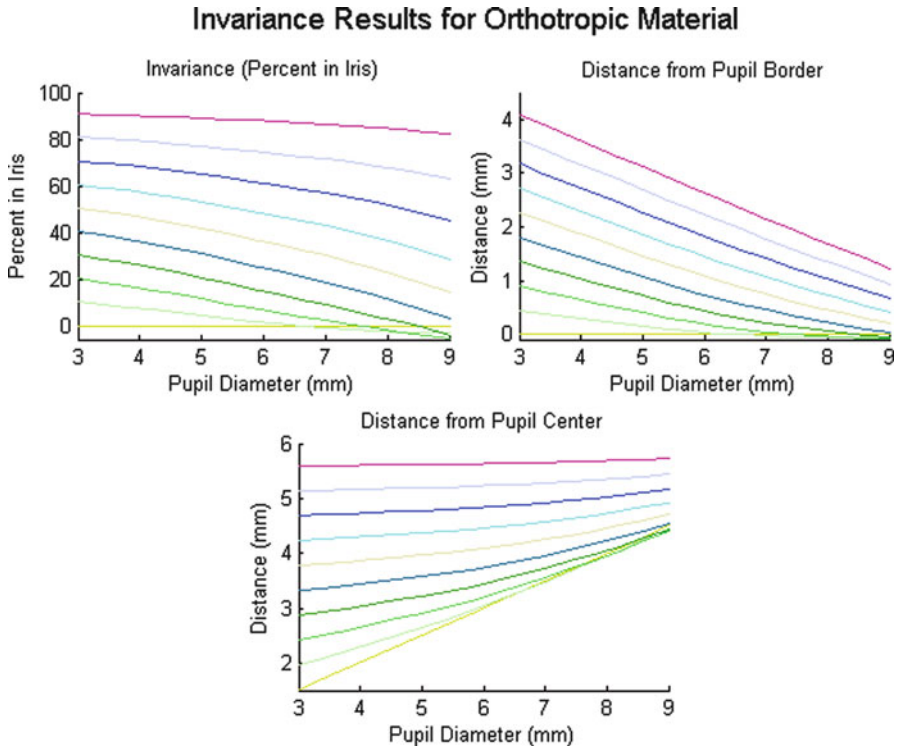


Fig. 7.6 Plot showing how the points interior to the iris move as the pupil dilates, given orthotropic material properties

As in the orthotropic case, we applied Pamplona’s metrics to track the interior points of the iris as the inner radius changes (due to pupil dilation) as shown in Fig. 7.10. The behavior is similar to the orthotropic case. However, it appears that the slopes in the invariance graph tend to be lower in this simulation. Furthermore, we see that the graphs exhibit a linear behavior when depicting the distance from the pupil border and the distance from the pupil center, respectively. Finally, looking at the invariance graph, our results show that folding occurs when the inner diameter is between 7.5 and 8 mm versus 6 and 6.5 mm in the orthotropic case.

7.4 Conclusions and Future Work

We used the knowledge of biomechanics in order to produce a theoretical two-dimensional mathematical model describing iris deformation. We believe that our model is an improvement from previous work because it takes into account the material properties of the iris as well as the cause of deformation due to the iris

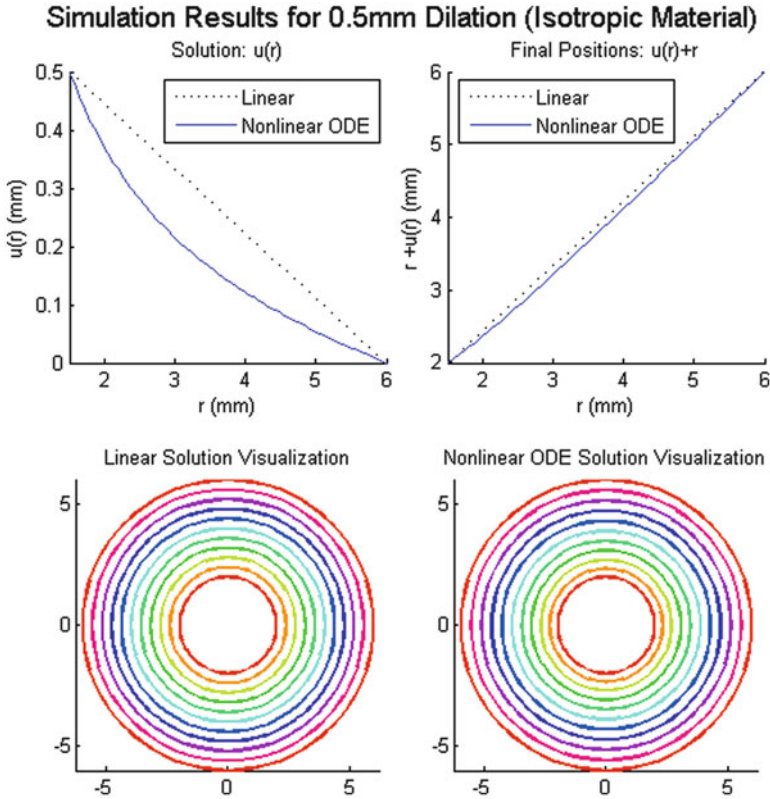


Fig. 7.7 Results in simulating a pupil with isotropic material properties dilating from 1.5 to 2 mm. Analogous to the orthotropic case, the nonlinear solution $u(r)$ to (7.21) with $\zeta = 1$ is nonlinear (*top left*); however, the difference in deformation between $u(r)$ and the linear approximation does not appear to be significant (*top right*). This is further shown in the visualizations of the iris deformation (*bottom*)

muscle activity. Also, by building our model, we were able to directly show that the deformation is nonlinear and that, therefore, there is more to the deformation than just a linear approximation. We were able to perform a mathematical analysis of our model via looking at the uniqueness conditions to show the behavior of $u(r)$ with various degrees of dilation. This analysis strengthens our model because we were able to relate the physiology to the mathematical theory. In our numerical experiments, we were able to visually compare the level of deformation against the linear assumption: this not only showed the level of nonlinearity but also the potential folds in the iris tissue. Furthermore, we were able to apply the metrics in Pamplona's [15] work to show that our model, as well as our approximation, behaves similarly to Pamplona's image model and is also consistent with Wyatt's results. However, through Pamplona's metrics, we were able to further show the nonlinearity as well as potential folds in the iris region. These folds are more

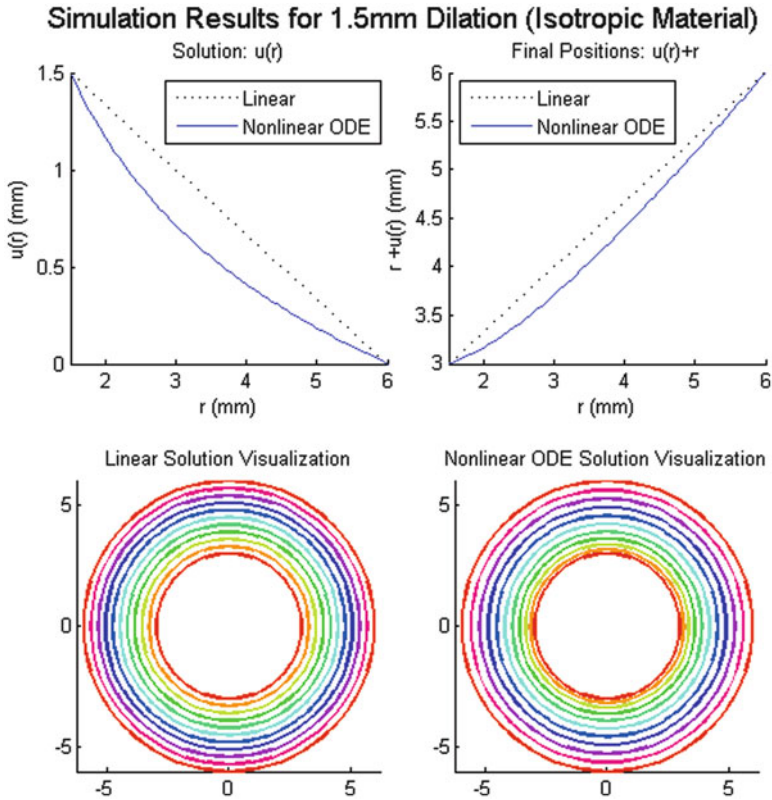


Fig. 7.8 Results in simulating a pupil with isotropic material properties dilating from 1.5 to 3 mm. Analogous to the orthotropic case, the nonlinear solution $u(r)$ to Eq. (7.21) with $\zeta = 1$ is nonlinear; however, the distinctions between $u(r)$ and the linear approximation become significant (top right). This is displayed further in the visualizations (bottom)

accurate in describing the iridal dynamics as the pupil dilates. Finally, through our mathematical analysis and numerical experiments, we were able to show the similarity in behavior between the isotropic and orthotropic case and can conclude that the isotropic formulation could serve as a practical approximation because the nonlinear dynamics are still preserved.

Our model can be useful to the biometrics and biologically related communities. In the biometrics community, this model can serve as an analytical description for completely describing the level of iris deformation. Thornton et al.'s [21] work showed through pattern recognition that if there is knowledge of the level of deformation, one should be able to correct its effects. Our model would be an improvement to their effort because it can be used to show the level of deformation in the iris region on both micro and macro levels. Our model can also serve to explore the possible enhancement of Daugman's rubber-sheet model, which uses

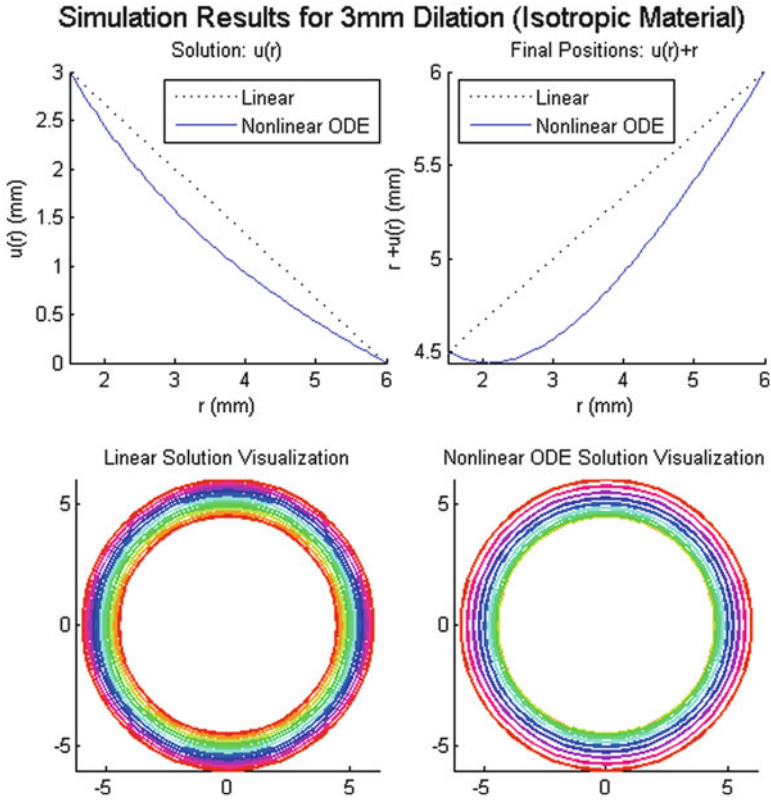


Fig. 7.9 Results in simulating a pupil with isotropic material properties dilating from 1.5 to 4.5 mm. Analogous to the orthotropic case, the iris tissue moves backward, behind the pupil border, suggesting the need for three-dimension simulations

a simple transformation to unwrap the iris while accounting for dilatation. Though Daugman’s model is used in commercial systems, the current challenge in iris recognition is accounting for variations in pupil dilatation. Therefore, our model can be applied to design an enhanced rubber-sheet model that counteracts iridal deformation. In biologically related communities, our model can generate interest for exploring other areas of iris deformation. Finally, our model can also be applied to the computer graphics arena to simulate realistic response in animated characters due to variations in illumination.

Though the numerical simulation of our model gives promising results, there are also a couple of areas for future work that need to be explored. First, there is a need to explore testing our model extensively on actual data to draw a more accurate comparison. Doing so will afford us the opportunity to explore the deformation pragmatically connecting our theoretical model to actual data. Pamplona [15] tested his model with actual data and concluded that the deformation is mostly linear when

Invariance Results for Isotropic Material




Fig. 7.10 Plot showing how the points interior to the iris move as the pupil dilates, given isotropic material properties

tracking certain feature points; however, our model shows that this conclusion is not necessarily complete. Second, for the extreme cases of dilation, we noticed the significant impact of the effects of folding. This is not completely shown in our two-dimensional model because we had assumed that the dynamics within the thickness of the iris is negligible. Therefore, there is a need to build a mathematical model that also takes into consideration the dynamics in the thickness of the iris as the pupil dilates. Doing so will enable us to have a more complete perspective of iris dynamics.

Acknowledgements This material is based upon work supported by the US Department of Homeland Security under Grant Award Number 2007-ST-104-000006. The views and conclusions contained in this chapter are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the US Department of Homeland Security. Arun Ross was supported by US National Science Foundation CAREER Grant No. IIS 0642554.

References

1. Antman, S.: *Nonlinear Problems of Elasticity*. Springer, New York (1994)
2. Carroll, W.F.: *A Primer for Finite Elements in Elastic Structures*. Wiley, New York (1999)
3. Daugman, J.G.: High confidence visual recognition of persons by a test of statistical independence. *IEEE Trans. Pattern Anal. Mach. Intell.* **15**(11), 1148–1161 (1993)
4. Ellis, C.J.: The pupillary light reflex in normal subjects. *Br. J. Ophthalmol.* **65**(11), 754–759 (1981)
5. Fung, Y.C.: *A First Course in Continuum Mechanics*. Prentice Hall, Englewood Cliffs (1994)
6. Fung, Y.C.: *Biomechanics*. Springer, New York (1997)
7. Henwood, D., Bonet, J.: *Finite Elements: A Gentle Introduction*. MacMillan, Houndmills (1996)
8. Hollingsworth, K., Bowyer, K.W., Flynn, P.J.: Pupil dilation degrades iris biometric performance. *Comput. Vis. Image Underst.* **113**(1), 150–157 (2009)
9. Jain, A.K., Flynn, P.J., Ross, A.: *Handbook of Biometrics*. Springer, Boston (2007)
10. Lei, Y., et al.: Experimental research on the mechanical properties of porcine iris. *Clin. Biomech.* **23**, 83–87 (2008)
11. Link, N., Stark, L.: Latency of the pupillary response. *IEEE Trans. Biomed. Eng.* **35**(3), 214–218 (1988)
12. Longtin, A., Milton, J.G.: Modelling autonomous oscillations in the human pupil light reflex using nonlinear delay-differential equations. *Bull. Math. Biol.* **51**(5), 605–624 (1989)
13. Moon, P., Spencer, D.: On the stiles-crawford effect. *J. Opt. Soc. Am.* **34**, 319–329 (1944)
14. Palazotto, A.: *Nonlinear Analysis of Shell Structures*. AIAA Education Services, Washington, DC (1992)
15. Pamplova, V.: Photorealistic models for pupil light reflex and iridal pattern deformation. *ACM Transactions on Graphics*, **28**(4) (2009)
16. Protter, H., Weinberger, H.F.: *Maximum Principles in Differential Equations*. Springer, New York (1984)
17. Reddy, J.N.: *Nonlinear Finite Element Analysis*. Oxford University Press (2004)
18. Rohen, H.: Der bau der regenbogenhaut beim menschen und einigen säugern. *Gegenbaur Morphol.* **91**, 140–181 (1951)
19. Ross, A.: Iris recognition: The path forward. *Computer* **43**, 30–35 (2010)
20. Stark, L.: The pupil as a paradigm for neurological control. *IEEE Trans. Biomed. Eng.* **31**, 1925–1939 (1984)
21. Thornton, J., Savvides, M., Kumar, B.V.V.: Bayesian approach to deformed pattern matching of iris images. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(4), 596–606 (2007)
22. Usui, S., Stark, L.: A model for nonlinear stochastic behavior of the pupil. *Biol. Cybern.* **45**(1), 13–21 (1982)
23. Wei, Z., Tan, T., Sun, Z.: Nonlinear iris deformation correction based on gaussian model. In: Lee, S.-W., Li, S.Z. (eds.) *Advances in Biometrics. Lecture Notes in Computer Science*, vol. 4642, pp. 780–789. Springer, Berlin/New York (2007)
24. Wyatt, H.J.: A ‘minimum-wear-and-tear’ meshwork for the iris. *Vis. Res.* **40**, 2167–2176 (2000)

Chapter 8

Iris Recognition in the Visible Wavelength

Hugo Proença

Abstract The human iris supports contactless data acquisition and can be imaged covertly. Thus, at least theoretically, the subsequent biometric recognition procedure can be performed without subjects' knowledge. The feasibility of this type of recognition has received increasing attention and is of particular interest for forensic and security purposes, such as the pursuit of criminals and terrorists and the search for missing children. Among others, one active research area sought to use visible wavelength (VW) light imagery to acquire data at significantly larger distances than usual and on moving subjects, which is a difficult task because this *real-world* data is notoriously different from the one used in the near-infrared (NIR) setup. This chapter addresses the feasibility of performing reliable biometric recognition using VW data acquired under dynamic lighting conditions and unconstrained acquisition protocols: with subjects at large distances (between 4 and 8 m) and on-the-move.

8.1 Introduction

The iris is known as one of the most valuable traits for the automatic identification of human beings, and growing attention has been paid to the development of this technology [3]. A number of reasons justify this interest: (1) it is a naturally protected internal organ that is visible from the exterior; (2) it has a near circular and planar shape that turns easier its segmentation and parameterization; and (3) its texture has a predominantly randotypic chaotic appearance that is stable over lifetime. The accuracy of the deployed iris recognition systems is remarkable: a study of 200 billion cross-comparisons conducted by Daugman [10] reported false acceptance rates of order 10^{-6} with false rejections of 1%. Other independent evaluations ([15] and [16]) confirmed these results.

H. Proença (✉)
University of Beira Interior, Covilhã, Portugal
e-mail: hugomcp@di.ubi.pt




Fig. 8.1 Comparison between (a) the quality of iris biometric images acquired in highly constrained conditions in the near-infrared wavelength (WVU database [23]) and (b) images acquired in the visible wavelength in unconstrained imaging conditions, acquired at a distance and on-the-move (UBIRIS.v2 database [26])

However, currently deployed systems—that are based on Daugman’s pioneering approach [7]—impose significant constraints on the subjects and acquisition environment, demanding that subjects stand relatively close and looking straight to the camera for several seconds while their eyes are illuminated by a near-infrared (NIR) source, enabling the capture of good quality data.

8.1.1 Why Use Visible Light?

Current systems require high illumination levels, sufficient to maximize the signal-to-noise ratio in the sensor and to capture images of the discriminating iris features with sufficient contrast. However, if similar processes were used to acquire iris images from a distance, acceptable depth-of-field values would demand significantly higher f-numbers for the optical system, corresponding directly (squared) with the amount of light required for the process. Similarly, the motion factor will demand very short exposure times, which again will require too high levels of light. The American and European standards councils ([1] and [6]) proposed safe irradiance limits for NIR illumination of near 10 mW/cm^2 . In addition to other factors that determine imaging system safety (blue light, nonreciprocity, and wavelength dependence), these limits should be taken into account, as excessively strong illumination can cause permanent eye damage. The NIR wavelength is particularly hazardous, because the eye does not instinctively respond with its natural mechanisms (aversion, blinking and pupil contraction). However, the use of visible light and unconstrained imaging setups can severely degrade the quality of the captured data (Fig. 8.1), increasing the challenges in performing reliable recognition.

The pigmentation of the human iris consists mainly of two molecules: brown-black eumelanin (over 90%) and yellow-reddish pheomelanin [24]. Eumelanin has most of its radiative fluorescence under the VW, which—if properly

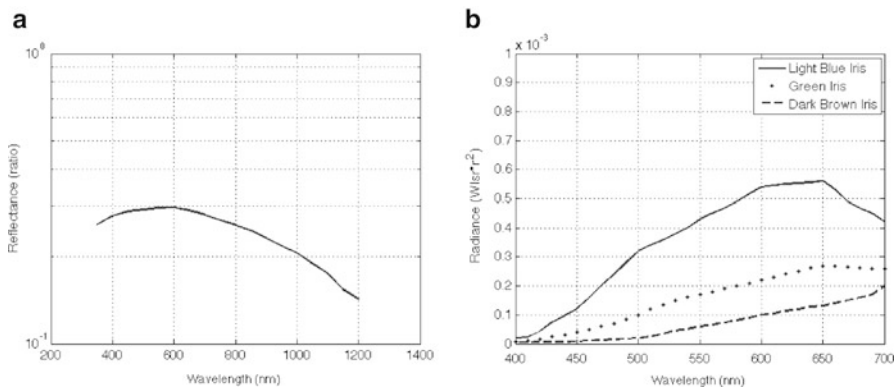


Fig. 8.2 Spectral reflectance and radiance of the iris and the sclera in respect of the wavelength. Spectral reflectance of the human sclera [20]. Spectral radiance of the human iris according to the levels of iris pigmentation [14]

imaged—enables the capture of a much higher level of detail, but also of many more noisy artifacts, including specular and diffuse reflections and shadows. Also, the spectral reflectance of the sclera is significantly higher in the VW than in the NIR (Fig. 8.2a), and the spectral radiance of the iris in respect of the levels of its pigmentation varies much more significantly in the VW than in the NIR (Fig. 8.2b). These optical properties are the biological roots behind the higher heterogeneity of the VW iris images, when compared with the traditional NIR data. Also, the types and number of noisy artifacts likely to appear in VW and NIR data are notoriously different, which justify the need to specialized recognition strategies.

The feasibility of the unconstrained VW iris recognition remains controversial, and several skepticisms remain, especially for high pigmented irises that constitute the majority of the world’s population. This chapter has two major parts: (1) to describe a data set of VW iris images captured in unconstrained conditions and (2) to discuss the major issues for VW iris recognition, namely describing the factors that likely degrade performance and giving results about the specificity and sensitivity that pattern recognition (PR) systems achieve in such challenging conditions.

8.2 Less Constrained Iris Recognition

The term *constraint* refers to one of the factors that currently deployed systems impose, in order to perform recognition with enough confidence: subjects distance, motion and gaze direction, and lighting conditions of the environment. These constraints motivate growing research efforts and became the focus of many recent proposals, among which the “Iris-on-the-move” project [17] should be highlighted: it is a major example of engineering an image acquisition system to make the recognition process less intrusive for subjects. The goal is to acquire NIR close-up

iris images as a subject walks at normal speed through an access control point. *Honeywell Technologies* applied for a patent [13] on a very similar system, which was also able to recognize irises at a distance. Previously, Fancourt et al. [11] concluded that it is possible to acquire sufficiently high-quality images at a distance of up to ten meters. Narayanswamy and Silveira [19] used a wavefront-coded optic to deliberately blur images in such a way that they do not change over a large depth of field. Removing the blur with digital image processing techniques makes the trade-off between signal-to-noise ratio and depth of field linear. Also, using wavefront coding technology, Smith et al. [25] examined the iris information that could be captured in the NIR and VW spectra, addressing the possibility of using these multispectral data to improve recognition performance. Park and Kim [21] acquired in-focus iris images quickly at a distance, and Boddeti and Kumar [4] suggested extending the depth of field of iris imaging frameworks by using correlation filters. He et al. [12] analyzed the role of different NIR wavelengths in determining error rates. More recently, Yoon et al. [28] presented an imaging framework that can acquire NIR iris images at a distance of up to three meters, based on a face detection module and on a light-stripe laser device used to point the camera at the proper scene region. Boyce et al. [2] studied the image acquisition wavelength of revealed components of the iris and identified the important role of iris pigmentation. Although concluding that illumination inside the 700–900 nm optimally reveals the richness of the iris structure, they observed that irises with moderate levels of pigmentation could be imaged in the visible light with good quality.

8.3 UBIRIS.v2

UBIRIS.v2 is a database of visible wavelength iris images captured on-the-move and at a distance. As described in [26], the major purpose of the UBIRIS.v2 data set is to constitute a new tool to evaluate the feasibility of visible wavelength iris recognition under far from ideal imaging conditions. The various types of nonideal images, imaging distances, subject perspectives, and lighting conditions existent on this database could be of strong utility in the specification of the visible wavelength iris recognition feasibility and constraints.

The setup of the used imaging framework is given in Table 8.1. This framework was installed on a lounge under both natural and artificial lighting sources. Several marks were placed on the floor (between three and 10 m away from the acquisition device), and volunteers were asked for the image acquisition processes. Two distinct image acquisition sessions were performed separated by an interval of 1 week. From the first to the second sessions the location and orientation of the acquisition device and artificial light sources were changed in order to increase heterogeneity. Volunteers were at large majority Latin Caucasian (around 90%) and also black (8%) and Asian people (2%). Around 60% of the volunteers performed both imaging sessions, while 40% performed exclusively one, either during the first or second acquisition period.

Table 8.1 Details of the UBIRIS.v2 images, of the image acquisition framework and setup, and of the subjects that offered themselves as volunteers to the imaging sessions

Image acquisition framework and setup		
Camera = Canon EOS 5D	Color representation = sRGB	Shutter speed = 1/197 s
Lens aperture = F/6.4 – F/7	Focal length = 400 mm	F-number = F/6.3 – F/7.1
Exposure time = 1/200 s	ISO speed = ISO-1600	Metering mode = Pattern
Details of the manually cropped resultant images		
Width = 800 pixels	Height = 600 pixels	Format = tiff
Horizontal resolution = 72 dpi	Vertical resolution = 72 dpi	Bit depth = 24 bit
Volunteers		
Totals = subjects 261; irises 522; images 11 102		
Gender = male: 54.4%; female: 45.6%		
Age = [0,20], 6.6%; [21,25], 32.9%; [26,30], 23.8%; [31,35], 21.0%; [36,∞], 15.7%		
Iris pigmentation = light, 18.3%; medium, 42.6%; heavy, 39.1%		

**Fig. 8.3** Examples of close-up iris images acquired at different distances (between 8 at the far left and 4 m at the far right), on a continuously moving subject and under dynamic lighting conditions

Subjects were required to walk toward the acquisition device and look at several lateral marks that obliged them to rotate head and eyes, enabling the manual capturing of three images per meter, between 8 and 4 m, giving a total of 15 images per eye and session. As it is illustrated in Fig. 8.3, the used acquisition setup and protocol enabled to acquire sequences of iris images with different scale, gaze, and under different types of illumination for each subject. The used acquisition hardware and setup (Table 8.1) enables the capturing of iris data with dimensions that in most cases is compliant with the standard for iris images (ISO/IEC, 2004) that recommends a minimum of 100 and preferable 200 pixels across the iris diameter.

A varying number of subjects offered as volunteers for the first, second, or for both imaging sessions. However, assuming that each iris image can be used to generate a biometric template, the other images from the same eye can be used to assess *match* variability, and all the remaining images can be used to assess *non-match* variability, it is possible to obtain a bound for the error that is possible to be tested with statistical significance.

The 11 102 images of the UBIRIS.v2 database enable respectively 127 746 match and 61 482 804 non-match comparisons. This guarantees statistical significance in experiments with an empirical error rate \hat{P} down to $1,623 \times 10^{-6}\%$. However, this value should be regarded as a lower bound, which is substantially increased if the independence between images is not assumed and error correlations are taken into account.

8.4 Less Constraining Acquisition Frameworks

Figure 8.4a gives a global perspective of the used acquisition framework and of the environment where data of the UBIRIS.v2 database was collected. In order to disburden the imaging sessions for volunteers and maximize the number of usable images per subject, we decided to perform data acquisition manually. However, a completely automated process could have been performed, using state-of-the-art techniques of human silhouette and face detection and object tracking.

In the meanwhile, a completely automated acquisition framework was devised, being composed by two commercial web cameras, a pan-and-tilt device and the same high-resolution camera as before (Fig. 8.4b). The process starts by a software module that performs the detection of human silhouettes, according to the data acquired from one of the web cameras. Using this information and a set of semantic rules, a region of the scene is cropped and given to the real-time face detector module (according to the well-known method of Viola and Jones [27]). This phase given the 2D position (x,y) of a face in the scene, which is then sent to a stereo vision module that collects data from both web cameras and infers the subject depth in the scene, that is, the distance z between the acquisition camera and the subject. Using the pan-and-tilt device, the acquisition camera is directed to the 3D scene point at coordinates (x,y,z) , and an image that contains approximately the region of the subject's head is captured. Finally, using a set of biologically based semantic rules, a region that contains the subject's eyes is cropped and used in the biometric recognition phases.

Such a completely automated framework would be continuously acquiring and processing data without any human supervision, which increases probability for the acquisition of extremely degraded data: images might be completely out of focus due to exaggerated movements and might be captured under unusable perspectives due to subjects gaze, and even out-of-iris images are probable due to failures in




Fig. 8.4 Overview of the image acquisition frameworks (a) used in collecting the UBIRIS.v2 data set (A, B), light sources (C, D), and subjects location (E) and (b) used to automatically perform image acquisition, with labels similar to the manual configuration




Fig. 8.5 Maximum amount of information possible to acquire at a distance and on-the-move through the used image acquisition framework and setup. (a) Average diameters of the captured irises as a function of the image acquisition distance and (b) average images resolution at the iris regions, as a function of the image acquisition distance

the human silhouette and face detection modules. Due to these probabilities, image quality assessment is fundamental and has as goal to quantify characteristics and fidelity of the segmented data, particularly in terms of its utility. Such assessment module should act as a filter for most of the acquired data and select the key frames that are more likely to be conveniently handled by the later recognition modules.

8.4.1 Image Assessments

The error rates of any biometric system mainly depend on the amount of information captured, as well as on the proportion of noise that obstructs the underlying texture and corrupts the discriminating information.

In order to establish a relationship between the image acquisition distance and the maximum amount of iris data that is possible to capture with the described image acquisition framework and setup, the average iris diameter of frontal images was measured, with respect to the acquisition distance. The obtained results are given in Fig. 8.5a and indicate that the diameter of the captured irises varies according to an inverse logarithmic function. Through trial-and-error interpolation, the following coarse approximation function $d(x) : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ was obtained:

$$d(x) \approx \frac{150}{\ln\left(\frac{x}{1.85}\right)} \quad (8.1)$$

where x is the image acquisition distance (meters) and $d(x)$ is the average diameter of the captured irises. This is confirmed in Fig. 8.5b that gives the images resolution in the iris regions, where the shaded areas represent the 95% confidence intervals.

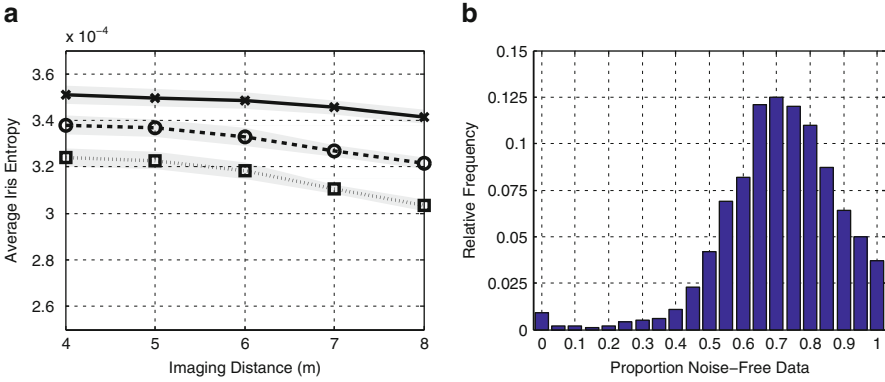


Fig. 8.6 Proportion of noise that occludes the iris texture (Fig. 8.6a) and average entropy of the noise-free iris regions as function of the levels of iris pigmentation (Fig. 8.6b). Light- and heavy-pigmented irises are represented by the continuous, dotted and dashed lines. (a) Proportion ([0,1]) of noise-free iris data within the captured iris rings and (b) average image entropy of the noise-free iris regions

Results were obtained through the division of the number of pixels that fall into the iris region by an iris area of 0.785 cm^2 , averaged from human eye's morphology studies that are publicly available. These observations confirm that the afore-described acquisition framework and setup enables the capturing of sufficient data to perform iris recognition, as near 50% of the database images have iris diameters close to the lower bound (140 pixels) recently proposed by Daugman [9].

The proportion of iris pixels that are occluded was obtained, and results are shown in the histogram of Fig. 8.6a. The horizontal axis denotes the proportion of noise-free pixels and the vertical axis the respective probability of occurrence in the database. It was observed that—on average—25–30% of the pixels that fall within the iris ring were corrupted by one of the local noise factors. Also, just about 3% of the images are completely noise-free, while around 0.9% contain full-noise data, correspondent to out-of-iris images or imaged in extremely poor lighting conditions.

Images texture was assessed by means of the *image entropy*, a statistical randomness measure widely used in the image processing domain. According to the conclusion reported by He et al. [12], we considered the levels of iris pigmentation as a factor that influences the amount of discriminating information captured for biometric purposes.

Let I be a grayscale image with g gray levels, and $p(k)$ be the probability of occurrence of the gray level k in I . The image entropy h is given by $h(I) = -\sum_0^{g-1} p(k) \log_2(p(k))$. Let B be the binary noise mask that gives the non-occluded iris regions of I . Thus, $B(x, y) = 0$ or 1 when the respective pixel (x, y) respectively belongs to the iris ring and is noise-free or not. The *entropy of the iris region* $ih(\cdot)$ of I was obtained taking exclusively into account pixels (x_i, y_i) such that $B(x_i, y_i) = 0$, that is, the noise-free iris data. Finally, for normalization purposes, the entropy value was divided by the area of the noise-free iris region

$$ih(I) = \frac{h(I(x,y)|B(x,y) = 0)}{\mathbb{I}_{\{B(x,y)=0\}}} \quad (8.2)$$

where $\mathbb{I}_{\{\cdot\}}$ is the characteristic function.

The obtained values are given in Fig. 8.6b. The horizontal axis corresponds to the imaging distance and the vertical axis to the entropy value. Data points give the observation average values for *light*-pigmented (continuous series)-pigmented, *medium* (dotted series), and *heavy*-pigmented irises (dashed series), as a function of the imaging distance. The shaded areas represent the 95% confidence intervals. It can be confirmed that the average entropy of the iris data decreases inversely with the imaging distance and, more evidently, with the levels of iris pigmentation.

8.4.2 Degradation Factors

Images of the UBIRIS.v2 data set are degraded by several factors and are high heterogeneous, regarding the lighting conditions of the environment. Through visual inspection, fourteen different factors were detected and classified into one of two major categories: *local* or *global*, as they affect exclusively image regions or the complete image. The *local* category comprises iris obstructions, reflections, and off-angle and partial images, while the *global* comprises poor-focused, motion-blurred, rotated, improper-lighting and out-of-iris images. Examples of the UBIRIS.v2 noise factors are given in Fig. 8.7.

8.4.3 Specificity

Several works about the iris recognition technology reported a very small—almost infinitesimal—probability of produce a false match in the comparison between signatures extracted from data with good quality (e.g., [5, 9, 15], and [16]). This is due to the chaotic appearance of the main components of the iris texture and is regarded as one of the technology’s major advantages, when compared to other biometric traits. Also, this is of special interest for the type of recognition discussed in this chapter: to guarantee that the comparison between templates (extracted from good quality data) and samples extracted from degraded iris data (or even from partial or non-iris regions due to failures on the eye detection and segmentation modules) will not frequently produce false matches.

This hypothesis was tested through a procedure illustrated in Fig. 8.8. Using the recognition method proposed by Daugman [9]—composed by iris segmentation, normalization (Daugman Rubber Sheet), encoding (bidimensional Gabor wavelets), and matching (Hamming distance)—we extracted 1,000 signatures from UBIRIS.v2 iris images with good quality. Further, we extracted a set of sample signatures from




Fig. 8.7 Comparison between a good quality image (Fig. 8.7) and several types of nonideal images of the UBIRIS.v2 database. These images resulted of less constrained imaging conditions, under varying lighting conditions, at-a-distance and on-the-move subjects. (a) Good quality iris image. (b) Off-angle iris image. (c) Poor focused iris image. (d) Rotated iris image. (e) Motion-blurred iris image. (f) Iris obstructions due to eyelids. (g) Iris obstructions due to eyelashes. (h) Iris obstructions due to glasses. (i) Iris obstructions due to contact lenses. (j) Iris obstructions due to hair. (k) Iris imaging in poor lighting conditions. (l) Iris with specular reflections. (m) Iris with lighting reflections. (n) Partial captured iris. (o) Out-of-iris image

1,000 iris images with very poor quality, 10,000 non-iris or partial iris images, and 10,000 natural and synthetic textures images. Finally, we performed a one-against-all comparison between each template and the set of samples, giving a total of 21,000,000 comparisons. During these tests, we did not get a single dissimilarity value close to the usual acceptance threshold (0.33), which means that not even a single false acceptance was observed if the traditional acceptance thresholds are used.

Figure 8.9 gives the histogram of the obtained dissimilarity values (vertical bars) and the approximated Gaussian distribution (line plot with $\mu = 0.49992$ and $\sigma = 0.02419$). We confirmed that, even on highly degraded data, the used iris encoding and comparison strategies produce a false match with almost null probability. Based on the parameters of the fitted Gaussian distribution, the probability of producing a dissimilarity value lower than 0.33 will be approximately 1.03923×10^{-12} . Once again, the role of this value for the type of recognition discussed in this chapter

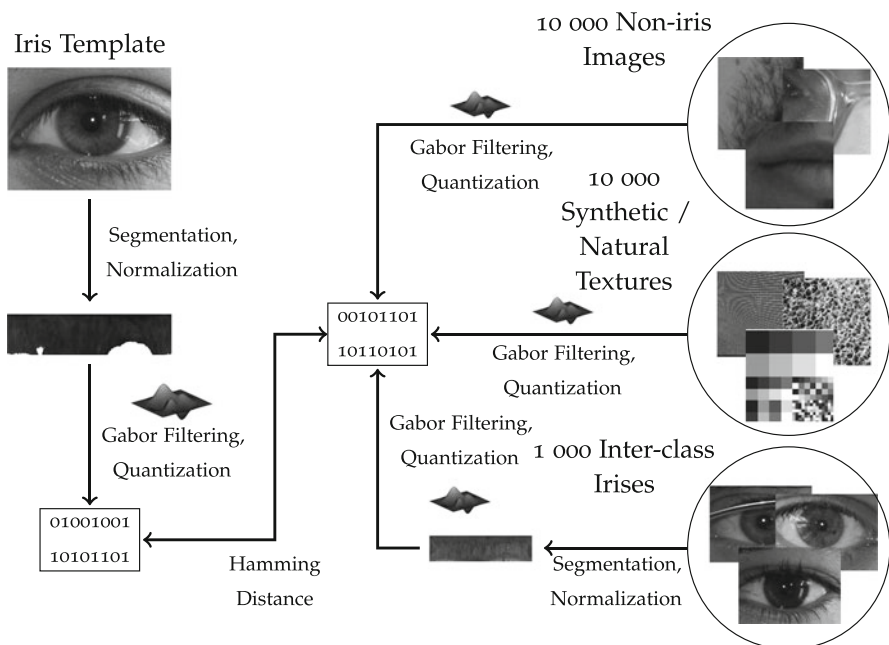


Fig. 8.8 Setup of the experiments performed to evaluate the probability of producing a false match in the comparison between iris signatures extracted from good quality data (“iris template”) and signature samples resultant from iris data with bad quality or even partial or non-iris data. We used the main recognition stages proposed by Daugman and successfully deployed in recognition systems to evaluate the probability of producing a false match in these situations

should be stressed: it can be assumed with extreme high confidence that such recognition systems will not produce false matches and, thus, any match reported has a full probability of being genuine. This means that, independently of how much false rejections will be frequent (due to lighting variations, movements, and gazes), any reported match is highly reliable.

8.4.4 Sensitivity

As expected, the uncontrolled imaging conditions and acquisition setups lead to data with notoriously heterogeneous quality, which will have evident impact in the recognition rates of VW iris biometric systems. In this scope, quality assessment is a fundamental task. The goal is to quantify characteristics and fidelity of the segmented data, particularly in terms of its utility. This is essential, as performing recognition in too much degraded data decreases matching accuracy and increases computational complexity.




Fig. 8.9 Histogram of the obtained dissimilarities when comparing signatures extracted from 1,000 templates with good quality and 21,000 signatures extracted from iris images with bad quality, partial irises, and non-iris data. R-square gives the goodness-of-fit of the plotted Gaussian distribution with $\mu = 0.499$ and $\sigma = 0.024$ to the obtained results

According to the afore observations, this section aims at give an approximation of the recognition rate that VW iris recognition systems can achieve, according to data quality. For such, we used the classical Daugman's recognition strategy [9], as it is the most well known and is functioning in the large majority of the commercially deployed systems. According to this choice, the segmented boundaries were normalized to dimensionless polar coordinates. Then, a bank of Gabor filters was used to analyze the iris texture and the angle of each phasor quantized to one of four quadrants. Finally, the fractional Hamming distance gave the dissimilarity between two irises. A subset of 10,427 UBIRIS.v2 images was selected, which under visual inspection, we verified that the segmentation method has accurately segmented. Then, for all subsequent experiments, this set was divided according to different criteria. For comprehensibility, we refer to a *recognition test* when each sample of a data set is matched against all the remaining images of the the same data set, resulting in two types of comparisons: intra-class (match) and interclass (non-match). As suggested by Daugman [8], for two-choice decisions, the decidability index d' measures how well separated are the two types of distributions and recognition errors correspond to their overlap area:

$$d' = \frac{|\mu_E - \mu_I|}{\sqrt{\frac{1}{2}(\sigma_I^2 + \sigma_E^2)}} \quad (8.3)$$

where μ_I and μ_E are the means of the two distributions and σ_I and σ_E their standard deviations.




Fig. 8.10 Comparison between the histograms of the fractional Hamming distances obtained (a) for a subset of 10,427 images of the UBIRIS.v2 database and (b) when samples of poor quality are not considered to the recognition test. Figure at far right gives the corresponding receiver operating characteristic curves

Figure 8.10 compares the histograms of the fractional Hamming distances for the match (light bars) and non-match (dark bars) comparisons obtained when all images were used in the recognition test (figure at the far left) and when the poorest quality samples (according to the visual perception of quality) were rejected (figure at the center). The line plots correspond to the fitted normal distributions, and the upper left corner gives the corresponding decidability index d' . As general considerations, we confirmed that values obtained for the non-match distributions do not significantly vary according to the quality of the data and are almost the same reported for the NIR-constrained recognition setups. Oppositely, there is a significant movement of the match distributions toward the non-matches, substantially decreasing the sensitivity of the system, if traditional acceptance thresholds are used. Due to this, the decidability of the VW recognition systems significantly varied. Figure 8.10c shows how the true and false matches in our system would change according to different decision thresholds, when no quality is considered (continuous line) and when only samples with good quality are considered for recognition (dashed line). Here, we plot the area under curve (AUC) for both setups, which significantly augments as the poorest quality samples are rejected.

The next question to consider is about the effect of each factor in the effectiveness of VW iris recognition systems, especially in its sensitivity. For such, we repeated the “one-against-all” recognition setup, but selecting several subsets, each one containing samples that are notoriously degraded by eac analyzed factor: scale, focus, gaze, and levels of iris pigmentation. It should be noted that although this experiment enables to perceive *how much* recognition effectiveness is degraded by each factor, it cannot be regarded as an effective evaluation. This will demand a data set with images degraded exclusively by one factor and UBIRIS.v2 images (as well as all the remaining available data sets) are simultaneously degraded by various factors.

The initial evaluation is about the effect of scale in the recognition sensitivity (Fig. 8.11). For such, two subsets with good quality images that were acquired at the farthest (8 m) and closest (4 m) distances were selected. Under visual inspection, the quality of each image was verified, having as main concern to exclusively select




Fig. 8.11 UBIRIS.v2 images contain images captured at notoriously different scales, between 4 and 8 m (a). Figures at the *center* and far *right* compare the match and non-match histograms of the fractional Hamming distances obtained for images captured at the farthest (b) and closest (c) scales




Fig. 8.12 The acquisition of a *small* moving target (iris) at relatively *large* distances propitiates very different levels of image focus (a). Figures at the *center* and far *right* compare the match and non-match histograms of the fractional Hamming distances obtained for blurred (b) and focused (c) images

focused images with practically unoccluded irises and gaze aligned with the optical axis of the camera. Figure 8.11 shows how the histograms of the match and non-match distributions would change according to the scale criterion. The observed values for the non-match distributions remained approximately equal, in opposition with the match distributions that appear to move toward the non-matches, as the acquisition distance augments.

Then, we assessed the impact that levels of image focus could have in the sensitivity of VW recognition systems. A subset composed by good quality images was selected, composed by focused images with occluded irises and gaze aligned with the optical axis of the camera. As defocus is equivalent to multiplying the 2D Fourier transform of a perfectly focused image by a 2D Fourier transform of a Gaussian kernel, an artificial set of defocused images was obtained by the convolution with the previous images with a Gaussian kernel of width 5 and $\sigma 2.0$. As expected, we observed a relatively minor impact of focus in recognition effectiveness (Fig. 8.12), as it is known that most of the iris-discriminating information spans between two and three octaves of the original scale and only severe defocus would notoriously affect the recognition process. In our opinion, this is one of the biggest advantages of the iris, when compared to other biometric traits: using the lowest and middle-low frequency components for recognition improves the robustness to noise (focus and artifacts) in less constrained acquisition setups.




Fig. 8.13 Acquiring data from moving subjects and without requiring them any type of cooperation increases the probability of capturing off-angle irises, when the subject’s gaze is not aligned with the optical axis of the camera (a). These variations cause perspective deformations in the iris texture and significant changes in phase. Taking into account that current state-of-the-art techniques are phase based, this may lead to significant deteriorations in the recognition effectiveness, as it is illustrated by the results given in (b) off-angle images and (c) aligned images

Due to the intrinsic properties of the used iris encoding and matching strategies, it is expected that variations in subjects gaze may have big influence in recognition effectiveness, when compared to other factors. We suspect that gaze propitiates significant perspective deformations in the iris texture, and these cause significant changes in phase, deteriorating the recognition effectiveness. This is reiterated in Fig. 8.13 that compares the results obtained when the recognition tests were performed on images that under visual inspection were considered as aligned and another set that contains images with moderate to severe deviations between the subjects gaze and the optical axis of the camera.

As stated before, most of the skepticism in the development of iris recognition systems that operate in the VW is for high-pigmented irises, corresponding to dark brown and almost black perceived colors, with a very low albedo. While they constitute the large majority of the world population, it is considered that these irises would demand very strong light sources, in order to be acquired with sufficient discriminating information. Here, we infer *how much* the levels of iris pigmentation increase the recognition challenges, by analyzing the separability between intra-class and interclass comparisons obtained for two subsets of good quality images (focused, without significant occlusions or gaze deviations) with very different levels of iris pigmentation. The “light” subset contains light-pigmented irises, corresponding to blue and light green eyes. The “dark” subset is composed by dark brown and black irises and includes almost all the images of the African and Asian volunteers of the UBIRIS.v2 data set. Figure 8.14 confirms that the levels of iris pigmentation constitute an obstacle to VW iris recognition, as the decidability values obtained for light-pigmented irises (right histograms) were notoriously higher than those obtained when heavily pigmented images were used. The justification for these result lies in the spectral absorbance of the predominant pigment of the human iris (brown-black eumelanin) that has a peak in the VW. Thus, the levels of iris pigmentation are in direct proportion of the tissue absorbance and inverse to the perceived albedo, which turns more difficult to capture the discriminating patterns of heavily pigmented irises (ridges, valleys, furrows, and crypts).




Fig. 8.14 The spectral absorbance of the human iris predominant pigment (eumelanin) has a peak in the VW, which turns more difficult to capture the discriminating patterns of high-pigmented irises. This is confirmed by the reduced separability between the intra-class and inter-class comparisons obtained for high-pigmented irises (c) when compared with the results obtained for light-pigmented ones (b)

8.5 Conclusions and Trends

The possibility of performing automatic recognition of human beings in uncontrolled environments and without requiring any type of cooperation is of evident interest for forensic and security purposes and represents a *grand challenge* for the pattern recognition community. To the best of our knowledge, no method was developed to achieve such purpose, which will have effective consequences in the everyday style of modern societies.

This chapter addressed the use of VW light to perform iris biometric recognition in scenarios that are far less constrained than those where currently deployed systems are operating with success: at larger image acquisition ranges (between 4 and 10 m) and on moving subjects. As expected, the acquired *real-world* data brings many challenges to the pattern recognition task, as it might be degraded by several factors and contains various types of noise artifacts that obstruct the discriminating iris texture and make the recognition process difficult.

We start by presenting the main characteristics of a dataset that is freely available for the research community, UBIRIS.v2. This dataset can provide basis for the fair comparisons between different proposals.

The optical properties of the iris, namely, its spectral absorbance and reflectance, were overviewed in order to give an idea of the major discriminating points between the appearance of data acquired in the well-known NIR scenario and in the VW. Special attention was paid to the factors that are likely to degrade the effectiveness of VW iris recognition systems and motivated serious skepticism about the possibility of performing reliable recognition in such challenging conditions. Even so, preliminary experiments appear to confirm that this task is possible, which is justified by two reasons: (1) an extremely low probability for the existence of false acceptances was observed, even when using the same recognition strategies used for NIR data, and (2) the observed sensitivity of recognition systems in such data, which points that an evident amount of information that discriminates between




Fig. 8.15 Periocular biometrics refers to the automatic recognition of human beings using not only the information of the iris texture but also of the surrounding information (eyelids, eyelashes, skin, eyebrow). This type of recognition can provide a significant improvement toward the development of recognition methods that perform surreptitiously and in unconstrained environments

different subjects is able to be acquired in the described conditions and used for recognition purposes.

Due to the extremely appealing potential applications of these systems, a growing number of researchers are working on this field and proposing specialized strategies to perform the automatic acquisition, segmentation, normalization, and encoding and matching of real-world VW data. Among them, there is an extremely promising new type of biometric recognition called *periocular biometrics* that refers to the regions in the immediate vicinity of the eye (Fig. 8.15) and attempts to perform recognition based not only in the iris but also to its neighborhood.

As argued by Park et al. [22] and Miller et al. [18], periocular recognition is a trade-off between using the entire face region or only the iris: it avoids the resolution problems of iris images captured at large distances and can be used for a wide range of distances. Also, face images acquired from unconstrained environments often suffer from poor illumination, motion blur, low resolution, and pose variations, that significantly degrade the effectiveness of face recognition techniques. To the best of our knowledge, few studies have been conducted on the use of the periocular region as a biometric trait. Park et al. [22] used both local and global image features to match periocular images acquired with visible light and established its utility as a soft biometric trait. Also, Miller et al. [18] used local binary pattern (LBP) to encode and match periocular images.

Acknowledgements The financial support given by FCT-Fundação para a Ciência e Tecnologia and FEDER in the scope of the PTDC/EIA-EIA/103945/2008 NECOVID (Negative Covert Biometric Identification) research project is acknowledged.

References

1. American National Standards Institute: American national standard for the safe use of lasers and LEDs used in optical fiber transmission systems ANSI Z136-2 (1988)
2. Boddeti, N., Kumar, V.: Extended depth of field iris recognition with correlation filters. In: Proceedings of Computer Vision and Pattern Recognition Workshop on Biometrics, New York, pp. 51–59 (2006)
3. Bowyer, K., Hollingsworth, K., Flynn, P.: Image understanding for iris biometrics: a survey. *Comput. Vis. Image Underst.* **110**, 281–307 (2008)
4. Boyce, C., Ross, A., Monaco, M., Hornak, L., Li, X.: Multispectral iris analysis: a preliminary study. In: Proceedings of the First IEEE International Conference on Biometrics: Theory, Applications, and Systems, Arlington, pp. 1–8 (2008)
5. Cambier, J.: Iridian large database performance. Tech. rep., Iridian Technologies (2007)
6. Commission International de l'Eclairage: photobiological safety standards for safety standards for lamps. Report of TC 6-38 CIE 134-3-99 (1999)
7. Daugman, J.: High confidence visual recognition of persons by a test of statistical independence. *IEEE Trans. Pattern Anal. Mach. Intell.* **15**(11), 1148–1161 (1993)
8. Daugman, J.: Biometric decision landscapes. Tech. rep. TR482, University of Cambridge Computer Laboratory (2000)
9. Daugman, J.: How iris recognition works. *IEEE Trans. Circuits Syst. Video Technol.* **14**(1), 21–30 (2004)
10. Daugman, J.: Probing the uniqueness and randomness of iriscodes: results from 200 billion iris pair comparisons. *Proc. IEEE* **94**, 1927–1935 (2006)
11. Fancourt, C., Bogoni, L., Hanna, K., Guo, Y., Wildes, R., Takahashi, N., Jain, U.: Iris recognition at a distance. In: Proceedings of the 2005 IAPR Conference on Audio and Video Based Biometric Person Authentication, Hilton Rye Town, pp. 1–13 (2005)
12. He, Y., Cui, J., Tan, T., Wang, Y.: Key techniques and methods for imaging iris in focus. In: Proceedings of the IEEE International Conference on Pattern Recognition, Hong Kong, pp. 557–561 (2006)
13. Honeywell International Inc.: A distance iris recognition. United States Patent 0036397 (2007)
14. Imai, F.: Preliminary experiment for spectral reflectance estimation of human iris using a digital camera. Tech. rep., Munsell Color Science Laboratories, Rochester Institute of Technology (2000)
15. International Biometric Group: Independent test of iris recognition technology (2005)
16. Mansfield, T., Kelly, G., Chandler, D., Kane, J.: Biometric product testing final report. National Physical Laboratory, Middlesex (2001)
17. Matey, J.R., Ackerman, D., Bergen, J., Tinker, M.: Iris recognition in less constrained environments. In: Ratha, N.K., Govindaraju, V. (eds.) *Advances in Biometrics: Sensors, Algorithms and Systems*, pp. 107–131. Springer-Verlag London (2007)
18. Miller, P.E., Rawls, A., Pundlik, S., Woodard, D.: Personal identification using periocular skin texture. In: Proceedings of the 2010 ACM Symposium on Applied Computing, ACM, New York 1496–1500 (2010)
19. Narayanswamy, R., Johnson, G., Silveira, P., Wach H.: Extending the imaging volume for biometric iris recognition. *Appl. Opt.* **44**(5), 701–712 (2005)
20. Nematí, B., Rylander III, H.G., Welch, A.J.: Optical properties of conjunctiva, sclera, and the ciliary body and their consequences for transscleral cyclophotocoagulation. *Appl. Opt.* **35**(19), 3321–3327 (1996)
21. Park, K., Kim, J.: A real-time focusing algorithm for iris recognition camera. *IEEE Trans. Syst. Man Cybern.* **35**(3), 441–444 (2005)
22. Park, U., Ross, A., Jain, A.K.: Periocular biometrics in the visible spectrum: a feasibility study. In: Proceedings of the Biometrics Applications and Systems, Washington DC (BTAS), pp. 153–158. IEEE, Piscataway (2009)

23. Ross, A., Crihalmeanu, S., Hornak, L., Schuckers, S.: A centralized web-enabled multimodal biometric database. In: Proceedings of the 2004 Biometric Consortium Conference (BCC), Arlington (2004)
24. Sarna, T., Meredith, P.: The physical and chemical properties of eumelanin. *Pigment Cell Res.* **19**, 572–594 (2010)
25. Smith, K., Pauca, V.P., Ross, A., Torgersen, T., King, M.: Extended evaluation of simulated wavefront coding technology in iris recognition. In: Proceedings of the First IEEE International Conference on Biometrics: Theory, Applications, and Systems, Crystal City, pp. 1–7 (2007)
26. Proenca, H., Filipe, S., Santos, R., Oliveira, J., Alexandre, L.A.: The ubiris v2: a database of visible wavelength iris images captured on-the-move and at-a-distance. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(8), 1529–1535 (2007)
27. Viola, P., Jones, M.: Robust real-time face detection. *Int. J. Comput. Vis.* **57**(2), 137–154 (2002)
28. Yoon, S., Bae, K., Ryoung, K., Kim, P.: Pan-tilt-zoom based iris image capturing system for unconstrained user environments at a distance. *Lect. Notes Comput. Sci.* **4642**, 653–662 (2007)

Chapter 9

Multispectral Iris Fusion and Cross-Spectrum Matching

Mark J. Burge and Matthew Monaco

Abstract Traditionally, only a narrow band of the Near-Infrared (NIR) spectrum (700–900nm) is utilized for iris recognition since this alleviates any physical discomfort from illumination, reduces specular reflections and increases the amount of texture captured for some iris colors. However, previous research has shown that matching performance is not invariant to iris color and can be improved by imaging outside of the NIR spectrum. Building on this research, we demonstrate that iris texture increases with the frequency of the illumination for lighter colored sections of the iris and decreases for darker sections. Using registered visible light and NIR iris images captured using a single-lens multispectral camera, we illustrate how physiological properties of the iris (e.g., the amount and distribution of melanin) impact the transmission, absorbance, and reflectance of different portions of the electromagnetic spectrum and consequently affect the quality of the imaged iris texture. We introduce a novel iris code, Multispectral Enhanced irisCode (MEC), which uses pixel-level fusion algorithms to exploit texture variations elicited by illuminating the iris at different frequencies to improve iris matcher performance and reduce Failure-To-Enroll (FTE) rates. Finally, we present a model for approximating an NIR iris image using features derived from the color and structure of a visible light iris image. The simulated NIR images generated by this model are designed to improve the interoperability between legacy NIR iris images and those acquired under visible light by enabling cross wavelength matching of NIR and visible light iris images.

M.J. Burge (✉)
The MITRE Corporation, McLean, VA, USA
e-mail: mburge.mitre.org

M. Monaco
Noblis, Falls Church, VA, USA
e-mail: Matthew.Monaco@Noblis.org

9.1 Introduction

Traditionally, only a narrow band of the Near-Infrared (NIR) spectrum (700–900 nm) is utilized for iris recognition [6] since this alleviates any physical discomfort from illumination, reduces specular reflections, and increases the amount of texture captured for some iris colors. However, previous research has shown [7] that matching performance is not invariant to iris color and can be improved by imaging outside of the NIR spectrum and that physiological properties of the iris (e.g., the amount and distribution of melanin) impact the transmission, absorbance, and reflectance of different portions of the electromagnetic spectrum and the ability to image well-defined iris texture [1].

Texture-based encoding of the iris is commonly used to create a biometric template that can be quickly and accurately matched against large datasets. Prior to recognition, the iris must be localized and segmented from the overall image. Localization is typically done by finding landmark features of the iris (e.g., the limbic boundary and the pupil) and removing background features (e.g., eyelids, eye lashes, sclera). Once the iris is segmented, it is unwrapped using a polar transformation from its natural circular form into a fixed-size Cartesian rectangle. Once in this form, a quantification of the textural content of the iris is computed using 2D Gabor wavelet filters [3] to create a template, commonly referred to as an IrisCode.

Phasors, which contain information on the orientation, spatial frequency, and position of texture within the image, are used to create templates and masks. Since an iris template can be represented by a fixed-size array of bits, efficient large-scale matching can be accomplished by computing a circularly shifted, unmasked, and normalized Hamming distance between two templates over some range. This Hamming distance measure can then be used as a test of the statistical independence between the two iris images to determine if they can be deemed to be from different subjects [4].

Gabor-based iris encoding algorithms are robust because they quantify the phase structure of the iris for encoding and discard the amplitude. Note that the use of the amplitude, and not the phase, in the figures in this paper is solely for illustrative purposes. The actual measurements used in our fusion and transformation algorithms are based on phase information extracted using multiscale log-Gabor filter banks.

9.2 Outline

Section 9.3 presents the design of single-lens and stereo rig multispectral cameras capable of simultaneously acquiring registered four channel (i.e., 460, 550, 670, and 800 nm) images of an iris. Section 9.4 demonstrates that iris texture increases with the frequency of the illumination for lighter-colored sections of the iris and

decreases for darker sections. Section 9.5 presents a model for approximating an NIR iris image using features derived from the color and structure of a Visible Light (VL) iris image. Section 9.6 introduces a novel iris code which uses pixel-level fusion algorithms to exploit texture variations elicited by illuminating the iris at different frequencies to improve iris matcher performance and reduce FTE rates. Section 9.7 summarizes our results and describes our ongoing evaluation and assessment.

9.3 Multispectral Iris Cameras

We are using a single-lens multispectral camera and stereo imaging rig capable of simultaneous acquisition across the VL and NIR portions of the spectrum. These cameras are being used to collect a database of pixel-level registered multispectral iris images. The multispectral iris images, Fig. 9.1, are being used to develop and evaluate supervised learning algorithms for the pixel-level transformation of VL iris images into a form suitable for matching against existing NIR iris images and fusion algorithms which exploit texture variations elicited by illuminating the iris at different frequencies to improve iris matcher performance and reduce FTE rates.

The multispectral acquisition system shown in Fig. 9.2a utilizes a single-lens camera where four images are simultaneously captured across four different wavelengths by two Sony ICX205AL sensors (with spectral response ranges between 400 and 1,000 nm with a peak of 550 nm) to capture the IR and red components and two Sony RGBICX205 sensors (blue response ranging from 400 to 550 nm with a peak at 460 nm and a green response from 400 to 650 nm with a peak at 550 nm) to acquire the green and blue channels. Optimal imaging parameters across all wavelengths of interest are obtained by illuminating the iris using a broadband (350–1,100 nm) quartz halogen ring-light located between the imaging device and the subject.

The stereo rig illustrated in Fig. 9.2b is built from two Hitachi KP-F120 cameras, selected for their sensitivity across the 400–1,000 nm band. Each camera is equipped with VL-cut filters enabling the selective acquisition of (NIR, NIR), (NIR, VL), and (VL, VL) stereo pairs. Iris specularities are limited to a small portion of the pupil by using a cluster of NIR (750–950 nm) LED illuminators centered between the cameras together with adaptive illumination which varies in intensity based on ambient lighting conditions. Since the intrinsic (i.e., image point from camera to pixel coordinates) and extrinsic (i.e., relative position and orientation of the two cameras) parameters of the stereo rig are known a priori, it is possible to solve the correspondence problem and reconstruct a 3-D depth map of the iris.

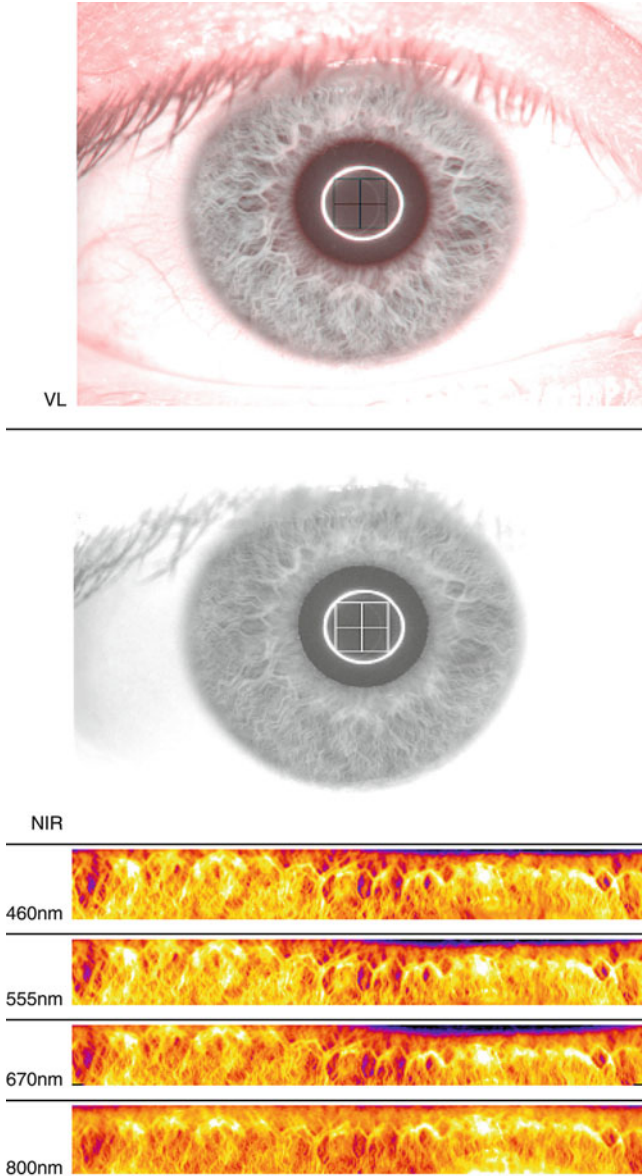


Fig. 9.1 A spatially registered four-channel iris image of a light blue eye

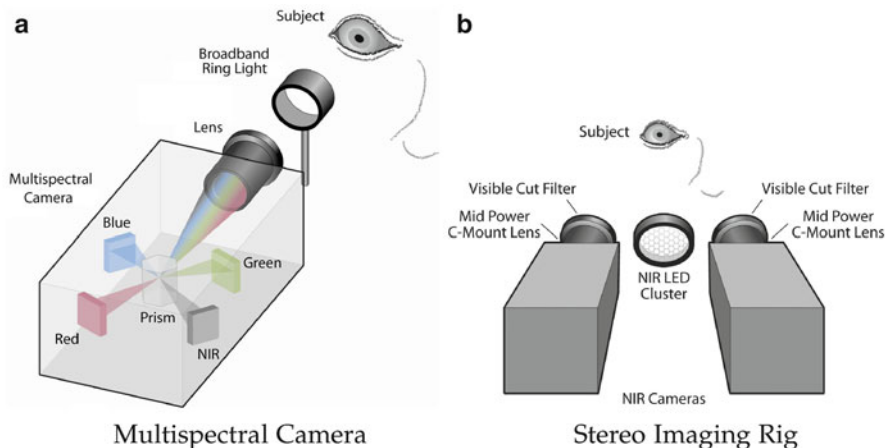


Fig. 9.2 (a) Simultaneous acquisition of spatially registered iris images is carried out using four optically aligned sensors, two Sony ICX205AL and two Sony RGBICX205 (n.b., internal optics simplified for illustrative purposes). Since these multispectral iris images are spatially registered at the sensor level, it is possible to observe, with the aim of modeling, the effect of wavelength on the imaging of iris structure. The stereo rig (b) provides known camera geometry for the simultaneous acquisition of (NIR, NIR), (NIR, VL), and (VL, VL) iris pairs with slightly different spatial orientations. Since the intrinsic (i.e., image point from camera to pixel coordinates) and extrinsic (i.e., relative position and orientation of the two cameras) parameters of the stereo rig are known a priori, it is possible to solve the correspondence problem and reconstruct a 3-D depth map of the iris

9.4 Multispectral Iris Texture

The amount of texture information extractable from a given iris color cluster is dependent on the wavelength under which it is illuminated. In the case of lighter-colored clusters, more information is available at shorter wavelengths than at longer wavelengths and conversely for darker-colored clusters. Figure 9.3 illustrates this with details from registered images of “green” (left) and “blue” (right) irises acquired at 460, 670, and 800 nm. In the darker “green” iris, the crypt is clearly visible at 800 nm but by 460 nm has deteriorated to noise, while the opposite occurs in the lighter “blue” iris where the crypt is clearly visible at 460 nm but is much less pronounced by 800 nm.

Our hypothesis is that across the portion of the spectrum under consideration, the texture of lighter-colored clusters of the iris increases with the frequency of the illumination and decreases in a similar way for darker clusters. This relation is demonstrated in Fig. 9.3 using the following texture measures derived from the gray-level co-occurrence matrix [5] where $P_{\theta,d}(a,b)$ [8] describes how frequently two pixels with gray-levels a, b appear in a window separated by distance d in direction θ :

$$T_{asm} = \sum_{a,b} P_{\theta,d}^2(a,b) \quad (9.1)$$

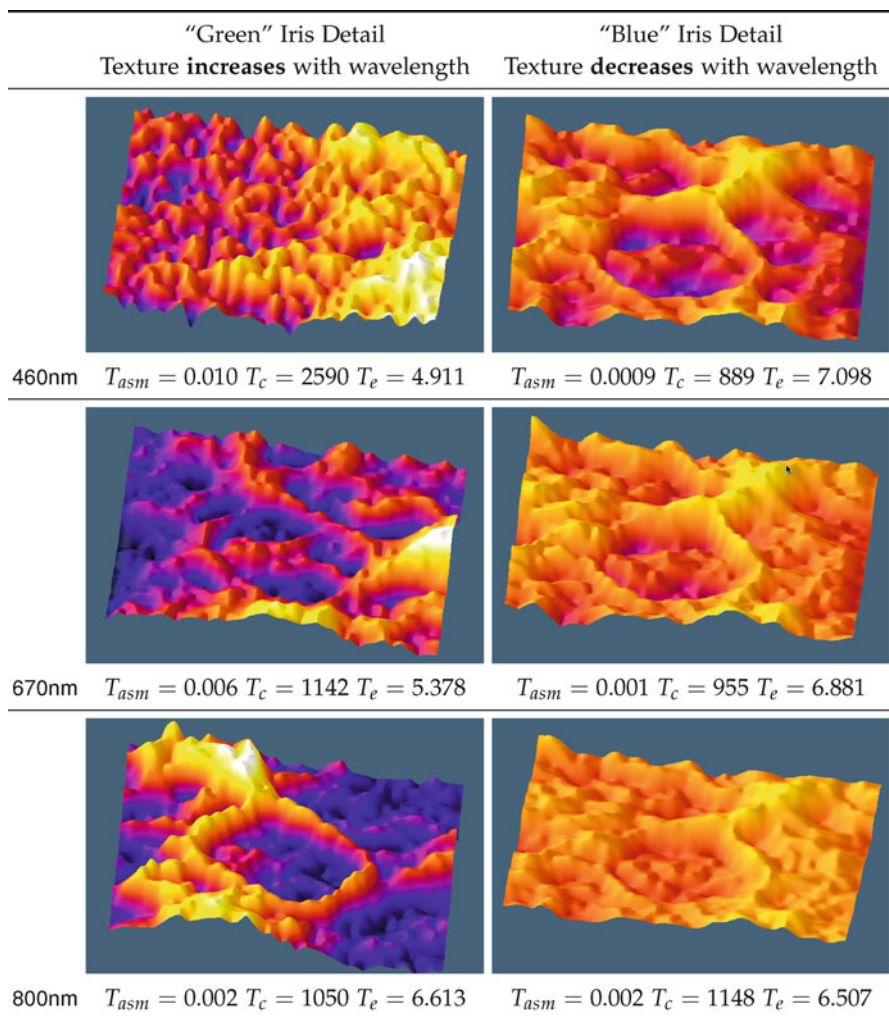


Fig. 9.3 In this example, details from “green” (*left*) and “blue” (*right*) irises from registered images acquired at 460, 670, and 800 nm are shown as heat-coded height maps where hotter colors (e.g., *white, yellow, and red*) indicate higher values and cooler colors (e.g., *purple, blue, black*) indicate lower values. In the darker “green” iris, the crypt is clearly visible at 800 nm but by 460 nm has deteriorated to noise, while the opposite occurs in the lighter “blue” iris where the crypt is clearly visible at 460 nm but is much less pronounced at 800 nm. Our hypothesis is that the information content of the image increases with the frequency of the illumination for lighter-colored sections of the iris and decreases for darker sections. This relation is demonstrated using the gray-level co-occurrence matrix [5] derived texture measures: angular second moment T_{asm} (Eq. 9.1), contrast T_c (Eq. 9.2), and entropy T_e (Eq. 9.3). Note that T_{asm} and T_c decrease with information content, while T_e increases

angular second moment,

$$T_c = \sum_{a,b} |a - b|^2 P_{\theta,d}^1(a,b) \quad (9.2)$$

contrast, and

$$T_e = \sum_{a,b} P_{\theta,d}(a,b) \log_2 P_{\theta,d}(a,b) \quad (9.3)$$

entropy. Note that the texture measures T_{asm} and T_c decrease with information content while T_e increases. We combine these, and other texture measures derived from multiscale log-Gabor filter banks, to create the texture function $T(I_w(x,y))$ defined in Eq. 9.5.

9.5 Modeling NIR Iris Images

Due to the reflectance properties of melanin, which is partially responsible for the color of the iris commercial-off-the-shelf, iris acquisition devices (e.g., HIIDE, IrisAccess) acquire iris images using NIR illumination. While NIR images enable improved matching performance, especially for darker eyes, it can be difficult to directly match the templates derived from these images to those derived from irises which have been imaged in the visible spectrum.

We present a model for approximating an NIR iris image using features derived from the color and structure of a VL iris image. The simulated NIR images generated by this model are designed to improve the interoperability between legacy NIR iris images and those acquired under VL by enabling cross wavelength matching of NIR and VL iris images.

Eye color does not vary smoothly across the iris, but instead is arranged in clusters of similar colors which can have abrupt boundaries that arise from physiological structures within the anterior and posterior layers of the iris. Given that overall eye color arises from the montage of small color patches found in the iris (e.g., hazel eyes are a combination of many color patches), building a single model, or even a small collection of models, for each “eye color” is not sufficient (see Fig. 9.4).

A multistage supervised learning process is used to develop a set of mathematical models for approximating NIR iris images from VL iris images. In the first stage of our approach, the cluster analysis in the $L^*a^*b^*$ color space [2] on the VL channels of the multispectral iris images is used to identify a set of n representative color clusters. Next, for each of the n clusters, a set of tuples

$$(L^*, a^*, b^*, NIR)$$




Fig. 9.4 Intra-eye color distribution. Eye color does not vary smoothly across the iris but instead is grouped in clusters of similar colors that correspond to physiological structures across the anterior and posterior layers of the iris. The color clusters used in the supervised learning phase are derived from a k -means clustering in the $L^*a^*b^*$ color space of a boosted sampling drawn from the registered multispectral images

is constructed that associates the $L^*a^*b^*$ and NIR values of the registered images. These tuples are then augmented with local structural information derived from k multiscale log-Gabor filters $G(w)$, where w is the filter's center frequency,

$$(G(w_1), G(w_2), \dots, G(w_k), L^*, a^*, b^*, \text{NIR})$$

such that, for each of the n clusters, a mathematical model,

$$f(G(w_1), G(w_2), \dots, G(w_k), L^*, a^*, b^*)_n \approx \text{NIR} \quad (9.4)$$

can be constructed using supervised learning to approximate the final NIR term given the first $k + 3$ terms.

Finally, each of the n learned functions which have been tuned to transform a small range of iris colors into approximations of their NIR values are combined to build a model for approximating an NIR iris image using features derived from the color and structure of a VL iris image.

9.6 Multispectral Iris Fusion

By analyzing the reflectance of the iris at each band of the multispectral image and selecting the areas of the iris image that have the highest probability of resulting in an accurate match, a single high-confidence image, or Multispectral Enhanced irisCode (MEC), can be fused. Figure 9.5 illustrates the generation of such a MEC by image-level fusion of multispectral iris images. For each wavelength w , a mask M_w is generated which is set if and only if the texture measure T_e , across all imaged wavelengths, is highest at that location in wavelength w . More formally, let I_w be an iris imaged at wavelength w and $T(I_w(x, y))$ be a monotonically increasing measure of the texture at a given position (x, y) in the image I_w ; then we define a MEC as an image where

$$I_{\text{mec}}(x, y) = \arg \max (T(I_w(x, y))) \quad (9.5)$$

and the mask M_w as

$$M_w(x, y) = \begin{cases} 1 & \iff \arg \max (T(I_w(x, y))) \\ 0 & \text{otherwise} \end{cases} \quad (9.6)$$

9.7 Conclusions

We illustrated the design of a single-lens multispectral camera and used spatially registered multispectral iris images acquired from it to demonstrate that iris texture




Fig. 9.5 Image-level fusion in multispectral iris images. The first three images show the original and unwrapped distribution within each wavelength of the image patches exhibiting the highest texture (Eq. 9.5) responses across all the imaged wavelengths. The last three images are of the channel and structural masks (Eq. 9.6) and the MEC (Eq. 9.5) generated by applying these masks and fusing the regions with the highest information (i.e., maximum texture) from each wavelength to create a single enhanced image. The 800 nm NIR image is not shown, as a very *light blue* iris was used to illustrate the point that stronger texture is available outside of the NIR band, and for this particular iris, the texture elicited under VL was more pronounced at every point than under NIR.

increases with the frequency of the illumination for lighter-colored sections of the iris and decreases for darker sections. We presented a model for approximating an NIR iris image using features derived from the color and structure of a VL iris image. The simulated NIR images generated by this model are designed to improve the interoperability between legacy NIR iris images and those acquired under visible light by enabling cross wavelength matching of NIR and visible light iris images. Finally, we introduced the MEC, a novel iris code which uses pixel-level fusion algorithms to exploit texture variations elicited by illuminating the iris at different wavelengths to improve iris matcher performance and reduce FTE rates. We are currently evaluating the performance of both the VL to NIR transform and the MEC iris code and will report a comprehensive assessment in a follow-on paper.

References

1. Boyce, C., Ross, A., Monaco, M., Hornak, L., Li, X.: Multispectral iris analysis: a preliminary study. In: Proceedings of IEEE Computer Society Workshop on Biometrics. IEEE Computer Society, Washington, DC (2006)
2. Burger, W., Burge, M.J.: Digital Image Processing: An Algorithmic Introduction Using Java. Springer, New York (2007)
3. Daugman, J.G.: Complete discrete 2-D gabor transforms by neural networks for image analysis and compression. *IEEE Trans. Acoust. Speech Signal Process.* **36**(7), 1169–1179 (1988)
4. Daugman, J.: How iris recognition works. *IEEE Trans. CVSVT* **14**(1), 21–30 (2004)
5. Haralick, R.M., Shanmugam, K., Dinstein, I.: Textural features for image classification. *IEEE Trans. Syst. Man Cybern.* **3**(6), 610–621 (1973)
6. Jain, A., Bolle, R., Pankanti, S.: Introduction to biometrics. In: Jain, A., Bolle, R., Pankanti, S. (eds.) *Biometrics: Personal Identification in a Networked Society*. Kluwer, Boston (1999)
7. Monaco, M.: Color space analysis for iris recognition. MSEE Dissertation, West Virginia University (2007)
8. Sonka, M., Hlavac, V., Boyle, R.: *Image Processing, Analysis, and Machine Vision*. Thomson, Toronto (2007)

Chapter 10

Robust and Secure Iris Recognition

Jaishanker K. Pillai, Vishal M. Patel, Rama Chellappa, and Nalini K. Ratha

Abstract Iris biometric entails using the patterns on the iris as a biometric for personal authentication. It has additional benefits over contact-based biometrics such as fingerprints and hand geometry. However, iris biometric often suffers from the following three challenges: ability to handle unconstrained acquisition, privacy enhancement without compromising security, and robust matching. This chapter discusses a unified framework based on sparse representations and random projections that can address these issues simultaneously. Furthermore, recognition from iris videos as well as generation of cancelable iris templates for enhancing the privacy and security is also discussed.

10.1 Introduction

Iris recognition deals with using the patterns on the iris as a biometric for identification or verification. Iris is the colored region of the eye, which controls the amount of light entering the eye. It encloses the darker pupil and is surrounded by a white region called the sclera. Iris and the pupil are protected by a transparent membrane called the cornea. Figure 10.1 shows the different parts of the human eye.

Human iris typically has a rich set of patterns, which are believed to be formed randomly during the fetal development of the eye. The iris patterns are observed to differ for different individuals and remain stable for long periods of time, making it a good biometric. Iris-based human authentication deals with using these iris patterns

J.K. Pillai (✉) • V.M. Patel • R. Chellappa
University of Maryland, College Park, MD, USA
e-mail: jsp@umiacs.umd.edu; pvishalm@umiacs.umd.edu; rama@umiacs.umd.edu

N.K. Ratha
IBM Watson Research Center, Hawthorne, NY, USA
e-mail: Ratha@us.ibm.com




Fig. 10.1 Different parts of the eye



Fig. 10.2 Parts of an iris recognition system

to recognize individuals. It is a non contact biometric, which is more hygienic than a contact biometric like finger print. Also it is protected from wear and tear by the cornea, which ensures its stability over years. However, an iris-based personal authentication system requires close proximity of the eye and a cooperative subject. When acquired correctly, iris-based biometric is highly accurate.

10.1.1 Components of an Iris Recognition System

A generic iris recognition system has five main components, as shown in Fig. 10.2. They are the acquisition unit, segmentation unit, quality estimation unit, feature extraction unit, and matching unit. Since the most well-known iris recognition system is the one by Daugman [9], we use it as an example to explain the different steps below.

The first step in an iris recognition system is the image acquisition, where the iris image of the user is obtained. Near infrared illumination is normally used as it reveals the detailed structure of the iris much better than visible light. Visible light gets absorbed by the melanin pigments on the iris where as the longer near infrared radiations get reflected by them. After the image is obtained, the segmentation algorithm extracts the region of the image corresponding to the iris, for further processing. In Daugman's method, the inner and outer boundaries of the iris (known as the pupillary and limbic boundaries, respectively) are obtained by approximating them by circles and finding their parameters. Each of the boundaries is obtained for an image $I(x, y)$ by integrodifferential operators, which find the parameters (r, x_0, y_0) having the highest blurred partial derivatives in a circular arc of width ds along a circle of center (x_0, y_0) and radius r , given by

$$\max_{(r, x_0, y_0)} \left| G_{\sigma}(r) * \frac{\partial}{\partial r} \oint_{r, x_0, y_0} \frac{I(x, y)}{2\pi r} ds \right|$$

where $G_\sigma(r)$ denotes a smoothing operator like the Gaussian with scale σ and $*$ denote convolution. After obtaining these boundaries, further processing is performed to remove the eyelids and eyelashes present in the iris region. To make the iris template invariant to changes in size, the extracted iris region is mapped to a normalized coordinate system, where the pupillary boundary is mapped to zero and the limbic boundary is mapped to one. So any pixel in the new coordinate system is defined by an angle between 0° and 360° and a radial coordinate from 0 to 1.

The quality estimation stage estimates the quality of the acquired iris sample and decides whether to further process the image for recognition or reject the sample. Energy of the high-frequency Fourier coefficients can be used as a measure of blur in the acquired image. A blurred image will not have high-frequency components and hence have lower energy in those coefficients.

If the acquired image is accepted by the quality estimation module, feature extraction unit computes features from the iris region, which capture the identity of the person. The texture patterns on the iris contain the information unique to each individual. To capture this, the extracted and normalized iris region is convolved with two-dimensional Gabor filters of the form

$$G(r, \theta) = \exp^{-i\omega(\theta-\theta_0)} \exp^{-\frac{(r-r_0)^2}{\alpha^2}} \exp^{-\frac{(\theta-\theta_0)^2}{\beta^2}}$$

where (r, θ) denote the polar coordinates. (r_0, θ_0) denote the location, ω the scale and (α, β) the width of the Gabor filter, respectively.

To improve the speed of matching, the extracted Gabor features are encoded into a binary vector. Each complex Gabor feature is represented by 2 bits—the first bit is one if the real part of the Gabor feature is positive and zero otherwise. Similarly the second bit captures the sign of the imaginary part of the Gabor feature. Normalized Hamming distance is then used as a measure of similarity between two iris images. It is defined as the fraction of the bits for which the two iris codes differ. The distance is computed only for those pixels for which both the iris codes are not occluded. To account for the in-plane rotation, the distance is computed for several rotations of the iris vector. The smallest among all those scores is used as the final similarity score.

The source code for a typical generic iris recognition system, similar to the one described above, has been released by Masek and Kovesei [21]. We refer the reader to [5] for an excellent survey of the recent efforts in iris recognition.

10.1.2 Cancelable Iris Biometric

An important aspect in iris biometrics is security and privacy of the users. When the texture features of one's iris are stored in a template dictionary, a hacker could possibly break into the dictionary and steal these patterns. Unlike credit cards,

which can be revoked and reissued, biometric patterns of an individual cannot be modified. So, directly using iris features for recognition is extremely vulnerable to attacks. To deal with this, the idea of cancelable iris biometrics has been introduced [3, 15, 28], which can protect the original iris patterns as well as revoke and reissue new patterns when the old ones are lost or stolen.

The concept of cancelable biometrics was first introduced by Ratha et al. in [3, 28]. A cancelable biometric scheme intentionally distorts the original biometric pattern through a revocable and non-invertible transformation. The objectives of a cancelable biometric system are as follows [15]:

- Different templates should be used in different applications to prevent cross matching.
- Template computation must be non-invertible to prevent unauthorized recovery of biometric data.
- Revocation and reissue should be possible in the event of compromise.
- Recognition performance should not degrade when a cancelable biometric template is used.

In [10], hash functions were used to minimize the compromise of the private biometric data of the users. Cryptographic techniques were applied in [13] to increase the security of iris systems. In [17], error correcting codes were used for cancelable iris biometrics. A fuzzy commitment method was introduced in [16]. Other schemes have also been introduced to improve the security of iris biometric. See [10, 13, 15–17, 33] and the references therein for more details.

The pioneering work in the field of cancelable iris biometric was done by Zuo et al. [36]. They introduced four non-invertible and revocable transformations for cancelability. While the first two methods utilized random circular shifting and addition, the other two methods added random noise patterns to the iris features to transform them. As noted by the authors, the first two methods gradually reduce the amount of information available for recognition. Since they are essentially linear transformations on the feature vectors, they are sensitive to outliers in the feature vector that arise due to eyelids, eyelashes and specular reflections. They also combine the good and bad quality regions in the iris image leading to lower performance. The proposed random projection-based cancelability algorithm works on each sector of the iris separately, so outliers can only affect the corresponding sectors and not the entire iris vector. Hence, it is more robust to common outliers in iris data when compared to [36].

10.1.3 Iris Recognition from Videos

Though research in iris recognition has been extremely active in the past decade, most of the existing results are based on recognition from still iris images [14]. Multiple iris images have been used in the past to improve performance. Du [12] demonstrated higher rank one recognition rates by using three gallery images

instead of one. Ma et al. [20] also enrolled three iris images and averaged the three Hamming distances to obtain the final score. Krichen et al. [18] used the minimum of the three Hamming distance as the final score. Schmid et al. [31] demonstrated that fusing the scores using log-likelihood ratio gave superior performance when compared to average Hamming distance. Liu and Xie [19] and Roy and Bhattacharya [30] used multiple iris images for training classifiers.

The distortions common in iris image acquisition like occlusion due to eyelids, eyelashes, blur, and specular reflections will differ in various frames of the video. So by combining the different frames in the video efficiently, the performance could be improved. The only work toward using the temporal continuity in iris videos for improving the performance is done by Hollingsworth et al. [14]. The authors introduced a feature level fusion by averaging the corresponding iris pixels and a score level fusion algorithm combining all the pairwise matching scores. Though averaging reduces the noise and improves the performance, it required images to be well segmented and aligned, which may often not be possible in a practical iris recognition system.

10.1.4 Publicly Available Datasets

When the iris acquisition conditions are not constrained, many of the acquired iris images suffer from defocus blur, motion blur, occlusion due to the eyelids, specular reflections, and off-angle distortions. Numerous datasets are now available for the research community, containing varying degrees of these artifacts.

One of the earliest ones was the CASIA version 1 [32]. It contains 108 subjects and seven images per person. The images in the database are clean, without significant occlusion, blur, or specular reflection on the iris. But the pupil has been edited in these images and replaced by a black region, which masks the specular reflections in the pupil region. A larger dataset containing some of the unedited images is now available under the name CASIA version 3 [32]. It has 1,500 subjects and more than 2,200 iris images.

Two large datasets have been collected by the University of Notre Dame for the Iris Challenge Evaluation (ICE) in 2005 and 2006. The ICE 2005 dataset [22] has 244 subjects and close to 3,000 iris images in all. The ICE 2006 dataset [22] has 480 subjects and 60,000 images in all. A super set of the ICE 2005 and ICE 2006 datasets is now available called the ND-IRIS-04-05 dataset [4] (ND dataset). It contains iris images with occlusion due to shadows and eyelids, blur due to motion of the subject and defocus, specular reflections, and in-plane rotations. Some of these images are shown in Fig. 10.3. Gazed iris images with off-angle distortions are available in the West Virginia University iris dataset [34].




Fig. 10.3 Poorly acquired iris images from the ND dataset. Note that image (a) has specular reflections on the iris and is difficult to be segmented correctly due to the tilt and noncircular shape. Images (b) and (d) suffer from blurring, whereas image (c) is occluded by the shadow of the eyelids

10.2 Sparsity-Motivated Selection and Recognition of Iris Images

Some of the challenges in unconstrained iris recognition have been addressed recently using the ideas of sparse representations and random projections [26]. This method provides a unified framework for image quality estimation, recognition, and privacy in iris biometrics.

The idea is to create a dictionary matrix of the training samples as column vectors. The test sample is also represented as a column vector. Different dimensionality reduction methods are used to reduce the dimension of both the test vector and the vectors in the dictionary. It is then simply a matter of solving an ℓ_1 minimization problem in order to obtain the sparse solution. Once the sparse solution is obtained, it can provide information as to which training sample the test vector most closely relates to.

Let each image be represented as a vector in \mathbb{R}^n , A be the dictionary (i.e., training set), and y be the test image. The general sparse representation-based classification (SRC) algorithm is as follows [26, 35]:

1. Create a matrix of training samples $A = [A_1, \dots, A_C]$ for C classes, where A_i are the set of images of each class.
2. Reduce the dimension of the training images and a test image by any dimensionality reduction method. Denote the resulting dictionary and the test vector as \tilde{A} and \tilde{y} , respectively.
3. Normalize the columns of \tilde{A} and \tilde{y} .
4. Solve the following ℓ_1 minimization problem:

$$\hat{\alpha} = \arg \min_{\alpha'} \|\alpha'\|_1 \quad \text{subject to } \tilde{y} = \tilde{A}\alpha'. \quad (10.1)$$

5. Calculate the residuals

$$r_i(\tilde{y}) = \|\tilde{y} - \tilde{A}\delta_i(\hat{\alpha})\|_2,$$

for $i = 1, \dots, C$, where δ_i is a characteristic function that selects the coefficients associated with the i^{th} class.

6. Identify $(y) = \arg \min_i r_i(\tilde{y})$.

The assumption made in this method is that given sufficient training samples of the k th class, \tilde{A}_k , any new test image y that belongs to the same class will approximately lie in the linear span of the training samples from the class k . This implies that most of the coefficients not associated with class k in $\hat{\alpha}$ will be close to zero. Hence, $\hat{\alpha}$ is a sparse vector. Note that ℓ_1 -norm used to solve the above problem is an approximation to the ℓ_0 -norm. This approximation is necessary as solving the above problem with ℓ_0 -norm is not practical. Furthermore, it has been shown that if the solution is sparse enough and the dictionary A satisfies some properties, then solving the above problem with either ℓ_1 -norm or ℓ_0 -norm is equivalent. This algorithm can also be extended to deal with occlusions and random noise. Furthermore, a method of rejecting invalid test samples can also be introduced within this framework [35]. In particular, to decide whether a given test sample is a valid sample or not, the notion of sparsity concentration index (SCI) has been proposed [35]. The SCI of a coefficient vector α is defined as

$$SCI(\alpha) = \frac{C \cdot \max_i \|\delta_i(\alpha)\|_1 - 1}{\|\alpha\|_1 - 1}. \quad (10.2)$$

SCI takes values between 0 and 1. SCI values close to 1 correspond to the case when the test image can be approximately represented by using only images from a single class. Thus, the test vector has enough discriminating features of its class and hence has high quality. If $SCI = 0$, then the coefficients are spread evenly over all classes. Hence, the test vector is not similar to any of the classes and hence is of poor quality.

10.2.1 Sector-Based Recognition of Iris Images

Different regions of the iris have different qualities [8]. So instead of recognizing the entire iris image directly, the idea proposed in [26] is to recognize the different regions separately and combine the results depending on the quality of the region. This reduces the computational complexity of the above method as the size of the dictionary is greatly reduced, and the recognition of the different regions can be done in parallel. Also, since occlusions affect only local regions on the iris which can only lower the quality of certain regions, the robustness of the recognition algorithm to occlusion due to eyelids and eye lashes is improved. A direct way of doing this would be to recognize the sectors separately and combine the results by voting [25]. This, however, does not account for the fact that different regions are recognized with different confidences. To deal with this, a score level fusion approach for recognition was proposed in [26], where the recognition results of different sectors are combined based on the recognition confidence using the corresponding SCI values. Figure 10.4 illustrates the different steps involved in the proposed approach.

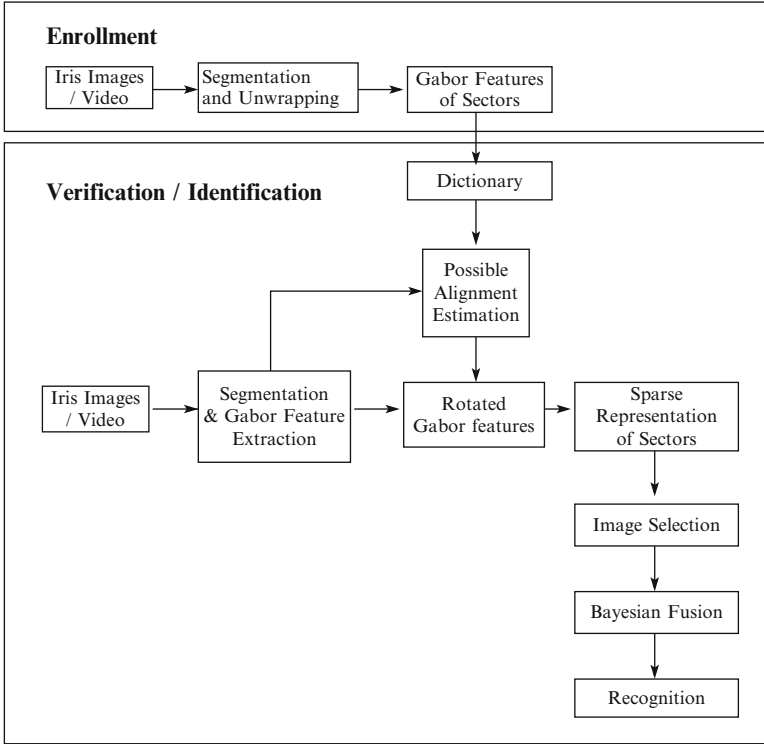


Fig. 10.4 A block diagram illustrating the Bayesian fusion-based image selection and recognition

The iris region is segmented using existing techniques and is divided into different sectors. Each sector is recognized separately using SRC. Here the Gabor features obtained from the sectors form the columns of the dictionary. SCI is used as a quality measure for each sector. Sectors with low SCI values are rejected.

A Bayesian fusion technique is then employed to combine the results from the remaining sectors, based on their quality. Let \mathbf{C} be the set of possible class labels and M be the number of sectors retained after rejection. Let d_1, d_2, \dots, d_M be the class labels of the retained sectors. The final class label is given by

$$\tilde{c} = \arg \max_{c \in \mathbf{C}} CSCI(c) \quad (10.3)$$

where

$$CSCI(c_l) = \frac{\sum_{j=1}^M SCI(d_j) \cdot \delta(d_j = c_l)}{\sum_{j=1}^M SCI(d_j)}. \quad (10.4)$$

CSCI of a class is the sum of the SCI values of all the sectors identified by the classifier as belonging to that class. Therefore, the optimal estimate is the class having the highest CSCI. Thus higher weighting is given to the labels predicted by

the better quality sectors when compared to the rest. To address the issue of in-plane rotation in the test iris image, the same formulation is extended by including sectors obtained after rotating the test iris image [26].

10.2.2 Recognition from Video

When test videos are available, the sectors of the different frames of the video can be combined based on their quality in a similar manner. Let $Y = \{y^1, y^2, \dots, y^J\}$ be the J -vectorized frames in the test video. As before, each frame is divided into \hat{M} sectors and recognized separately by the SRC algorithm. Let M_i be the number of sectors retained by the selection scheme in the i th frame. Let y_j^i be the j th retained sector in the i th frame. Similar to the sector-based recognition approach, one can derive the MAP estimate for video as

$$\tilde{c} = \arg \min_{c \in \mathbf{C}} \sum_{i=1}^J \sum_{j=1}^{M_i} \text{SCI}(d_j^i) \cdot \delta(c = d_j^i) \quad (10.5)$$

where d_j^i is the class label assigned by the classifier to y_j^i . Equation (10.5) can be alternatively written as

$$\tilde{c} = \arg \min_{c \in \mathbf{C}} \text{CSCI}(c)$$

where CSCI of a class c_l is given by

$$\text{CSCI}(c_l) = \frac{\sum_{i=1}^J \sum_{j=1}^{M_i} \text{SCI}(d_j^i) \cdot \delta(d_j^i = c_l)}{\sum_{i=1}^J \sum_{j=1}^{M_i} \text{SCI}(d_j^i)}.$$

As before, the MAP estimate consists of selecting the class having the highest cumulative SCI value, with the difference that the sectors of all the frames in the test video will be used while computing the CSCI of each class. Note that unlike existing feature level and score level fusion methods available for iris recognition, the CSCI incorporates the quality of the frames into the matching score. Hence, when the frames in the video suffer from acquisition artifacts like blurring, occlusion, and segmentation errors, the CSCI-based matching score gives higher weights to the good frames, at the same time, suppressing the evidence from the poorly acquired regions in the video. The different modes of operation of the proposed algorithm are illustrated in Fig. 10.5. Both the probe and the gallery can be separate iris images or iris videos. The iris images are segmented and unwrapped to form rectangular images. The Gabor features of the different sectors are computed, and SRC algorithm is used to select the good iris images. The good sectors are separately recognized and combined to obtain the class of probe image or video as described above.

for any v such that $\|v\|_0 \leq K$. When RIP holds, ΦA approximately preserves the Euclidean length of K -sparse vectors. When A is a deterministic dictionary and Φ is a random matrix, we have the following theorem on the RIP of ΦA .

Theorem 10.1 ([29]). *Let $A \in \mathbb{R}^{N \times (n.L)}$ be a deterministic dictionary with restricted isometry constant $\delta_K(A)$, $K \in \mathbb{N}$. Let $\Phi \in \mathbb{R}^{m \times N}$ be a random matrix satisfying*

$$P(\left| \|\Phi v\|^2 - \|v\|^2 \right| \geq \zeta \|v\|^2) \leq 2e^{-c\frac{n}{2}\zeta^2}, \quad \zeta \in \left(0, \frac{1}{3}\right) \quad (10.9)$$

for all $v \in \mathbb{R}^N$ and some constant $c > 0$ and assume

$$m \geq C\delta^{-2}(K \log((n.L)/K) + \log(2e(1 + 12/\delta))) + t \quad (10.10)$$

for some $\delta \in (0, 1)$ and $t > 0$. Then, with probability at least $1 - e^{-t}$, the matrix ΦA has restricted isometry constant

$$\delta_K(\Phi A) \leq \delta_K(A) + \delta(1 + \delta_K(A)). \quad (10.11)$$

The constant satisfies $C \leq 9/c$.

The above theorem establishes how the isometry constants of A are affected by multiplication with a random matrix Φ . Note that one still needs to check the isometry constants for the dictionary A to use this result. However, for a given dictionary, A , it is difficult to prove that A satisfies a RIP. One can empirically check the equivalence between the ℓ_0 -norm and ℓ_1 -norm by plotting the phase transition diagrams [2, 11].

The following are some matrices that satisfy (10.9) and hence can be used as random projections for cancelability:

- $m \times N$ random matrices Φ whose entries $\phi_{i,j}$ are independent realizations of Gaussian random variables $\phi_{i,j} \sim N(0, \frac{1}{m})$.
- Independent realizations of ± 1 Bernoulli random variables

$$\phi_{i,j} \doteq \begin{cases} +1/\sqrt{m}, & \text{with probability } \frac{1}{2} \\ -1/\sqrt{m}, & \text{with probability } \frac{1}{2}. \end{cases}$$

- Independent realizations of related distributions such as

$$\phi_{i,j} \doteq \begin{cases} +\sqrt{3/m}, & \text{with probability } \frac{1}{6} \\ 0, & \text{with probability } \frac{2}{3} \\ -\sqrt{3/m}, & \text{with probability } \frac{1}{6}. \end{cases}$$

- Multiplication of any $m \times N$ random matrix Φ with a deterministic orthogonal $N \times N$ matrix D , i.e., ΦD .

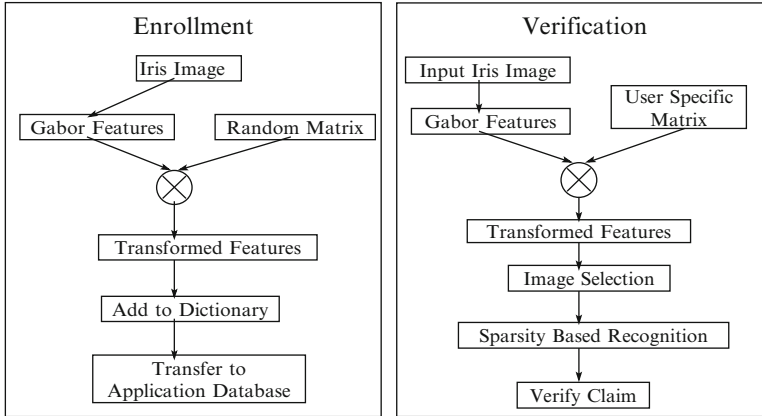


Fig. 10.6 Block diagram of the random projection-based cancelable system

Note that RPs meet the various constraints required for cancelability, mentioned in Sect. 10.1.2. By using different RP matrices, one can issue different templates for different applications. If a transformed pattern is compromised, one can reissue a new pattern by applying a new random projection to the iris vector. The RIP properties together with the sparsity of α ensure that the recognition performance is preserved. In the application database, only the transformed dictionary ΦA is stored. If a hacker illegally obtains the transformed dictionary ΦA and the transformed iris patterns of the user, a , he or she will have access to the person's identity. However, it is extremely difficult to obtain the matrix A from ΦA , and without A one cannot obtain the original iris patterns y . Hence, RP-based cancelable scheme is non-invertible as it is not possible to obtain the original iris patterns from the transformed patterns. Furthermore, since this method is based on pseudo-random number generation, one only needs to consider the state space corresponding to the value taken by the seed of the random number generator. Hence, instead of storing the entire matrix, only the seed used to generate the RP matrix needs to be stored (Fig. 10.6).

10.2.3.1 Random Permutation-Based Cancelable Iris Biometric

As discussed earlier, when the iris image has good quality, only the training images corresponding to the correct class will have high coefficients. If the training images of different classes are randomly arranged as columns of the dictionary, both the dictionary and the order of the training images are required for correct recognition. This can be used to enhance the security of the proposed iris recognition system [26].

When a new user is enrolled, his or her training images are divided into sectors and placed at random locations in the dictionary. In Fig. 10.7, we show the dictionary for a trivial example of four users. Note that the different sectors of each training image of the user are kept at different random locations in the dictionary. Without prior knowledge of these locations, it is impossible to perform recognition.

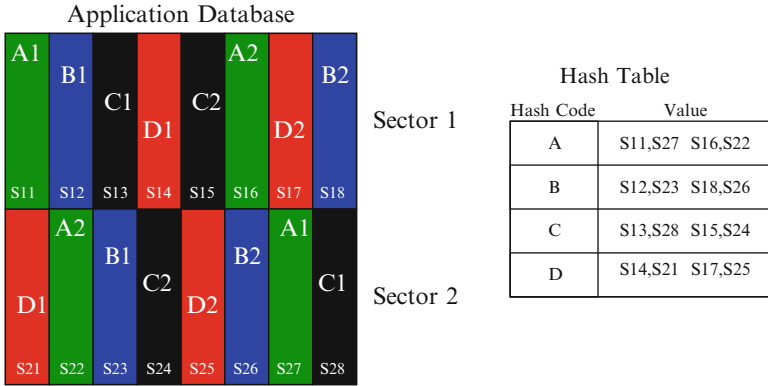


Fig. 10.7 Sample dictionary and hash table for a four-user example. The four users *A*, *B*, *C*, and *D* are indicated by colors *green*, *blue*, *black*, and *red*, respectively. *A1* and *A2* are the two training images corresponding to the first user. *S_{ij}* denote that the *jth* location and the *ith* sector. *D1* at *S14* means that the first sector of the user *D* is at location *S14*

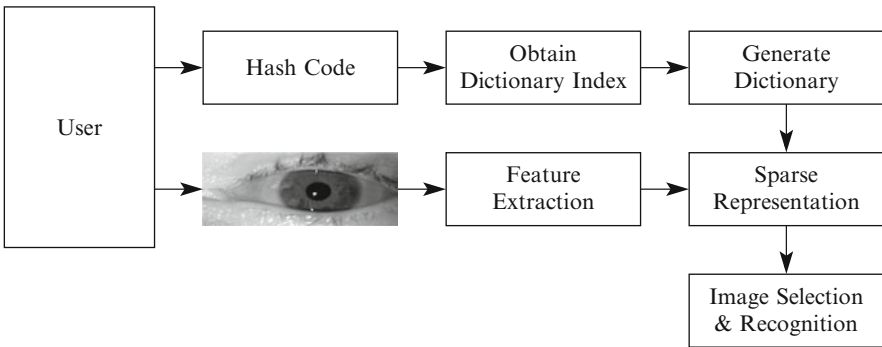


Fig. 10.8 Block diagram of the proposed cancelability scheme using random permutations

An array indicating the column numbers of the training images of the correct class is generated for each user. This array is stored in a hash table, and the corresponding hash code is given to the user during enrollment. During verification, the system acquires the iris image of the person and extracts the features. For each sector of the iris vector, the sparse coefficients are obtained using this shuffled dictionary. The user also has to present the hash code to the system. Using the hash code, the indices of training images are obtained from the hash table, and the coefficients belonging to different classes are grouped. Then, SCI is computed and used to retain or reject the images. If the image is retained, the CSCI values of the different classes are computed, and the class having the lowest CSCI value is assigned as the class label of the user. A block diagram of the security scheme is presented in Fig. 10.8.

If the hash code presented is incorrect, then the obtained indices of the training images for each class will be wrong. So the coefficients will be grouped in a wrong way, and all the classes will have similar energy leading to a low SCI value and the subsequent rejection of the image. Even if by chance, one of the classes happened to have high energy and the image is retained, the probability of that class being the correct class is very low ($\frac{1}{N}$). Thus, with high probability, the user will not be verified. Hence, if a hacker illegally acquires the iris patterns of a legitimate user, without having the hash code, he or she will not be able to access the system. Also, even if the hacker obtains the iris dictionary stored in the application database, the iris patterns of the user cannot be accessed without knowing the correct hash codes, because different sectors of an iris patterns reside at different random locations. If the hash code is compromised, the dictionary indices of the user can then be stored at a new location, and a new hash code can be issued to the user. Also, different applications can have different dictionaries. Thus, the user will have a different hash code for each application, preventing cross matching.

It should be noted that the additional security and privacy introduced by these techniques come at the expense of storing additional seed values. In applications requiring higher security, this can be stored with the user, so that a hacker will not get the original templates even if he or she gets hold of the cancelable patterns in the template database. For applications with greater emphasis on usability, the seed can be stored securely in the template database, so that the user will not have to carry it.

10.3 Experimental Results

To show the effectiveness of the sparsity-promoting methods for iris recognition, we highlight some of the results presented in [26] on the ICE2005 dataset [22], ND-IRIS-0405 (ND) dataset [4], and the MBGC videos [24].

To illustrate the robustness of the algorithm to occlusion due to eyelids and eyelashes, a simple iris segmentation scheme was performed, detecting just the pupil and iris boundaries and not the eyelids and eyelashes. Publicly available code of Masek and Kovese [21] for detecting these boundaries was used. To study the variation of SCI in the presence of common distortions during image acquisition like occlusion and blur, they were simulated on the clean iris images from the ND dataset. Fifteen clean iris images of the left eye of eighty persons were selected. Twelve such images per person formed the gallery, and distortions were simulated on the remaining images to form the probes. Seven different levels of distortion for each case were considered with level one indicating no distortion and level seven indicating maximum distortion. The dictionary was obtained using the gallery images, and evaluated the SCI of the various sectors of the test images.

Figure 10.9 shows some of the simulated images from the ND dataset. The first column includes images with distortion level one (no distortion). The middle column contains images with distortion level three (moderate distortions). The right most column contain images with distortion level five (high distortion). The first




Fig. 10.9 Simulated distortions on the images from the ND dataset. The detected pupil and iris boundaries are indicated as *red circles*

row contains images with blur, while the second contains images with occlusion. Images with simulated segmentation error and specular reflections are shown in the third and fourth rows, respectively.

Figure 10.10 illustrates the variation of SCI with the common acquisition distortions. It can be observed that good images have high SCI values whereas the ones with distortion have lower SCI values. So by suitably thresholding the SCI value of the test image, one can remove the bad images before the recognition stage. The relatively lower decrease in SCI with occlusion and specular reflection demonstrates the increased robustness attained by the proposed algorithm by separately recognizing the individual sectors and combining the results, as discussed earlier.

Fig. 10.10 Plot of the variation in SCI values with common distortions in iris image acquisition. Note that the SCI falls monotonically with increasing levels of blur and segmentation errors in the iris images. It is also robust to occlusions and specular reflections




Table 10.1 Recognition rate on the ND dataset

Image quality	NN	Masek’s implementation	Proposed method
Good	98.33	97.5	99.15
Blurred	95.42	96.01	98.18
Occluded	85.03	89.54	90.44
Seg. error	78.57	82.09	87.63

10.3.1 Recognition Performance

In the first set of experiments, 80 subjects were selected from the ND dataset. Fifteen clean images of the left iris were hand picked for each person. Of these fifteen images per person, 12 were randomly selected to form the gallery, and the remaining three images per person were used as probes. Image selection was not performed for this experiment.

The proposed algorithm was compared to a nearest neighbor-based recognition algorithm (NN) that uses the Gabor features and the Masek’s implementation. Since we use tough segmentation conditions retaining the eyelids and eye-lashes in the iris vector, direct application of NN and Masek’s method produced poor results. For fair comparison, the iris images were divided into different sectors, obtained the results using these methods separately on each sectors and combined the results by voting. Recognition rate of 99.15% is achieved compared to 98.33% for the NN and 97.5% for the Masek’s method.

To evaluate the recognition performance of the proposed algorithm on poorly acquired images, images with blur, occlusion, and segmentation errors were hand picked. Fifteen clean images per person were used to form the gallery. Probes containing each type of distortion were applied separately to the algorithm. Image selection was performed followed by recognition. The recognition rates are reported in Table 10.1.

Table 10.2 Verification rate at an FAR of 0.001 on the ICE 2005 dataset

Method	Verification rate (%)
Pelco	96.8
WVU	97.9
CAS 3	97
CAS 1	97.8
CMU	99.5
SAGEM	99.8
Proposed method	98.13

In the second set of experiments, we used images from the ICE 2005 dataset. The proposed algorithm is compared with the existing results on the ICE 2005 dataset corresponding to Experiment 1. Experiment 1 has 1,425 iris images corresponding to 120 different classes. Ten images per class were used in the gallery, and remaining iris images were used as the test images. Segmentation was performed using Masek’s code, and the Gabor features of the segmented iris images were applied to the proposed algorithm. No image selection was performed. The performance of the proposed algorithm was compared with the existing results [23] in Table 10.2, where the verification rates are indicated at a false acceptance rate of 0.001.

10.3.2 Cancelability Results

We present the cancelability results on the clean images from the ND dataset. The iris region obtained after segmentation was unwrapped into a rectangular image of size 10×80 . The real part of the Gabor features was obtained and concatenated to form an iris vector of length 800. The random Gaussian matrix was used to generate random projections. In [27], it was shown that the application of the random projections separately on each sector performed better performance when compared to the application of a single random projection on the entire iris vector. The real part of the Gabor features of each sector of the iris image is first vectorized, random projections were then applied, and finally the random projected vectors are concatenated to obtain the cancelable iris biometric. Either the same random Gaussian matrix for all the users or different random matrices for different users were applied to obtain the RP “same matrix” and “different matrix” vectors, respectively. Having obtained the random vectors from the Gabor features of the iris image, the sparsity-based recognition algorithm was performed.

Figure 10.11a plots the receiver operating characteristic (ROC) for the iris images in the ND dataset for the original and transformed iris patterns. As demonstrated, using different matrices for each class performs better than using the same matrix for all classes. In the “different matrix” case, it is assumed that the user provided the correct matrix assigned to him or her. So the performance exceeds even the original performance as class-specific random projections increases the interclass distance, still retaining the original intra-class distance. In Fig. 10.11b,

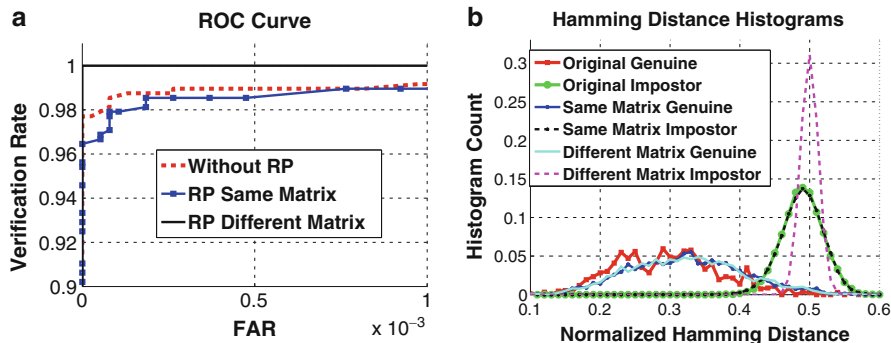


Fig. 10.11 (a) ROC for the ND dataset. The same matrix performance is close to the performance without cancelability. Using different matrices for each class gives better performance. (b) Comparison of the distribution of the genuine and impostor normalized Hamming distances for the original and transformed patterns

Table 10.3 Statistics of the normalized Hamming distance

Methods	Mean	Standard deviation
Without RP	0	0
Same matrix	0.5002	0.0123
Different matrix	0.4999	0.013
Dictionary permutations	0.4913	0.0254

the distribution of the genuine and impostor normalized Hamming distance for the original and transformed iris patterns is compared. It can be seen that the distribution of the genuine Hamming distance remains almost the same after applying the random projections. The original and same matrix cases have similar impostor Hamming distance distributions. However, the different matrix case has an impostor distribution that is more peaked and farther from the genuine distribution, indicating superior performance.

Table 10.3 provides the statistics of the normalized Hamming distance between the original and the transformed iris vectors. As can be seen, the mean of the normalized Hamming distance is very close to 0.5 with a very low standard deviation.

To evaluate the performance of the proposed cancelable method using dictionary permutations, the following three possible scenarios on the clean images from the ND dataset are considered. In the first case, the user provides the iris image and the correct hash code. In this case, the recognition performance was the same as that of the original method on the ND dataset, which is 99.17%. In the second case, the user provides the iris image but a wrong hash code. Here the recognition performance dropped to 2%, which is only slightly better than chance. This is equivalent to the case when a hacker illegally obtains the iris image of a valid user and tries to gain access into the system with a guess about the hash code. The low recognition performance clearly reflects the additional security introduced by the permutations, as a hacker needs to now have not only the iris image but also the hash code of a valid user to gain access. In the third experiment, the closeness between




Fig. 10.12 Plot of the histograms of the normalized Hamming distance between the original and transformed vectors. Note that the histogram peaks around 0.5 indicating that the original and transformed iris codes are significantly different

the Gabor features of the original iris images and the new feature vectors obtained by permutations of the Gabor features in the dictionary is found. As before, the normalized Hamming distance between the iris codes obtained from these vectors is used as the measure of similarity. The histogram of the normalized Hamming distance between the original and the randomly permuted iris vectors is shown in Fig. 10.12. The mean and standard deviation of the Hamming distance histogram are indicated in the last row of the Table 10.3. Note that the mean is close to 0.5, indicating that the permutations differ significantly different from the original iris images. Even if a hacker can use the dictionary from the application database, he or she will be unable to extract information about the original iris images without knowing the hash code of each user.

10.3.3 Results on Iris Video

In this section, the results on the MBGC videos [24] are presented [26]. Given the 30 classes, 28 classes that contained at least five good images were used in the experiments. Five clean images from the iris videos in the training set were hand-picked which formed the dictionary. In the test videos, batches of five frames were given as a probe to the algorithm. Using 28 available videos and 60 frames from each test video, one could form 336 probes. Only a basic segmentation of the iris and pupil using the Masek's code was performed. Also, the poorly segmented iris images were not removed before performing the recognition.




Fig. 10.13 ROC plots for video-based iris recognition. Method 1 treats each frame in the video as a different probe. Method 2 averages all the frames in the probe video. Methods 3 and 4 use the average and minimum of all the pairwise Hamming distance between the frames of the probe and gallery videos, respectively. The proposed method uses CSCI as the matching score. Note that the introduced quality based matching score outperforms the existing fusion schemes, which do not incorporate the quality of the individual frames in the video

The performance of the proposed algorithm was compared with four other methods. The ROC plots for the different methods are displayed in Fig. 10.13. In method 1, each frame of the video is considered as a different probe. It gave the worst performance, indicating that using multiple frames available in a video can improve the performance. Method 2 averages the intensity of the different iris images. Though it performs well when the images are clean, a single image which is poorly segmented or blurred could affect the entire average. In methods 3 and 4, all possible pairwise Hamming distances between the video frames of the probe videos and the gallery videos belonging to the same class are computed. Method 3 uses the average of these Hamming distance as the score. In Method 4, the minimum of the pairwise Hamming distance was used as the score. In the proposed method, the CSCI values were computed for each class for each probe video, and the probe video is assigned to the class having the highest CSCI value. For a fair comparison of the proposed quality measure in videos, none of the frames were rejected. The proposed method performed better than other methods. One of the reasons for the superior performance could be the fact that the algorithm is incorporating the quality of the different frames while computing the CSCI. Frames which are poorly segmented or blurred will have a low SCI value and hence will not affect the score significantly. In all the other methods, the image quality was not effectively incorporated into the matching score, so all frames are treated equally irrespective of their quality.

10.4 Discussion and Conclusion

In this chapter, we introduced a general framework for iris image selection and recognition based on sparse representation and random projections. The introduced method has numerous advantages over existing techniques when sufficient number of training data are available. It was shown that the selection algorithm can handle common distortions in iris image acquisition such as occlusions, blur, and segmentation errors. Also, a quality-based matching score was introduced and its effectiveness was demonstrated on the MBGC iris video dataset. Finally, random projection and random permutation were incorporated into the algorithm to prevent the compromise of sensitive biometric information of the subjects.

Acknowledgements This work was partially supported by a MURI grant N00014-08-1-0638 from the Office of Naval Research.

References

1. Baraniuk, R.: Compressive sensing. *IEEE Signal Process. Mag.* **24**(4), 118–121 (2007)
2. Blanchard, J.D., Cartis, C., Tanner, J.: The restricted isometry property and ℓ_q -regularization: phase transition for sparse approximation (preprint)
3. Bolle, R.M., Connel, J.H., Ratha, N.K.: Biometrics perils and patches. *Pattern Recognit.* **35**(12), 2727–2738 (2002)
4. Bowyer, K.W., Flynn, P.J.: The ND-IRIS-0405 iris image dataset. nd.edu/~cvrl/papers/ND-IRIS-0405.pdf
5. Bowyer, K.W., Hollingsworth, K., Flynn, P.J.: Image understanding for iris biometrics: a survey. *Comput. Vis. Image Underst.* **110**(2), 281–307 (2008)
6. Candes, E., Romberg, J., Tao, T.: Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory* **52**(2), 489–509 (2006)
7. Candes, E., Romberg, J., Tao, T.: Stable signal recovery from incomplete and inaccurate measurements. *Commun. Pure Appl. Math.* **59**, 1207–1223 (2006)
8. Chen, Y., Dass, S.C., Jain, A.K.: Localized iris image quality using 2-d wavelets. In: *Proceedings of International Conference on Biometrics, Hong Kong*. Springer Lecture Notes in Computer Science, vol. 3832, pp. 373–381 (2006)
9. Daugman, J.: High confidence visual recognition of persons by a test of statistical independence. *IEEE Trans. Pattern Anal. Mach. Intell.* **15**, 1148–1161 (1993)
10. Davida, G.I., Frankel, Y., Matt, B.J.: On enabling secure applications through off-line biometric identification. In: *IEEE Symposium on Security and Privacy, Oakland*, pp. 148–157 (1998)
11. Donoho, D.L.: High-dimensional centrally symmetric polytopes with neighborliness proportional to dimension. *Discret. Comput. Geom.* **35**(4), 617–652 (2006)
12. Du, Y.: Using 2d log-gabor spatial filters for iris recognition. *SPIE Biometric Technology for Human Identification III* **6202**, pp. 62020F1–62020F8 (2006)
13. Hao, F., Anderson, R., Daugman, J.: Combining crypto with biometrics effectively. *IEEE Trans. Comput.* **55**(9), 1081–1088 (2006)
14. Hollingsworth, K.P., Bowyer, K.W., Flynn, P.J.: Image averaging for improved iris recognition. *Lect. Notes Comput. Sci.* **5558**, 1112–1121 (2009)
15. Jain, A.K., Nandakumar, K., Nagar, A.: Biometric template security. *EURASIP J. Adv. Signal Process. Spec. Issue Biom.* **2008**(113), 1–17 (2008)

16. Juels, A., Wattenberg, M.: A fuzzy commitment scheme. In: ACM Conference on Computers and Communications Security, Singapore, pp. 28–36 (1999)
17. Kanade, S., Petrovska-Delacretaz, D., Dorizzi, B.: Cancelable iris biometrics and using error correcting codes to reduce variability in biometric data. *IEEE Comput. Vis. Pattern Recognit.* 120–127 Miami, Florida (2009)
18. Krichen, E., Allano, L., Garcia-Salicetti, S., Dorizzi, B.: Specific texture analysis for iris recognition. *Lect. Notes Comput. Sci.* **3546**, 23–30 (2005). Springer, Heidelberg
19. Liu, C., Xie, M.: Iris recognition based on DLDA. In: International Conference on Pattern Recognition, Hong Kong, pp. 489–492 (2006)
20. Ma, L., Tan, T., Wang, Y., Zhang, D.: Efficient iris recognition by characterizing key local variations. *IEEE Trans. Image Process.* **13**(6), 739–750 (2004)
21. Masek, L., Kovesi, P.: Matlab source code for a biometric identification system based on iris patterns (2003) www.csse.uwa.edu.au/~pk/studentprojects/libor/sourcecode.html
22. Newton, E.M., Phillips, P.J.: Meta-analysis of third-party evaluations of iris recognition. *IEEE Trans. Syst. Man Cybern.* **39**(1), 4–11 (2009)
23. Phillips, P.J.: Frgc and ice workshop (2006) biometrics.nist.gov/cs_links/iris
24. Phillips, P.J., Flynn, P.J., Beveridge, J.R., Scruggs, W.T., O’Toole, A.J., Bolme, D.S., Bowyer, K.W., Draper, B.A., Givens, G.H., Lui, Y.M., Sahibzada, H., Scallan, J.A., Weimer, S.: Overview of the multiple biometrics grand challenge. In: The Third IAPR/IEEE Conference on Biometrics. *Advances in Biometrics, Lecture Notes in Computer Science #5558*, pp. 705–714 Springer-Verlag Berlin, Heidelberg (2009)
25. Pillai, J.K., Patel, V.M., Chellappa, R.: Sparsity inspired selection and recognition of iris images. The Third IAPR/IEEE Conference on Biometrics, Alghero (2009)
26. Pillai, J.K., Patel, V.M., Chellappa, R., Ratha, N.K.: Secure and robust iris recognition using random projections and sparse representations. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(9), 1877–1893 (2011)
27. Pillai, J.K., Patel, V.M., Chellappa, R., Ratha, N.K.: Sectored random projections for cancelable iris biometrics. In: International Conference on Acoustics, Speech, and Signal Processing, Taipei (2009)
28. Ratha, N.K., Connel, J.H., Bolle, R.: Enhancing security and privacy in biometrics-based authentication systems. *IBM Syst. J.* **40**(3), 614–634 (2001)
29. Rauhut, H., Schnass, K., Vandergheynst, P.: Compressed sensing and redundant dictionaries. *IEEE Trans. Inf. Theory* **54**(5), 2210–2219 (2008)
30. Roy, K., Bhattacharya, P.: Iris recognition with support vector machines. In: International Conference on Biometrics, Hong Kong, pp. 486–492 (2006)
31. Schmid, N.A., Ketkar, M.V., Singh, H., Cukic, B.: Performance analysis of iris based identification system at the matching score level. *IEEE Trans. Inf. Forensics Secur.* **1**(2), 154–168 (2006)
32. Specification of CASIA Iris Image Database (ver 1.0), Chinese Academy of Sciences. <http://www.nlpr.ia.ac.cn/english/irids/irisdatabase.htm>. (2007)
33. Teoh, A., Goh, A., Ngo, D.: Random multispace quantization as an analytic mechanism for biohashing of biometric and random identity inputs. *IEEE Trans. Pattern Anal. Mach. Intell.* **28**(12), 1892–1901 (2006)
34. West virginia university iris dataset. www.csee.wvu.edu/~xinl/demo/nonideal_iris.html
35. Wright, J., Yang, A.Y., Ganesh, A., Sastry, S.S., Ma, Y.: Robust face recognition via sparse representation. *IEEE Trans. Pattern Anal. Mach. Intell.* **31**(2), 210–227 (2009)
36. Zuo, J., Ratha, N., Connell, J.: Cancelable iris biometric. In: Proceedings of International Conference on Pattern Recognition, Tampa, pp. 1–4 (2008)

Chapter 11

Template Aging in Iris Biometrics

Sarah E. Baker, Kevin W. Bowyer, Patrick J. Flynn, and P. Jonathon Phillips

Abstract Using a data set with approximately 4 years of elapsed time between the earliest and most recent images of an iris (23 subjects, 46 irises, 6,797 images), we investigate template aging for iris biometrics. We compare the match and non-match distributions for short-time-lapse image pairs, acquired with no more than 120 days of time lapse between them, to the distributions for long-time-lapse image pairs, with at least 1,200 days of time lapse. We find no substantial difference in the non-match, or impostor, distribution between the short-time-lapse and the long-time-lapse data. We do find a difference in the match, or authentic, distributions. For the image data set and iris biometric systems used in this work, the false reject rate increases by about 50% or greater for the long-time-lapse data relative to the short-time-lapse data. The magnitude of the increase in the false reject rate varies with changes in the decision threshold and with different matching algorithms. Our results demonstrate that iris biometrics is subject to a template aging effect.

11.1 Introduction

The term “template aging” refers to degradation of biometric performance that occurs with increased time between the acquisition of an enrollment image and acquisition of the image compared to the enrollment. Template aging effects are known to exist for biometrics such as face and fingerprint [7, 15, 21, 23, 27].

The iris biometrics community has long accepted the premise that the iris is “essentially stable” throughout a person’s life and that this means that template

S.E. Baker (✉) • K.W. Bowyer • P.J. Flynn
University of Notre Dame, Notre Dame, IN, USA
e-mail: sbaker3@cse.nd.edu; kwb@cse.nd.edu; flynn@cse.nd.edu

P.J. Phillips
National Institute of Standards and Technology, Gaithersburg, MD, USA
e-mail: jonathon@nist.gov

aging does not occur for iris biometrics. Daugman stated the core assumption this way – “As an internal (yet externally visible) organ of the eye, the iris is well protected and stable over time” [8]. This assumption is commonly repeated in research publications dealing with iris biometrics: “(the iris is) stable over an individual’s lifetime” [25], “(the iris is) essentially stable over a lifetime” [18], “the iris is highly stable over a person’s lifetime” [19]. The commercial iris biometrics literature explicitly connects this to the idea of lifetime enrollment – “only a single enrollment in a lifetime” [24].

Note that claims about stability of the iris texture and “lifetime enrollment” are never presented as dependent on the particular sensor, algorithm, length of time lapse, or any other condition. They are presented as universal claims about iris biometrics in general. Thus, a single counterexample is sufficient to disprove the universal claim.

It is well known in the medical literature that the eye and iris undergo a variety of changes with age [2, 5, 6, 14, 29, 30]. Any of these effects could in principle alter details of the imaged iris texture. It is also possible that a template aging effect could be due to aging of the sensor, changes in how a person uses the biometric system, or other factors. The essential question for iris biometrics is – does the quality of a match between two images of the same iris change with increased time between the enrollment image and the image to be recognized? That is, does a template aging effect exist? We present results of the first systematic investigation of this question.

We use an image data set involving 23 persons (46 irises) with approximately 4 years of time lapse between the earliest and latest images of a given iris. We consider image pairs in a short-time-lapse group, representing no more than 120 days of time lapse between the two images, and in a long-time-lapse group, representing at least 1,200 days of time lapse. We experiment with three iris biometric systems: our modification of the IrisBEE baseline matcher [22], Neurotechnology’s VeriEye system [28], and the Cam-2 submission to the Iris Challenge Evaluation 2006 [20]. We find that, for each of the three systems, there is no significant difference in the non-match, or “impostor,” distributions between the short-time-lapse and the long-time-lapse data. We also find that, for each of the three systems, the match distribution for the long-time-lapse data is different from that for the short-time-lapse data in a way that results in an increased false reject rate. Thus, we observe clear evidence of a template aging effect for iris biometrics.

11.2 Previous and Related Work

We do not know of *any* experimental study that supports the conclusion that template aging does not occur for iris biometrics. Claims about the stability of iris texture appear to be based on subjective human visual perception of iris texture in visible-light images of the iris. However, it has been shown that humans are able to perceive similarities in iris texture that do not result in closer iris biometric matches [13]. Thus, human perception of the general iris texture pattern does not automatically or necessarily imply anything about iris biometric operation.

Gonzalez et al. [26] report an effect of time lapse on iris recognition that may initially seem similar to our results. However, Gonzalez et al. compare matches between images acquired at the same acquisition session with those acquired with at most 3 months time lapse. They report a better match statistic for images from the same session than for those across sessions. However, they show little change in match statistics when comparing matches with short time lapses, between 2 weeks and 3 months. In our results presented here, we do not consider matches between images acquired in the same acquisition session, as we expect that this is not representative of a real-world biometric scenario. We expect that “same session” images will generally result in atypically good matches. Like Gonzalez et al., we do not find any significant difference in match scores for images with a few months time lapse. However, when considering a longer time lapse than that examined in Gonzalez et al., we do observe a statistically significant degradation in match scores.

This chapter expands upon our initial results [3] in several ways. First, we have increased the number of subjects from 13 to 23 and the number of irises from 26 to 46. Second, in [3] we only considered images from spring 2004 and spring 2008 and the matches within one semester and matches across the 4 years. In this work, we now consider all images acquired from 2004 through 2008 and have set two time thresholds in defining our short-time-lapse and long-time-lapse matches. Third, we have tested the time lapse effect on two additional iris biometric algorithms: Neurotechnology’s VeriEye [28] and the Cam-2 submission to the Iris Challenge Evaluation 2006 from the University of Cambridge [20]. We also test for various possible causes of match score degradation with increased time lapse. Finally, we present ROC curves for short-time-lapse and long-time-lapse matches for each of the three algorithms and explicitly show the difference in the false reject rates.

11.3 Image Data Set and Algorithms

All of the iris images used in this study were acquired with the same LG 2200 iris imaging system [16], located in the same studio throughout the 4 years of image acquisition. The system had no hardware or software modifications during the 4 years. The LG 2200 model is now discontinued. However, current state-of-the-art iris imaging systems of course did not exist at the time that data acquisition for this experiment started. We are currently pursuing additional work with images acquired using a newer model sensor, and initial results [10] are generally consistent with results of this study.

Image acquisition sessions were held at multiple times in each academic semester across the 4 years. At a given acquisition session, for a given subject, six images were acquired of each eye. The image acquisition protocol was the same as that used in the Iris Challenge Evaluation (ICE) 2005 and 2006 [20, 22]. However, it is important to note that while the protocol for the ICE acquisitions allowed for some images that did not pass the normal built-in quality control checks of the LG 2200 [20], *all images used in this study were manually screened for image quality.*




Fig. 11.1 Example iris images of a subject taken in 2004 and 2008 (subject 04233). *Upper left:* right iris from 2004; *upper right:* right iris from 2008; *lower left:* left iris from 2004; and *lower right:* left iris from 2008

Images of noticeably poor quality were excluded from this study, e.g., out-of-focus irises, major portions of the iris occluded, and obvious interlace artifacts were all excluded. Also, images that resulted in a noticeably poor iris segmentation by the IrisBEE algorithm were excluded from the study. (The detailed segmentation was not available from the other systems.)

A total of 23 persons participated in data acquisitions from 2004 through 2008. See Fig. 11.1 for examples of iris images. There are images from both irises of the 23 subjects over the 4 years. Subject age ranges from 22 to 56 years old at the end of the 4-year period. Sixteen subjects are male and seven are female. Sixteen subjects are Caucasian and seven are Asian. The repeated 16 by 7 breakdown is a coincidence; the ethnicity division does not follow the gender division. None of the subjects wore glasses for any of the data acquisition. Five subjects wore contact lenses at all acquisition sessions, and 18 subjects did not wear contact lenses at any acquisition session. The total number of iris images selected for use in this study was 6,797.

We created two sets of image pairs, a short-time-lapse set and a long-time-lapse set. The short-time-lapse set consists of image pairs where the two images were acquired with no more than 120 days of time lapse between them. The average time lapse in this group is 44 days. The long-time-lapse set consists of image pairs

acquired with no less than 1,200 days of time lapse. The average time lapse in this group is 1,405 days. A given iris image can participate in multiple short-time-lapse pairs and multiple long-time-lapse pairs.

11.4 Iris Matching Algorithms

To investigate the generality of any observed effects, three different iris biometric algorithms were included in the study. First, we used our own modified version of the IrisBEE system distributed through the ICE program [20]. This system represents an iris as a $240 \times 10 \times 2$ -bit iris code generated from the complex-valued responses of one-dimensional log-Gabor wavelet filters applied to the normalized iris image [17]. For the IrisBEE matcher, the output of matching two iris images is a fractional Hamming distance. The range of the fractional Hamming distance is $[0, 1]$, with zero being a perfect match and 0.5 a random level of match. Second, we used the commercial VeriEye 2.2 Iris SDK from Neurotechnology [28]. This system produces match scores on a different scale and with a different polarity than systems employing fractional Hamming distance. For the analysis in this chapter, we negated the match scores so that lower scores represented better matches. The third system was the Cam-2 submission to the ICE 2006 from the University of Cambridge [20]. The output of the Cam-2 matcher is nominally a fractional Hamming distance. Thus, we have used three different algorithms. One is based on a “baseline” source code that was made available to the research community, one is a readily available commercial product, and one was a best performer in the ICE 2006 results.

11.5 False Reject Rates for Short and Long Time Lapse

We computed the authentic and impostor distributions for each of the three algorithms. The impostor distributions showed no apparent difference between the short-time-lapse data and the long-time-lapse data. However, the authentic distributions for long-time-lapse data were shifted in the direction of the impostor distribution. For each of the three algorithms, the shift in the authentic distribution is such that it causes an increase in the false reject rate (FRR) for any practical choice of decision threshold.

Graphs that zoom in on the “tails” of the long-time-lapse and short-time-lapse authentic distributions for each algorithm are shown in Fig. 11.2. These graphs show the tails of the distributions across a range of possible values for the decision threshold. Recall that for the IrisBEE and Cam-2 algorithms, a smaller value (of fractional Hamming distance) represents a better match, while for the VeriEye algorithm, a larger value of different units represents a better match.

This figure shows that for all three algorithms, across a broad range of possible threshold values, *the long-time-lapse authentic distribution has a higher false reject rate than the short-time-lapse authentic distribution*. The IrisBEE algorithm shows

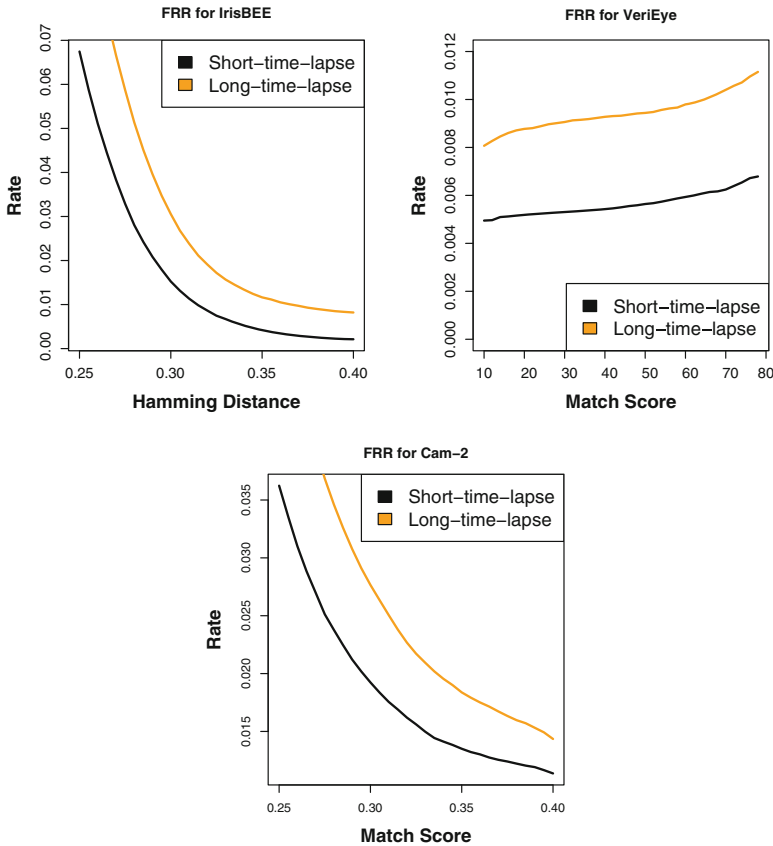


Fig. 11.2 Authentic distributions across a range of match scores, showing increased false reject rates

approximately 150% increase in the false reject rate across the range of decision thresholds, the VeriEye algorithm shows an approximately 70% increase, and the Cam-2 algorithm shows an approximately 40% increase. Thus, we observe clear and consistent evidence of a template aging effect for each of three algorithms considered in this study.

11.6 Frequency of Authentic Distribution with Worse Mean Score

We also performed a one-sided sign test to check for statistical significance of the frequency, across the 46 irises, of the long-time-lapse authentic distribution having a worse mean match score than the short-time-lapse authentic distribution. A worse

Table 11.1 Sign test for frequency of worse mean match score with longer time lapse

Algorithm	No. irises	Test statistic	p-value
IrisBEE	42	5.75	2.55×10^{-9}
VeriEye	41	5.46	2.20×10^{-8}
Cam-2	38	4.57	4.62×10^{-6}

Table 11.2 Overlap in number of irises for which the mean of the long-time match scores is greater than the mean for the short-time match scores. The overlap is reported for all combinations of the three algorithms and for all three algorithms

Algorithms	N of 46 irises in common
IrisBEE-VeriEye	38
IrisBEE-Cam2	35
VeriEye-Cam2	35
All three	34

mean score is one closer to the impostor distribution. If time lapse has no effect, then we would expect that the long-time-lapse mean is worse for half of the irises and the short-time-lapse mean is worse for half. This is the null hypothesis for the test. The sign test does not make any distributional assumptions about the means of similarity scores. The one-sided test was selected because we are interested in the alternative hypothesis that the longer-time-lapse data has a larger mean score.

The sign test results are presented in Table 11.1, including the test statistic, p-value, and number of irises for which the mean of the long-time match scores is worse than the mean of the short-time-lapse match scores ($\mu_L(i) > \mu_S(i)$). The results show that we can easily reject the null hypothesis for all three algorithms. The frequency of a worse match score occurring for the long time lapse is statistically significant. This indicates that the increased FRR seen in Fig. 11.2 is not the result of a small number of unusual irises in the data set but is characteristic of the data set in general.

Table 11.1 shows that for IrisBEE, there are 42 of 46 irises for which the long-time-lapse mean HD is worse; for VeriEye, there are 41 irises for which the long-time-lapse mean match score is worse; and for Cam-2, there are 38 irises for which the long-time-lapse mean HD is worse. One natural question is as follows: how many of these irises are in common? The answers are presented in Table 11.2, which shows the number of irises in common. The last row reports that 34 irises have the time-lapse effect for all three algorithms. A one-sided sign test for 34 of 46 irises showing an effect across all three algorithms produces a test statistic of 3.391 with a p-value of 8.207×10^{-4} . Thus, even if we use the criteria that all three algorithms must agree on the movement of the means, the null hypothesis is rejected.

11.7 Possible Causes of an Increased False Reject Rate

We considered a variety of factors that could conceivably contribute to causing the observed result. For example, it is known that the presence of contact lenses can adversely affect match quality [4]. If the short-time-lapse data contained image pairs where a subject did not wear contact lenses and the long-time-lapse data contained image pairs where the same subject was wore contacts, this could conceivably cause an increased FRR for long time lapse relative to short time lapse. Similarly, if a person was wearing the same type of contacts in short-time-lapse image pairs, but a different type in long-time-lapse image pairs, this could conceivably cause an increased FRR.

We manually checked for the presence of contact lenses in all images included in this study. We found that each subject in this study either wore contacts for all acquisition sessions or did not wear contacts to any acquisition session. Also, for the subjects who wore contacts, none appear to have changed the type of contacts worn. Thus, we conclude that the wearing of contact lenses is not an appreciable factor in our observed results.

Hollingsworth et al. [12] showed that the degree of the pupil dilation, and the difference in pupil dilation between two images, can affect the match distribution. We performed an analysis of the changes in pupil dilation and its possible effect on the difference between long-time-lapse and short-time-lapse data.

The first step in the analysis was to compute the ratio of the pupil diameter to the iris diameter for each image. The second step was to compute the difference in the pupil-to-iris ratio for the iris images in each match pair. Then, for each subject, we computed the average change in the pupil-to-iris ratio over all short-time-lapse match pairs. We denote this by $\rho_S(i)$. Similarly, we computed the average change in the pupil-to-iris ratio for all long-time match pairs, denoted by $\rho_L(i)$. Then for each iris, we computed the difference between the average short-time-lapse change in the pupil-to-iris ratio and the average long-time-lapse change in the pupil to iris ratio, denoted by $\rho_L(i) - \rho_S(i)$. For the IrisBEE algorithm, we created a scatterplot of the change in the pupil-to-iris ratio between long time lapse and short time lapse match pairs and change in match score between long time lapse and short-time-lapse. Figure 11.3 is a scatterplot of $\mu_L(i) - \mu_S(i)$ versus $\rho_L(i) - \rho_S(i)$. The corresponding Kendall correlation coefficient is 0.217. If the observed increase in false reject rate could be attributed to a change in pupil dilation, then $\mu_L(i) - \mu_S(i)$ versus $\rho_L(i) - \rho_S(i)$ would be substantially correlated. If $|\rho_L(i)| > |\rho_S(i)|$, then there is a greater difference in diameters of the pupils for long-time match pairs than for short-time match pairs. In turn, this implies that match scores should degrade. However, our analysis shows minimal correlation between $\mu_L(i) - \mu_S(i)$ versus $\rho_L(i) - \rho_S(i)$. Thus, we conclude changes in pupil dilation are not an appreciable factor in our observed result.

The percentage of an iris that is occluded can affect iris matching performance [9]. The more of the iris that is observable, the better the expected performance. Thus, one possible factor contributing to the observed increase in the




Fig. 11.3 Scatterplot of the change in match score between long time and short time lapse for each iris versus the change in the pupil-to-iris ratio between long-time- and short-time-lapse match pairs ($\mu_L(i) - \mu_S(i)$ versus $\rho_L(i) - \rho_S(i)$). The *horizontal axis* is the change in mean match scores for the long-time- and short-time lapse iris pairs. The *vertical axis* is the change in the average short-time change in the pupil-to-iris ratio and the average long-time change in the pupil-to-iris ratio. Each *red circle* is an iris

false reject rate is that the percentage of the iris that is observable decreased in the long-time-lapse data relative to the short-time-lapse data.

In the IrisBEE algorithm [20], the fraction of the iris that is visible is indicated by the fraction of the iris code bits that are marked in the iris code mask as representing non-occluded portions of the iris. To determine if there is a change over time in the fraction of the iris that is occluded, we divided the time period over which the data was collected for this study into 30-day intervals. We computed the average number of bits marked as non-occluded in the mask for all images collected in each 30-day interval. We then computed Kendall’s correlation coefficient between the average number of bits marked as non-occluded and time. The resulting Kendall’s correlation coefficient is -0.131 . This indicates that there is no substantial correlation between number of bits marked as non-occluded and elapsed time. Thus, we conclude that change in the amount of iris occluded does not account for the increase in the false reject rate observed in our results.

The iris images in the time-lapse study were collected with the same LG 2200 sensor [16]. It is conceivable that the sensor properties of the LG 2200 could have changed over time in such a way as to cause an increased false reject rate in the long-time-lapse data. To test for this, in fall 2008 we collected iris images with a second rarely used LG 2200 camera. We collected approximately 3,000 images from 77 subjects (154 irises) who attended three separate acquisition sessions (labeled “session one,” “session two,” and “session three”). There was approximately 2 weeks elapsed time between each session. During sessions one and three, iris images were collected with the original camera; during session two, the iris images were collected with the second rarely used camera. The first step in our sensor aging analysis was to compute the match and non-match score distributions between iris images collected in session one and session three, both sessions using the original

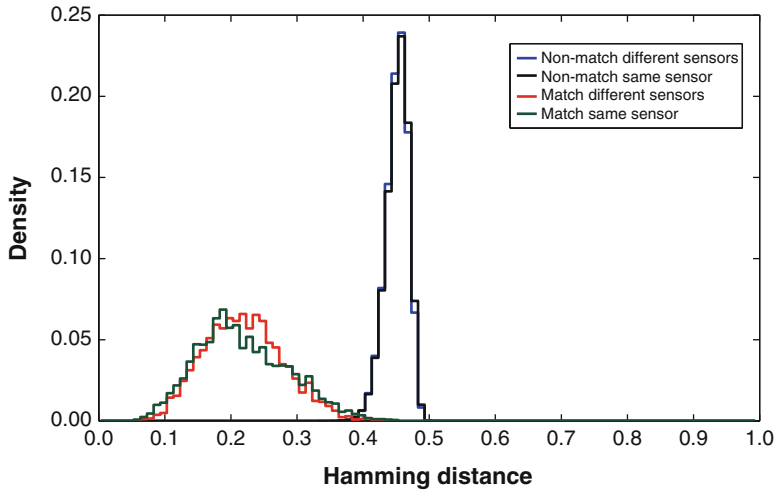


Fig. 11.4 The match and non-match distributions for the within and between sensors experiments. The match and non-match distributions are for the Hamming distance from the IrisBEE algorithm. The mean Hamming distance for match scores collected with the same sensor is 0.2153 and for match scores collected with different sensors is 0.2167. The mean Hamming distance for non-match scores collected with the same sensor is 0.4483 and for non-match scores collected with different sensors is 0.4478

sensor. The second step was to compute the match and non-match score distributions between iris images collected in session one and session two. In session two, the images were collected with the second rarely used sensor. If the sensor age affects match quality, we would expect a significant degradation in match scores between images collected from the two different sensors compared to image pairs collected with the original sensor. The average match score for image pairs collected with the original sensor is 0.215; the average match score for image pairs collected with the two different sensors was 0.217. Figure 11.4 shows a histogram for the match and non-match distributions for both within and between sensor comparisons. Based on this analysis, we conclude that a sensor aging effect cannot account for the increase in false reject rate that is seen in our results.

The LG 2200 camera actively illuminates the iris using three infrared light emitting diodes (LED) positioned on the left, right, and top of the sensor. When acquiring images, the camera is designed to take three images, one with each LED. In commercial applications, the camera will save the best quality image and discard the other two. For our acquisitions, the system had the capability to save all three images (for a detailed explanation see Phillips et al. [20, 22]). It is conceivable that if there were more matches between images acquired with the same LED in the short-time-lapse group, and more matches between images acquired with different LEDs in the long-time-lapse group, this could result in an increased false reject rate for the long-time-lapse group.

We grouped the matches into those in which the two images were taken with the same LED and those in which the two images were taken with different LEDs. For both groups, we observed an increased false reject rate of about 50% across all feasible decision threshold values for the long-time-lapse data over the short-time-lapse data. Thus, we conclude that variations in the particular LED illuminating the images are not the cause of the increased false reject rate seen in our results.

11.8 Conclusions and Discussion

For three different matching algorithms, and across the range of practical decision threshold values for each matching algorithm, we found that the false reject rate increases with longer time lapse between enrollment and verification. This is seen clearly in the difference in the tails of the authentic distributions. Also, the frequency of irises with a worse mean match score for long time lapse compared to short time lapse is statistically significant. Thus, our experimental results show clear and consistent evidence of a template aging effect for iris biometrics. The magnitude of the template aging effect varies between algorithms, with the value of the decision threshold, and other factors.

We were able to test for a variety of factors that could potentially contribute to observing an increased false reject rate with increased time lapse. We concluded that factors such as varying pupil dilation, wearing of contact lenses, differences in amount of iris occluded, and sensor aging are not an appreciable factor in our experimental results.

It is possible that the template aging effect observed in our experimental results is caused by normal aging of the eye. One well-known example of age-related change in the normal eye involves pupil size. Winn et al. studied factors affecting light-adapted pupil size and found that “of the factors investigated, only chronological age had a significant effect on the size of the pupil” [29]. They concluded “the results of this study are consistent with previous reports suggesting that pupil size becomes smaller in an almost linear manner with increasing age” [29]. The iris, of course, controls the pupil size, and so this change in average pupil size reflects a change in the functioning of the iris tissue. As the Merck Manual of Geriatrics describes it, “The iris comprises two sets of muscles that work together to regulate pupillary size and reaction to light. With aging, these muscles weaken and the pupil becomes smaller (more miotic), reacts more sluggishly to light, and dilates more slowly in the dark” [5].

There are also age-related changes in the melanocytes, the cells that produce melanin, in the iris. Eye color is largely determined by the melanocytes in the anterior layer of the iris. For some segments of the population, aging can lead to a noticeable change in the melanocytes and so the eye color. Bito et al. report that “Most individuals had stable eye color after early childhood. However, there was a subpopulation of white subjects with eye color changes past childhood. Approximately 17% of twins and 11% of mothers experienced a change in eye

color of 2 U or more. (...) Thus, eye color, and hence, iridial pigmentation, seems to change in some individuals during later years” [6]. They found that the changes in eye color were more similar for identical twins than fraternal twins, indicating a genetic link to this particular element of aging. One element of melanocyte aging can, in rare cases, lead to a cancer. “The melanocytes in the iris are constantly exposed to UV radiation, and this leads to the malignant transformation of these cells to form a specific type of malignant tumor, the uveal melanoma” [14].

Also connected with the melanocytes, iris freckles and nevi can arise in the iris and can grow over time. “Iris freckles are the most common iris tumors found in children as well as adults. They are collections of benign, but abnormal melanocytes that vary in size and shape. Although congenital, they tend to become more prominently pigmented with age. Iris freckles are clusters of normal melanocytes and have no malignant potential. Nevi efface the iris architecture and may cause clinical structural alterations. . .” [30].

In addition, it is known that the cornea undergoes age-related changes. “The shape and aberrations of the cornea change with age. It is well known that the radius of curvature slightly decreases with age, and the asphericity also changes. On average, the cornea becomes more spherical with age and, as a consequence, spherical aberrations tend to increase” [1]. The iris is imaged through the cornea; thus, corneal changes may affect iris images.

Small, incremental changes in imaged iris texture over time should be considered normal, as “. . . age related changes take place in all ocular tissues of the human eye. . .” [2]. The relevant question for iris biometrics is the timescale at which normal aging has an appreciable effect on the biometric template computed from the imaged iris texture. To underscore this point, we quote from the Flom and Safir iris recognition patent [11] – “The basic, significant features of the iris remain extremely stable and do not change over a period of many years. Even features which do develop over time, such as the atrophic areas discussed above, usually develop rather slowly, so that an updated iris image will permit identification for a substantial length of time.” In this quote, it is clear that Flom and Safir anticipated the possibility that small, incremental changes in iris texture could potentially result in the need for an “updated image” and reenrollment of the iris template. One interpretation of our results is that they confirm that the possibility that Flom and Safir envisioned is in fact true.

In an attempt to identify the regions of the iris that changed, degrading the match quality, we visually examined the iris images. Visual examination of the iris image pairs with the poorest match scores for the IrisBEE algorithm revealed no drastic or obvious changes in the irises or their textures. This suggests that, if the template aging effect is due to normal aging of the eye, humans may not be able to easily perceive the subtle changes that are involved.

Much additional research remains to be done in the area of template aging for iris biometrics. While we have experimentally observed a template aging effect, and have ruled out several factors as primary causes of the observed effect, we have not conclusively identified a primary cause of the observed template aging. It is important to understand the cause of the observed template aging effect, so

that techniques can be developed to mitigate the effect. It would also be valuable to know whether or not iris biometric template aging is constant across different demographic groups and whether it occurs at a faster or slower rate as a person ages. Studies that collect new and larger data sets, involve a larger pool of subjects, different sensors, a longer time period, and/or a sample of subjects that represent a greater range of demographics would all be important.

Acknowledgements SEB, KWB, and PJF were supported by the National Science Foundation under grant CNS01-30839, by the Central Intelligence Agency, by the Intelligence Advanced Research Projects Activity, and by the Technical Support Working Group under US Army contract W91CRB-08-C-0093. PJP acknowledges the support of the Biometric Task Force, the Department of Homeland Security's Directorate for Science and Technology, the Intelligence Advanced Research Projects Activity (IARPA), the Federal Bureau of Investigation (FBI), and the Technical Support Working Group (TSWG).

The opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of our sponsors. The identification of any commercial product or trade name does not imply endorsement or recommendation by the authors, the University of Notre Dame, or the National Institute of Standards and Technology.

References

1. Artal, P.: Chapter 3: aging effects on the optics of the eye. In: Cavallotti, C.A.P., Cerulli, L. (eds.) *Age-Related Changes of the Human Eye*. Humanan Press, Totowa (2008)
2. Atchison, D., Markwell, E., Kasthurirangan, S., Pope, J., Smith, G., Swann, P.: Age-related changes in optical and biometric characteristics of emmetropic eyes. *J. Vis.* **8**(4), 1–20 (2008)
3. Baker, S., Bowyer, K., Flynn, P.: Empirical evidence for correct iris match score degradation with increased time-lapse between gallery and probe matches. In: *Proceedings of the Third International Conference on Biometrics*, Alghero, pp. 1170–1179 (2009)
4. Baker, S., Hentz, A., Bowyer, K.W., Flynn, P.J.: Degradation of iris recognition performance due to non-cosmetic prescription contact lenses. *Comput. Vis. Image Underst.* **114**(9), 1030–1044 (2010)
5. Beers, M.H. (ed.): Chapter 126: aging and the eye. *The Merck Manual of Geriatrics*, 3rd edn. <http://www.merck.com/mkgr/mm/sec15/sec15.jsp> (2008)
6. Bito, L., Matheny, A., Cruickshanks, K., Nondahl, D., Carino, O.: Eye color changes past early childhood. *Arch. Ophthalmol.* **115**, 659–663 (1997)
7. Carls, J., Raines, R., Grimaila, M., Rogers, S.: Biometric enhancements: template aging error score analysis. In: *8th IEEE International Conference on Automatic Face and Gesture Recognition*, Amsterdam, pp. 1–8 (2008)
8. Daugman, J.: How iris recognition works. *IEEE Tran. Circuit Syst. Video Technol.* **14**(1), 21–30 (2004)
9. Daugman, J.: New methods in iris recognition. *IEEE Trans. Syst. Man Cybern.* **37**(5), 1167–1175 (2007)
10. Fenker, S., Bowyer, K.W.: Experimental evidence of a template aging effect in iris biometrics. In: *IEEE Computer Society Workshop on Applications of Computer Vision*, Kona (2011)
11. Flom, L., Safir, A.: Iris recognition systems. US Patent No 4641394 (1987)
12. Hollingsworth, K., Bowyer, K.W., Flynn, P.J.: Pupil dilation degrades iris biometric performance. *Comput. Vis. Image Underst.* **113**(1), 150–157 (2009)
13. Hollingsworth, K., Bowyer, K.W., Flynn, P.J.: Similarity of iris texture between identical twins. In: *Computer Vision and Pattern Recognition Biometrics Workshop*, San Francisco (2010)

14. Hu, D.: Photobiology of the uveal tract. In: Smith, K.C. (ed.) *Photobiological Sciences Online*. <http://www.photobiology.info/> (1987)
15. Lanitis, A.: A survey of the effects of aging on biometric identity verification. *Int. J. Biom.* **2**(1), 34–62 (2010)
16. Lg: <http://www.lgiris.com/> (2009). Accessed April 2009
17. Liu, X., Bowyer, K.W., Flynn, P.: Experiments with an improved iris segmentation algorithm. In: *Proceedings of the Fourth IEEE Workshop on Automatic Identification Technologies*, Buffalo, pp. 118–123 (2005)
18. Miyazawa, K., Ito, K., Aoki, T., Kobayashi, K., Nakajima, H.: An effective approach for iris recognition using phase-based image matching. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(10), 1741–1756 (2008)
19. Monro, D., Rakshit, S., Zhang, D.: Dct-based iris recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **4**(29), 586–595 (2007)
20. Phillips, P.J., Scruggs, W.T., O’Toole, A.J., Flynn, P.J., Bowyer, K.W., Schott, C.L., Sharpe, M.: FRVT 2006 and ICE 2006 large-scale experimental results. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(5), 831–846 (2010)
21. Phillips, P., Grother, P., Michaels, R., Blackburn, D., Tabassi, E., Bone, M.: *Face Recognition Vendor Test 2002: Overview and Summary*. National Institute of Technology and Standards, Gaithersberg, Maryland (2000). http://www.nist.gov/manuscript-publication-search.cfm?pub_id=50767
22. Phillips, P., Bowyer, K., Flynn, P., Liu, X., Scruggs, T.: The iris challenge evaluation 2005. In: *Proceedings of the Second IEEE Conference on Biometrics: Theory, Applications, and Systems*, Washington, DC (2008)
23. Ryu, J., Jang, J., Kim, H.: Analysis of effect of fingerprint sample quality in template aging. In: *NIST Biometric Quality Workshop II*, pp. 7–8. National Institute of Technology and Standards, Gaithersberg, Maryland (2007). http://www.nist.gov/itl/iad/ig/bio_quality_wkshopii.cfm
24. LGE Iris Tech Win In India Redefines Biometric Scalability: <http://www.findbiometrics.com/article/115> (2009). Accessed April 2009
25. Thornton, J., Savvides, M., Kumar, V.: A bayesian approach to deformed pattern matching of iris images. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(4), 596–606 (2007)
26. Tome-Gonzalez, P., Alonso-Fernandez, F., Ortega-Garcia, J.: On the effects of time variability in iris recognition. In: *Proceedings of the Second IEEE Conference on Biometrics: Theory, Applications and Systems*, Washington, DC (2008)
27. Uludag, U., Ross, A., Jain, A.: Biometric template selection and update: a case study in fingerprints. *Pattern Recognit.* **37**, 1533–1542 (2008)
28. VeriEye Iris Recognition Technology. <http://www.neurotechnology.com/verieye.html> (2008). Accessed Nov 2008
29. Winn, B., Whitaker, D., Elliot, D., Phillips, N.: Factors affecting light-adapted pupil size in normal human subjects. *Investig. Ophthalmol. Vis. Sci.* **35**(3), 1132–1137 (1994)
30. Wright, K., Spiegel, P.: *Pediatric Ophthalmology and Strabismus*, p. 438. Springer, New York (2003)

Chapter 12

Fusion of Face and Iris Biometrics

Ryan Connaughton, Kevin W. Bowyer, and Patrick J. Flynn

Abstract This chapter presents a system which simultaneously acquires face and iris samples using a single sensor, with the goal of improving recognition accuracy while minimizing sensor cost and acquisition time. The resulting system improves recognition rates beyond the observed recognition rates for either isolated biometrics.

12.1 Introduction

The practice of using more than one biometric modality, sample, sensor, or algorithm to achieve recognition, commonly referred to as *multi-biometrics*, is a technique that is rapidly gaining popularity. By incorporating multi-biometrics into the recognition process, many of the shortcomings of traditional single-biometric systems can be alleviated, and overall recognition accuracy can be improved. Multi-biometrics can inherently increase system robustness by removing the dependency on one particular biometric approach. Further, a system that utilizes more than one biometric feature or matcher may be more difficult to deliberately spoof [17]. Systems that make use of multiple biometric features can also provide redundancy that may lower failure-to-acquire rates. Though multi-biometrics offers many potential advantages over traditional biometric systems, inefficient system design can greatly increase sensor cost, computation time, and data acquisition time.

While research into multi-biometrics has received a large increase in attention over recent years, the task of fusing multiple biometric modalities from a single sensor remains an understudied challenge. Due to a lack of available multimodal data, many current experiments in multi-biometrics create “chimeric” datasets, in

R. Connaughton (✉) • K.W. Bowyer • P.J. Flynn
University of Notre Dame, Notre Dame, IN, USA
e-mail: rconnaug@nd.edu; kwb@cse.nd.edu; flynn@cse.nd.edu

which samples of one biometric modality from one set of subjects are arbitrarily paired with a second biometric modality from a separate set of subjects in order to simulate a multi-biometric scenario [1]. This approach, though useful for preliminary experimentation, may mask unknown dependencies between modalities. Further, chimeric datasets simulate a multi-biometric scenario in which samples of each modality are acquired independently. In practice, it is much more desirable to simultaneously acquire multiple modalities from a single sensor if possible for cost and usability reasons.

This chapter presents a system which simultaneously acquires face and iris samples using a single sensor, with the goal of improving recognition accuracy while minimizing sensor cost and acquisition time. The resulting system improves recognition rates beyond the observed recognition rates for either isolated biometric.

12.2 Characteristics of Multi-biometric Systems

The term *multi-biometrics* encompasses a wide range of fusion techniques, and its precise meaning is somewhat inconsistent in the literature [1, 16, 21]. In the simplest, traditional single-biometric system, one sensor images a particular body part (i.e., iris, face, or fingerprint) to produce a single image. The image is then processed and matched against a gallery using a specific algorithm to obtain a verification or identification result. A multi-biometric system aims to improve recognition rates (or address some other drawbacks of traditional systems) by providing redundancy at one or more of the steps in this recognition process.

In general, there are five types of multi-biometric systems [15]:

1. *Multi-sample*: Multi-sample systems collect and process multiple images of the same biometric. Such systems benefit from some of the advantages of multi-biometrics while minimizing sensor cost.
2. *Multi-instance*: Similar to multi-sample, multi-instance systems collect and process images of several distinct instances of the same biometric trait. Examples of multi-sample systems include systems that consider multiple fingerprints or both irises for recognition. Alternatively, multi-instance systems may collect multiple images of the same trait with some controlled variation; for example, a system may collect face images with smiling and neutral expressions.
3. *Multi-sensor*: A multi-sensor system images the same biometric trait using more than one sensor. Multi-sensor systems may be considered implicitly multi-sample as well. The incorporation of multiple sensors naturally leads to an increase in system cost, but this approach may help to address a particular bias or shortcoming in a specific sensor by obtaining a cross-sensor consensus.
4. *Multi-algorithm*: Multi-algorithm systems use more than one matching algorithm on the same biometric sample and then fuse the results to improve system performance. Because this approach can make use of the same biometric sample for each matcher, multi-algorithm systems can be cost-effective and help to reduce algorithmic biases.

5. *Multi-modal*: Multimodal systems consider more than one biometric trait, or modality, in the recognition process. Ko [7] suggests that multi-modal fusion benefits the most when the biometric modalities are orthogonal. Modalities can be considered orthogonal when the match performance of one modality does not predict the performance of the other. In the ideal scenario, all of the biometrics would be orthogonal, simultaneously imaged with the same sensor, and captured at high quality.

While these five classifications can be used to describe many multi-biometric approaches, there are naturally some systems which represent hybrids of more than one multi-biometric approach. Nonetheless, it is useful to have some method of categorizing multi-biometric systems, and understanding the advantages and disadvantages associated with each approach is crucial to good system design.

12.3 Levels of Fusion

In multi-biometric systems, the term *fusion* is often used to describe the process of combining information from more than one source in the recognition process. The previous section described the stages at which multi-biometric systems may use redundancy to improve performance; fusion is used to combine the results of the redundancy so that a single output can be produced. There are five levels at which fusion can occur in a multi-biometric system:

1. *Signal level*: Using signal-level fusion, multiple samples may be combined together to create one superior sample. An example of signal fusion is a super-resolution technique which combines multiple images of the same iris to achieve a higher-quality image.
2. *Feature level*: In a system that uses feature-level fusion, matching features are first extracted from each biometric sample, and fusion is used to condense all of the features into a single biometric signature.
3. *Score level*: With score-level fusion, the match scores are combined to produce a final result. Examples include a multi-sample approach in which each sample is matched separately and the resulting scores are fused or a multi-algorithm approach in which the same sample is matched using multiple matchers and the results of all of matchers are combined.
4. *Rank level*: Similar to score-level fusion, rank-level fusion combines match rankings, rather than the actual scores, into a final ranking to determine the best match.
5. *Decision level*: Decision-level fusion applies a matcher to each biometric sample (or the same matcher to multiple samples) to obtain a Boolean response indicating whether or not each comparison is a match. The outputs are then fused using Boolean operators, a voting scheme, or some similar method.

It has been suggested in the literature that systems which incorporate fusion at an early stage of the recognition process (e.g., signal- or feature-level fusion) have the potential to be more effective than systems which use fusion later in the pipeline [17]. Despite this, many researchers believe that score-level fusion offers the best trade-off between potential performance gain and ease of implementation [11].

12.4 Related Work

The fusion of face and iris modalities is a biometric approach that has gained increasing attention over the past decade, likely due to the popularity of the individual modalities as well as the natural connection between them. Despite this recent trend, very few studies have been done on fusion of face and iris biometrics from a single sensor.

The most common method of multi-biometric fusion is score-level fusion. Zhang et al. approach the problem of fusing face and iris biometrics under near-infrared lighting using a single sensor [24]. Frontal face images are acquired using a 10-megapixel CCD camera. Eye detection and face alignment are performed using Local Binary Pattern histogram matching as described in Li et al. [9]. The eigenface algorithm and Daugman's algorithm are used to perform face and iris recognition, respectively, and score-level fusion is accomplished via the sum and product rules after min-max normalization. Numerous other score-level fusion approaches have been tested on chimeric datasets. Chen and Te Chu use an unweighted average of the outputs of matchers based on neural networks [4]. Wang et al. test weighted average, linear discriminant analysis, and neural networks for score-level fusion [22].

Another common approach to biometric fusion is feature-level fusion through concatenation. Rattani and Tistarelli compute SIFT features for chimeric face and iris images and concatenate the resulting feature vectors [14]. The number of matching SIFT features between two vectors (measured by Euclidean distance) is used as a match score for that comparison. Son and Lee extract features for face and iris images based on a Daubechies wavelet transform [18]. Concatenation is used to form a joint feature vector, and Euclidean distance between feature vectors is used to generate match scores.

The Multiple Biometrics Grand Challenge (MBGC) provided a collection of face and iris data to researchers in order to provide a standard test bed for comparing matching and fusion techniques [12, 13]. The MBGC data included a subset of the near-infrared videos used in the experiments being presented in this chapter, as well as face stills, high-quality color face video, iris stills, and iris video. In general, results showed that fusion of face and iris biometrics offered improved accuracy over either biometric alone. The near-infrared videos released as part of the MBGC are also used by Yang et al. [23]. Yang et al. investigate the use of SIFT features to perform alignment between the partial faces present in the dataset in order to facilitate face matching, but do not incorporate these results into a multi-biometric experiment.

The work presented in this chapter differs from previous work in the fusion of face and iris biometrics in several facets. First, this chapter uses only genuine multimodal data, rather than chimeric data for experimentation. Additionally, the fusion is accomplished using a single sensor. Though Zhang et al. also use a single sensor, the authors also manually acquire each image to guarantee high-quality face and iris samples. In the experiments presented in this chapter, an on-the-move and at-distance sensor is used to acquire data for a high-throughput scenario. The resulting dataset consists of a much wider range of sample quality with incomplete data for some subjects, making the dataset a practical but challenging test bed for fusion experiments. These experiments also differ from work presented on the MBGC data; the near-infrared videos used in the MBGC dataset were manually selected to guarantee the presence of a subject in the field of view, whereas in the experiments shown in this chapter, this process is done automatically. Finally, this work uses multi modal, multi-sample, and multi-instance approaches to improve system accuracy and robustness.

12.5 Approach

To facilitate the fusion of face and iris biometrics from a single sensor, the Iris on the Move (IOM) sensor was selected for data acquisition. The IOM, shown in Fig. 12.1, is a sensor designed for high-throughput stand-off iris recognition [10]. The IOM features a portal which subjects walk through at normal walking pace.



Fig. 12.1 Picture of the Iris on the Move sensor designed by Sarnoff Corporation. The IOM was used for all probe data collection (Picture reprinted from [2] with permission from Elsevier)



Fig. 12.2 Example of corresponding frames from the IOM as the subject passes through the in-focus region of the portal. The *left* image shows a frame from the *top* camera, the *middle* image shows a frame from the *middle* camera, and the *right* shows a frame from the *bottom* camera

As a subject passes through the portal, the subject is illuminated with near-infrared (NIR) LEDs, and frontal video is captured by an array of three vertically arranged, fixed-focus cameras equipped with NIR filters. The presence of multiple cameras allows the system to handle a larger range of subject heights, and the sensor can be extended to include more than three cameras to support an even larger range of subject heights. Though the sensor is intended for iris image acquisition, the face is typically captured as well. While the sides of the portal help to direct subjects into the field of view of the cameras, it is possible for subjects to stray partially out of the video frames, leading to frames with partial faces or only one iris visible. Figure 12.2 shows corresponding frames from each of the three IOM cameras while a subject passes through the in-focus region of the IOM. Each frame captured by one of the IOM cameras is a 2048 by 2048 pixel gray scale image. A typical iris acquired by the system is approximately 120 pixels in diameter.

The general steps used in this work to combine face and iris biometrics from the IOM sensor are outlined in Fig. 12.3. As previously described, when a subject passes through the IOM portal, three videos are collected, with one video coming from each of the IOM cameras. In a preprocessing step, the corresponding frames of the three videos are stitched together to create one virtual video. Next, a series of detection phases are used to locate whole faces and eyes in each frame. Matching is then performed on each face sample and iris sample independently, and the results are fused using several different techniques.

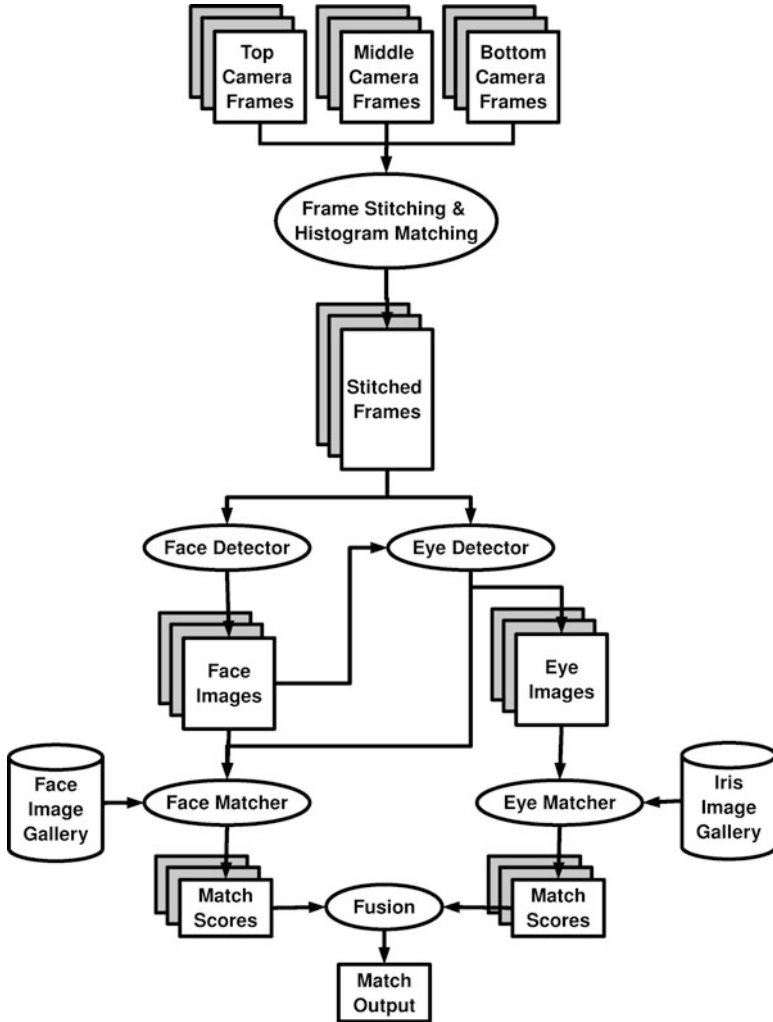


Fig. 12.3 A diagram of the pipeline used in the proposed multi-biometric system

12.5.1 Preprocessing

In order to increase the likelihood of a whole face being captured for each subject, the three videos from each IOM acquisition are “stitched” together to combine corresponding frames. As can be seen in Fig. 12.2, there is significant vertical overlap between the top and middle cameras, as well as between the middle and bottom cameras. Due to imperfect calibration of the individual cameras, some horizontal misalignment between the cameras is also present.




Fig. 12.4 An example of the progression during alignment between corresponding frames from the *top* to *middle* camera. The *top left* image is the frame from the *top* camera with the template marked as a *rectangle*. The *bottom left* image is the frame from the *middle* camera, with the matched region indicated. The *middle* image is the composite image, with the frame from the *top* camera cropped and padded. The overlapping region is indicated. The *right* image shows the final stitching results after histogram matching. A similar approach is used to stitch the frame from the *bottom* camera to the *middle* frame

A template-matching approach is taken to determine the desired translation to align frames from adjacent cameras. Specifically, the bottom portion of the top frame is cropped and used as a template. This template is then matched against the upper half of the middle frame, and the best match is selected as the desired alignment. This process is repeated for the bottom camera, where the template is created from the top portion of the bottom frame and matched against the lower half of the middle frame.

Finally, noticeable illumination differences were observed between corresponding frames from different cameras. To account for this discrepancy, histogram matching is used to match the top and bottom frame to the illumination observed in the middle frame. Figure 12.4 shows the intermediate and final results of the stitching procedure for an example frame.

12.5.2 Face Detection

Once the frame stitching is completed, the next step in the preprocessing phase is to detect a face in each frame. To accomplish this task, the OpenCV implementation of the Viola-Jones cascade face detector is used [3, 20]. The detector was trained on whole faces and thus may or may not detect faces which lie only partially within the field of view of the camera.

12.5.3 Eye Detection

The purpose of the eye detection phase is twofold. The primary goal is to detect any eyes present in each frame for iris matching. However, the locations of the eyes that are detected in the faces produced by the face detector are also used for an alignment phase during face matching. A template-matching approach is adopted for eye detection. The template used to search for eyes in each frame is based on the specular highlights generated by the reflection of the IOM LEDs.

The eye detection is completed in two phases. First, the template matching is performed on the upper left and upper right quadrants of each face detected by the face detector. This approach guarantees that each detected face will have two eye locations estimated as well.

Because it is possible for eyes to be detected in frames where whole faces were not present (or in frames where the face detector failed to detect the face), a second round of template matching is performed on any stitched frame where a face was not detected. In these frames, the location of the partial face can be crudely estimated by computing the sums of the rows and columns of the image and comparing these sums to appropriate thresholds. This partial face detection step is not required, but reduces the likelihood of false eye detections by limiting the search space to the region of the image that is likely to contain the eyes. An example of a face region being estimated in this manner is shown in Fig. 12.5. Once the partial face region has been estimated, the template matching is performed twice to identify the two best eye locations. Finally, the detected eyes are cropped from the corresponding location in the *original* frames to remove any possible artifacts caused by the histogram matching in the stitching phase. In cases where the detected eye is located in the overlapping region between two cameras, the eye is cropped from *both* camera frames.

12.5.4 Face Matching

In this work, Colorado State University's implementation of the eigenface algorithm is used for face matching [5, 19]. To achieve alignment with the training set, the probe face images are normalized using the eye centers detected by the eye detector. The Mahalanobis cosine metric is used to compute the distance between two feature vectors. Using this metric, match scores can range from -1.0 to 1.0 , with -1.0 being a perfect score. The output of the face matcher stage of the pipeline is a distance for every comparison between each probe face image and gallery face image.

12.5.5 Iris Matching

For the iris matcher, a modified version of Daugman's algorithm is used to compare each probe iris image to the gallery [6]. The normalized fractional Hamming




Fig. 12.5 Example of the image projection technique used to estimate the location of the face during eye detection. The graphs on the *right* and *bottom* of the image represent the summations of the pixel values in each *row* or *column*, respectively. The projected lines represent the face boundaries determined using appropriate thresholds

distance, referred to simply as the Hamming distance in the rest of this work, ranges from 0.0 to 1.0, with 0.0 being a perfect match. The Hamming distance is normalized to adjust low Hamming distances that occur for comparisons that used relatively few bits. The output of the iris matcher stage of the pipeline is a Hamming distance for every comparison between each probe eye image and gallery iris image.

12.5.6 Fusion

In this framework, the problem is multi-sample (i.e., several faces from each video), multimodal (i.e., both iris and face samples from each video), and multi-instance (i.e., both left and right irises from each video). Consequently, there are many methods which could be used to combine the face and iris biometrics from each video. Several fusion techniques are considered at both the score and rank levels.

The first method considers only one biometric modality in the fusion process and makes use only of the multi-sample and multi-instance dimensions of the problem by taking the minimum score for a given modality. For example, in the MinIris approach, the minimum score for all of the iris comparisons from a given video is reported as the best match. Similarly, the MinFace approach takes the minimum score for all of the face comparisons from a given video to determine the best match. Equations (12.1) and (12.2) express the MinIris and MinFace fusion rules, respectively, for a given probe video:

$$\text{MinIris} = \text{Min}\{I_{i,j} | i = 1 \dots n, j = 1 \dots G\} \quad (12.1)$$

$$\text{MinFace} = \text{Min}\{F_{i,j} | i = 1 \dots m, j = 1 \dots G\} \quad (12.2)$$

where n and m are the number of irises and faces detected in the video, respectively; G is the number of gallery subjects; $I_{i,j}$ is the Hamming distance between the i -th iris and the j -th gallery subject; and $F_{i,j}$ is the score for the comparison between the i -th face and the j -th gallery subject.

The next type of fusion method considered is rank-level fusion and can incorporate face, iris, or both modalities into the decision process. A Borda count is used to determine a best match across the desired biometric modalities. In a Borda count, the scores for all comparisons from a given sample are sorted such that the first rank corresponds to the best score for that sample. Each sample then casts votes for the top v ranked subjects, where the weight of each vote is inversely proportional to rank number. Each sample votes in this manner, and the gallery subject with the most votes is taken to be the best match. In these experiments, the BordaIris method considers only the iris scores to perform fusion, and the BordaFace method considers only face scores. The BordaBoth method allows both face and iris samples to vote, with v votes being cast by each iris and face sample.

Two vote weighting schemes are tested for the BordaIris, BordaFace, and BordaBoth fusion methods. In the linear approach, the vote weight is linearly proportional to the rank; specifically, the weight associated with the rank- n match is described by the equation

$$\text{VoteWeight}_n = v + 2 - n \quad (12.3)$$

and v represents the total number of votes cast by each biometric sample. In the exponential approach, the weight of the vote is exponentially related to the

rank. Specifically, the weight associated with the rank- n match is described by the equation

$$\text{VoteWeight}_n = 2^{v-n} \quad (12.4)$$

The third fusion method again uses score-level fusion, implementing a weighted summation of the iris and face scores. The summation rule can be expressed as Eq. (12.5) for a given probe video:

$$\text{SumScore}_k = \frac{\alpha * \sum_{i=1}^n (1 - \text{FNorm}_{i,k}) + \beta * \sum_{j=1}^m (1 - \text{INorm}_{j,k})}{\alpha * n + \beta * m} \quad (12.5)$$

where n and m are the number of irises and faces detected in the video, respectively; $\text{INorm}_{j,k}$ is the normalized Hamming distance between the j -th iris and the k -th gallery subject; and $\text{FNorm}_{i,k}$ is the normalized score for the comparison between the i -th face and the k -th gallery subject. Each face and iris score is normalized using min-max normalization, according to the expression

$$\text{Score}' = \frac{\text{Score} - \text{Min}}{\text{Max} - \text{Min}} \quad (12.6)$$

where Min and Max are the minimum and maximum possible values for each score metric, so that all normalized scores fall between 0.0 and 1.0, with 1.0 representing a perfect match. In Eq. (12.5), α and β are coefficients used to weight the face and iris biometrics, respectively. In the presented work, $\alpha = 1 - \beta$ for simplicity. In Eq. (12.5), SumScore_k represents the final match score for the given probe video with gallery subject k ; the best match score can be determined by finding the maximum SumScore_k for all k . SumIris is the special case where $\alpha = 0$ and $\beta = 1$, which corresponds to summing only the iris scores to determine the best match. Similarly, SumFace is the case where $\alpha = 1$ and $\beta = 0$ and equates to summing only the normalized face scores.

12.6 Experiments

The previously described multi-biometric system was tested on a probe dataset of 1,886 IOM video sets. Note that here a video “set” refers to the corresponding videos from each of the three IOM cameras, so the dataset is comprised of 5,658 videos in total. The 1,886 videos spanned 363 unique subjects, with an average of about five videos per subject. The most frequently occurring probe subject had 15 videos in the probe set, and the least frequently occurring had one probe video.

The iris gallery contained one left eye and one right eye for each of the 363 gallery subjects. The gallery images were acquired using the LG IrisAccess 4000 (LG4000) [8], a high-quality iris acquisition camera, and the gallery was manually




Fig. 12.6 Images of the same iris image using the LG4000 (*left*) and the IOM (*right*). The IOM image shown on the right represents a well-focused IOM iris image

screened for good quality and segmentation. For comparison, Fig. 12.6 shows an example of an image of the same iris acquired from both the LG4000 and the IOM.

The face gallery contained one full face image for each of the 363 subjects. The gallery images were acquired using the IOM. Each of the 363 subjects in the study had an additional IOM video set acquired in which the presence of a whole face was verified manually. The frames were stitched using the process previously described, and then, the best frame was manually selected, and the coordinates of the eye centers were manually annotated for alignment. The PCA training was performed on the face image gallery.

12.6.1 Detection Results

Across the entire dataset, 14,829 left irises and 14,711 right irises were detected and successfully segmented, and 9,833 faces were detected with valid eye locations for alignment. In this context, “successful segmentation” simply means that the iris segmentation routine returned pupil and limbic boundaries; it does *not* guarantee correctness. On average, 15.7 ($\sigma = 8.1$) irises, 5.2 ($\sigma = 3.7$) faces, and 20.9 ($\sigma = 20.9$) of either biometric samples were found in each video.

Table 12.1 provides a breakdown of the detection results by frame and video. The 1,886 videos were composed of a total of 28,381 frames. From Table 12.1, it can be seen that while a large number of frames (44.1%) contained no detected features, a much larger percentage of the probe *videos* (99.3%) had at least one biometric feature detected. Further, the majority (80.6%) of the probe videos contained samples of face and both iris features.

Table 12.1 Detailed detection results

Modalities detected	Frame count	Video count
Left iris (only)	1,447 (5.1%)	35 (1.9%)
Right iris (only)	2,104 (7.4%)	46 (2.4%)
Face (only)	900 (3.2%)	2 (0.1%)
Left and right irises (only)	2,495 (8.8%)	209 (11.1%)
Face and left iris (only)	1,411 (5.0%)	34 (1.8%)
Face and right iris (only)	724 (2.6%)	27 (1.4%)
Face, left, and right irises	6,798 (24.0%)	1,522 (80.6%)
None	12,502 (44.1%)	11 (0.6%)




Fig. 12.7 The match and non-match score distributions for the face features from the entire probe dataset

12.6.2 Matching Results

Figure 12.7 shows the match and non-match score distributions for all 9,833 detected faces. The mean match score was -0.281 with a standard deviation of 0.213 , while the mean non-match score was 0.000 with a standard deviation of 0.676 . If each face were treated independently, the rank-one recognition achieved for the 9,833 probes faces would be 51.6% ($5,073/9,833$) recognition.

The results from the left and right irises were aggregated, and Fig. 12.8 shows the match and non-match score distributions. The mean match score was 0.398 with




Fig. 12.8 The match and non-match score distributions for the *left* and *right* iris features from the entire probe dataset

a standard deviation of 0.053, while the mean non-match score was 0.449 with a standard deviation of 0.013. Figure 12.8 shows a significant number of match comparisons with fairly high scores. Upon examination of the data, it was found that most of these scores arise from incorrect segmentation. In some cases, these high match scores were caused by severe image defocus. Additionally, there are some false positives from the eye detector (non-eye regions) that contain features that resemble pupil and limbic boundaries according to the segmentation routine. If each iris image were treated independently, the rank-one recognition achieved for all of the probe irises would be 46.6% (13,556/29,112) recognition.

12.6.3 Fusion Results

The results of the iris and face matchers were combined using each of the methods previously described. The rank-one recognition rates achieved by each fusion approach are shown in Table 12.2. In the fusion methods based on Borda counts, the number of votes given to each sample was varied between 1 and 363 (though all samples were given the same number of votes for any given fusion experiment), and the best results for each approach are presented. Similarly, results from the optimal tested values of α and β are presented.

Table 12.2 Rank one recognition rates for fusion approaches

Approach	Fusion parameters	Rank one (raw)
MinIris		86.7% (1,635/1,886)
MinFace		62.6% (1,180/1,886)
BordaIris-linear	$\nu = 3$	86.4% (1,629/1,886)
BordaIris-exponential	$\nu = 20$	86.8% (1,637/1,886)
BordaFace-linear	$\nu = 3$	58.9% (1,110/1,886)
BordaFace-exponential	$\nu = 5$	59.3% (1,118/1,886)
BordaBoth-linear	$\nu = 10$	91.7% (1,729/1,886)
BordaBoth-exponential	$\nu = 10$	92.0% (1,735/1,886)
SumIris	$\alpha = 0.0, \beta = 1.0$	87.8% (1,656/1,886)
SumFace	$\alpha = 1.0, \beta = 0.0$	61.3% (1,156/1,886)
SumBoth	$\alpha = 0.3, \beta = 0.7$	93.2% (1,757/1,886)




Fig. 12.9 ROC curves for the various fusion methods using the optimal tested parameters for each. The BordaBoth method shown is the BordaBoth-exponential method

In summary, the best single-modality fusion approach was the SumIris approach, which achieved an 87.8% rank-one recognition rate. The SumBoth approach achieved the overall highest recognition rate (93.2%), and all multi-modal fusion approaches achieved higher recognition rates than the fusion methods based on a single modality.

Figure 12.9 shows the ROC curves for the best SumBoth and BordaBoth approaches, as well as the MinIris, MinFace, SumFace, and SumIris results

for comparison. From this graph, it is clear that the BordaBoth and SumBoth approaches outperform the single-modality fusion methods. Interestingly, while SumBoth achieved the highest rank-one recognition rate, Fig. 12.9 shows that the BordaBoth fusion technique performs better at false-positive rates less than 0.06.

In general, the videos that failed to match correctly typically had relatively few face and iris features detected. While the iris proved to be the more accurate of the two modalities in the multi-sample fusion scenarios, Fig. 12.8 indicates that many of the iris features detected are of poor quality, represent false detections from the eye detector, or failed to segment correctly. While the fusion techniques in these experiments were able to overcome these challenges when enough samples were present, videos in which a small number of faces and iris are detected are much less likely to be correctly matched.

12.7 Conclusions

This chapter presents an investigation into the fusion of face and iris biometrics from a single sensor, a surprisingly understudied problem in current literature. The previously described multi-biometrics framework utilizes multi-sample, multi-instance, and multimodal fusion techniques to improve recognition rates from a single sensor. The multi-biometric system is tested on a non-chimeric dataset of over 1,886 videos spanning 363 subjects. This represents one of the largest genuine multimodal experiments that has been conducted to date. Face and iris biometric samples extracted from videos produced from the Iris on the Move sensor were combined using several different fusion methods. In these experiments, the combination of face and iris biometrics via match score summation yielded a 5.4% increase in recognition rate over the best single-modality approach that was tested, while a modified Borda count approach performed best at lower false-positive rates (< 0.06).

The multi-biometrics system proposed exploits the face information collected by the IOM, a sensor that is intended for iris recognition purposes, with no modifications to the sensor and no increase in probe data acquisition time. The resulting system is less likely to experience failures to acquire, and the use of multiple modalities could allow the system to identify subjects with incomplete gallery data. This approach could be extended to operate on other stand-off iris sensors, which often detect the face as a preliminary step to iris image acquisition.

Acknowledgements Datasets used in this work were acquired under funding from the National Science Foundation under grant CNS01-30839, by the Central Intelligence Agency and by the Technical Support Working Group under US Army Contract W91CRB-08-C-0093. The authors were supported by a grant from the Intelligence Advanced Research Projects Activity.

References

1. Bowyer, K.W., Chang, K.I., Yan, P., Flynn, P.J., Hansley, E., Sarkar, S.: Multi-modal biometrics: an overview. In: Presented at the Second Workshop on Multi-Modal User Authentication (MMUA 2006), Toulouse (2006)
2. Bowyer, K.W., Hollingsworth, K., Flynn, P.J.: Image understanding for iris biometrics: a survey. *Comput. Vis. Image Underst.* **110**, 281–307 (2008)
3. Bradski, G., Kaehler, A.: *Learning OpenCV*. O'Reilly Media, Sebastopol (2008)
4. Chen, C.H., Te Chu, C.: Fusion of face and iris features for multimodal biometrics. In: Zhang, D., Jain, A. (eds.) *Advances in Biometrics*. Lecture Notes in Computer Science, vol. 3832, pp. 571–580. Springer, Berlin/Heidelberg (2005)
5. Colorado State University: Evaluation of Face Recognition Algorithms. <http://www.cs.colostate.edu/evalfacerec/algorithms5.html> (2010)
6. Daugman, J.: How iris recognition works. In: 2002 International Conference on Image Processing, Rochester, vol. 1, pp. 33–36 (2002)
7. Ko, T.: Multimodal biometric identification for large user population using fingerprint, face and iris recognition. In: 34th Applied Imagery and Pattern Recognition Workshop, Washington, DC, pp. 218–223 (2005)
8. LG Iris: LG Iris Products and Solutions. <http://www.lgiris.com/ps/products/irisaccess4000.htm> (2010)
9. Li, S.Z., Chu, R., Ao, M., Zhang, L., He, R.: Highly accurate and fast face recognition using near infrared images. In: International Conference on Biometrics (ICB 2006), Hong Kong, pp. 151–158 (2006)
10. Matey, J., Naroditsky, O., Hanna, K., Kolczynski, R., LoIacono, D., Mangru, S., Tinker, M., Zappia, T., Zhao, W.: Iris on the move: acquisition of images for iris recognition in less constrained environments. *Proc. IEEE* **94**, 1936–1947 (2006)
11. Morizet, N., Gilles, J.: A new adaptive combination approach to score level fusion for face and iris biometrics combining wavelets and statistical moments. In: Bebis, G., Boyle, R., Parvin, B., Koracin, D., Remagnino, P., Porikli, F., Peters, J., Klosowski, J., Arns, L., Chun, Y., Rhyne, T.M., Monroe, L. (eds.) *Advances in Visual Computing*. Lecture Notes in Computer Science, vol. 5359, pp. 661–671. Springer, Berlin/Heidelberg (2008)
12. National Institute of Standards and Technology (NIST): Portal Challenge Problem – Multiple Biometric Grand Challenge, Preliminary Results of Version 2. http://biometrics.nist.gov/cs_links/face/mbgc/2009/PORTAL_V2_FINAL.pdf (2009)
13. Phillips, P.J., Flynn, P.J., Beveridge, J.R., Scruggs, W.T., O'Toole, A.J., David, B., Bowyer, K.W., Draper, B.A., Givens, G.H., Lui, Y.M., Sahibzada, H., Scallan Iii, J.A., Weimer, S.: Overview of the multiple biometrics grand challenge. In: Proceedings of the Third International Conference on Advances in Biometrics, ICB '09, pp. 705–714. Springer, Berlin/Heidelberg (2009)
14. Rattani, A., Tistarelli, M.: Robust multi-modal and multi-unit feature level fusion of face and iris biometrics. In: Tistarelli, M., Nixon, M. (eds.) *Advances in Biometrics*. Lecture Notes in Computer Science, vol. 5558, pp. 960–969. Springer, Berlin/Heidelberg (2009)
15. Ross, A.: An introduction to multibiometrics. In: 15th European Signal Processing Conference (EUSIPCO), Poznan, pp. 20–24 (2007)
16. Ross, A., Jain, A.K.: Multimodal biometrics: an overview. In: 12th European Signal Processing Conference (EUSIPCO), Vienna, pp. 1221–1224 (2004)
17. Ross, A.A., Nandakumar, K., Jain, A.K.: *Handbook of Multibiometrics*. Springer Science and Business Media. Springer, New York (2006)
18. Son, B., Lee, Y.: Biometric authentication system using reduced joint feature vector of iris and face. In: Kanade, T., Jain, A., Ratha, N. (eds.) 6th International Conference on Audio- and Video-Based Biometric Person Authentication (AVBPA'03). Lecture Notes in Computer Science, vol. 3546, pp. 513–522. Springer, Berlin/Heidelberg (2005)

19. Turk, M., Pentland, A.: Face recognition using eigenfaces. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '91), Lahaina, pp. 586–591 (1991)
20. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2001), Kauai, vol. 1, pp. 511–518 (2001)
21. Volner, R., Bores, P.: Multi-biometric techniques, standards activities and experimenting. In: 2006 International Baltic Electronics Conference, Tallinn, pp. 1–4 (2006)
22. Wang, Y., Tan, T., Jain, A.K.: Combining face and iris biometrics for identity verification. In: 4th International Conference on Audio- and Video-Based Biometric Person Authentication (AVBPA'03), pp. 805–813. Springer, Berlin/Heidelberg (2003)
23. Yang, J., Liao, S., Li, S.: Automatic partial face alignment in nir video sequences. In: Tistarelli, M., Nixon, M. (eds.) *Advances in Biometrics. Lecture Notes in Computer Science*, vol. 5558, pp. 249–258. Springer, Berlin/Heidelberg (2009)
24. Zhang, Z., Wang, R., Pan, K., Li, S., Zhang, P.: Fusion of near infrared face and iris biometrics. In: Lee, S.W., Li, S. (eds.) *Advances in Biometrics. Lecture Notes in Computer Science*, vol. 4642, pp. 172–180. Springer, Berlin/Heidelberg (2007)

Chapter 13

Methods for Iris Segmentation

Raghavender Jillela and Arun A. Ross

Abstract Under ideal image acquisition conditions, the iris biometric has been observed to provide high-recognition performance compared to other biometric traits. Such a performance is possible by accurately segmenting the iris region from the given ocular image. This chapter discusses the various challenges associated with the segmentation process, along with some of the prominent iris segmentation techniques proposed in the literature. Furthermore, methods to refine and evaluate the output of the iris segmentation routine are presented. The goal of this chapter is to provide a brief overview of the progress made in iris segmentation.

13.1 Introduction

An iris acquisition system typically captures an image of the eye that, besides the iris, includes the pupil, eyelids, eyelashes, and sclera. Thus, the iris is perceived to be located in the vicinity of these structures in a 2D image of the eye. The process of locating and isolating the iris from such an image is known as iris *localization* or *segmentation*. The primary task of segmentation is to determine pixels in the image that correspond to the iris region. In Fig. 13.1, the pupillary boundary refers to the boundary separating the pupil (the black region in the center of the eye) from the iris (the textured region surrounding the pupil) while the limbus boundary refers to the boundary separating the iris from the sclera (the white of the eye). Typically, segmentation is accomplished by detecting the pupillary boundary and the limbus boundary as well as the eyelids and eyelashes that can interrupt the contour of the limbus boundary. Iris segmentation is a critical component of any iris biometric system; inaccuracies in localizing the iris can severely impact the matching accuracy of the system, thereby undermining its utility.

R. Jillela (✉) • A.A. Ross
West Virginia University, Morgantown, WV, USA
e-mail: Raghavener.Jillela@mail.wvu.edu; arun.ross@mail.wvu.edu




Fig. 13.1 The anterior surface anatomy of a human iris




Fig. 13.2 (a) Input iris image. (b) Close-up view of the iris. (c) Contour indicating the pupillary boundary. Notice that the pupillary boundary is elliptical and not circular


To provide an accurate output, an iris segmentation algorithm is required to consider the following characteristics of the iris boundaries¹:

1. *Circularity of the iris boundaries*: Some of the earliest iris segmentation algorithms [4, 27] operated under the assumption that the iris boundaries are circular. However, the iris boundaries are not always circular. Figure 13.2 shows a sample iris image in which the pupillary boundary is elliptical in shape. A wide number of recently proposed algorithms [7, 23] accommodate noncircular (or even non-conic) boundaries.
2. *Nasal inclination of the pupil*: The medical literature suggests that in most cases, the pupil is nasally inclined. Therefore, the center (or centroid) of the pupillary boundary may not always coincide with that of the limbus boundary. Figure 13.3 shows an iris image with non-concentric iris boundaries.
3. *Difference in the variation of intensities at the boundaries*: Typically, the variation in image intensities across the pupillary boundary is much stronger than the variation across the limbus boundary.² This property lends to the fact

¹The term “iris boundaries” is used in this chapter to collectively refer to both the pupillary and limbus boundaries.

²This is true for images obtained in the near-infrared spectrum.

Fig. 13.3 An image showing the *right* eye of an individual. Notice that the pupil is slightly inclined toward the subject's nose (on the *right* side of the image). The distance from the pupillary boundary to the limbus boundary on the *left* side of the image, d_1 , is greater than the distance on the *right* side, d_2



that, in most cases, the pupillary boundary can be easily detected using intensity thresholding operation. On the other hand, determining the limbus boundary can be a comparatively difficult task.

13.2 Challenges in Iris Segmentation

The performance of an iris segmentation algorithm is affected by the image acquisition process [30]. A majority of iris recognition systems require a significant amount of user cooperation during image acquisition to provide good recognition performance. For an iris image acquired under near-ideal conditions (e.g., good illumination, cooperative subject, un-occluded iris), iris segmentation can be accomplished using simple image-processing techniques. This is based on the observation that both the iris boundaries show significant variation in pixel intensities across their contours (e.g., iris versus the pupil and iris versus the sclera). However, when an iris image is acquired under nonideal conditions, segmentation becomes a challenging task. Some of the factors that render iris segmentation challenging are listed below:

1. *Occlusions caused by the anatomical features of the eye:* One of the most common and significant challenge faced during iris segmentation is the occlusion caused by the eyelids and/or the eyelashes.
 - *Eyelids:* The eyelids are thin folds of skin that cover and protect the eye from foreign bodies and extreme lighting. The movement of eyelids can be both voluntary (e.g., closing eyelids when tired) and involuntary (e.g., blink caused by a reflex). To obtain an un-occluded image of the iris, the user is required to hold the eyelids wide open for a brief period of time during image acquisition. However, under normal conditions, a minor portion of the human eye is typically occluded on the top and the bottom by the upper and the lower eyelid, respectively. Figure 13.4 shows an iris image exhibiting eyelid occlusion. In such cases, the contour of the limbus boundary is no longer circular or elliptical.

Fig. 13.4 An iris image showing occlusions caused by the eyelids





Fig. 13.5 An iris image showing occlusions due to eyelashes



- *Eyelashes:* Eyelashes are the hair at the end of the eyelids. Like eyelids, eyelashes also provide protection to the eye from external debris. Although the occlusions caused by eyelashes are minimal, accurate detection of the limbus boundary becomes very difficult in the presence of eyelashes. This is due to the fact that eyelashes can cause uneven interruptions at the limbus boundary. Empirical observations reveal that eyelash occlusion is typically more pronounced in Asian subjects, due to the presence of the epicanthic fold. Figure 13.5 shows an iris image with eyelash occlusions.

2. *Illumination:* The level and the type of illumination (lighting) used, along with the eye region on which it is focused, play an important role in the quality of an iris image.

- *Poor illumination:* Segmenting an iris image acquired under poor illumination is extremely difficult because the image may offer minimal or no information about the boundaries of the iris. Furthermore, the texture of the iris in such images may not be adequately highlighted, resulting in poor recognition performance. Figure 13.6 shows an iris image acquired under poor illumination conditions.

Fig. 13.6 An iris image acquired under poor illumination conditions. Although the iris boundaries can be distinguished by humans, automatic segmentation of such images is extremely difficult





Fig. 13.7 An iris image containing specular reflections on the pupillary boundary



- *Specular reflections*: Specular reflections are small regions in an iris image characterized by pixels of high-intensity values that are typically caused by improper focusing of the light source. If specular reflections are present on (or even close to) the iris boundaries, iris segmentation becomes difficult. Figure 13.7 shows an iris image with specular reflections on the pupillary boundary.
3. *User cooperation*: Most iris image acquisition systems require a considerable amount of user cooperation to record a good quality image. In cases where the user cooperation is low, the acquired images can be of poor quality, thereby affecting the segmentation performance.
- *Off-angled iris*: Off-angled iris images are caused when the sensor is not orthogonal to the plane of the iris. In such cases, the surface area of the iris is reduced, and the detection of the boundaries becomes difficult. Off-angled iris images are typically caused in situations where the image acquisition is nonideal, i.e., when the subject is in motion or not aware of the image acquisition process. Figure 13.8 shows an off-angled iris image.
 - *Motion blur*: Motion blur in iris images can occur mainly due to three reasons: (a) when the image is acquired from a moving subject, (b) movement of the camera, and (c) movement of the subject's eye while adjusting to the device and the environment. In images containing motion blur, the intensity variations across the iris boundaries may be reduced, thereby impacting segmentation. Figure 13.9 shows an image containing motion blur.

Fig. 13.8 An off-angled iris image




Fig. 13.9 An iris image containing motion blur



- *Eye glasses or contact lenses:* If a user wears an eye glass or contact lens, the acquired iris images may suffer from additional reflection artifacts due to these entities. In a cooperative iris recognition system, this problem can be minimized by requesting the user to avoid wearing eye glasses during image acquisition. However, if a subject wears contact lenses (cosmetic/non-cosmetic), it may not be convenient for the user to remove them, even in a cooperative iris recognition system. Research has shown that contact lenses can impact the performance of iris segmentation and recognition [1]. Figure 13.10 shows the iris image of a user wearing a contact lens (non-cosmetic, hard lens).
4. *External factors:* Some external factors play a major role in the quality of the input iris images, thereby impacting the segmentation performance.
- *Standoff distance:* Standoff distance refers to the distance of the camera from the subject. If the standoff distance is large, the resolution (number of pixels occupied by the iris region in an image) can be low. In such cases, the iris boundaries may not be clearly distinguishable. Figure 13.11 shows an iris image acquired at a large standoff distance.

Fig. 13.10 Iris image of a subject wearing a non-cosmetic, hard contact lens. An automated segmentation scheme may fail to distinguish the iris boundaries from the lens boundary (Image source: Baker et al. [1] ©Elsevier)




Fig. 13.11 Close-up of an iris image acquired at a large standoff distance



- *Image sensors*: The following factors related to image sensors play a significant role in acquiring a good quality iris image:
 - (a) Resolution of the sensor: A high-resolution sensor that can capture iris images with a minimum diameter of 200 pixels are preferred over other low-resolution sensors.
 - (b) Positioning of the sensor: The positioning of the sensor plays an important role in acquiring an iris image of good quality. For example, if the iris sensor is placed above or below the eye level of a subject, the acquired image may contain an off-angled iris.
 - (c) Sensor noise: Although not seen as a major factor, sensor noise can produce artifacts in an image, thereby affecting iris segmentation.
- *Nature of the interacting population*: To acquire a good quality iris image, it is required for the target population to be cooperative and habituated with the iris image acquisition system.
- *Outliers*: In very rare cases, abnormalities in the shape of the iris can cause problems with segmentation. An example of a subject suffering from congenital iris abnormality is shown in Fig. 13.12. In such cases, it can be noticed that the iris boundaries are neither circular nor elliptical.

Fig. 13.12 Iris image of a subject suffering from congenital iris abnormality. Notice that the pupil boundary is neither circular nor elliptical (Image source: <http://sciencerooll.files.wordpress.com/2007/06/congenitalirisabnorm.jpg>)



13.3 Classical Iris Segmentation Algorithms

13.3.1 Integro-Differential Operator

The technique proposed by John Daugman [4] in the early 1990s is considered to be the pioneering work in the field of automated iris recognition. A vast majority of commercial iris recognition systems worldwide employ Daugman's approach. In this approach, iris segmentation is carried out using an integro-differential operator. The segmentation process is performed by approximating the iris boundaries as perfect circles.

Given an iris image $I(x,y)$, it is first convolved with an image smoothing function (e.g., a Gaussian filter). This process of smoothing the image helps in (a) attenuating the effect of noise (e.g., sensor noise) in the image and (b) eliminating undesired weak edges (e.g., boundaries within the iris), while retaining the desired strong edges (e.g., iris boundaries, eyelid boundaries). An integro-differential operator is then used to search for the maximum value of a normalized integral along circular contours of varying radii and center coordinates. The search process over the image domain (x,y) using an integro-differential operator can be mathematically expressed as


$$\mathbf{max}(r, x_0, y_0) \left| G_\sigma(r) * \frac{\partial}{\partial r} \oint_{r, x_0, y_0} \frac{I(x,y)}{2\pi r} ds \right|, \quad (13.1)$$

where

$$G_\sigma(r) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(r-r_0)^2}{2\sigma^2}\right) \quad (13.2)$$

represents the radial Gaussian with a center r_0 and standard deviation (scale) σ , which is used for image smoothing. The symbol $*$ denotes the convolution operation, and r represents the radius of the circular arc ds , centered at the location (x_0, y_0) . The division by a factor of $2\pi r$ normalizes the circular integral with respect to its perimeter. In other words, the integro-differential operator behaves as a circular

Fig. 13.13 Iris segmentation using the integro-differential operator



edge detector that searches iteratively for the maximum response of a contour path defined by the parameters (x_0, y_0, r) . Depending on the values of the radii considered, the optimal parameters of the integro-differential operator are treated as either the pupillary or limbus boundaries. Figure 13.13 illustrates the search process using an integro-differential operator.


The value of σ , which controls the amount of blurring of the iris image, can be varied when searching for the pupillary and the limbus boundaries. As the pixel intensity variation across the pupillary boundary is more pronounced, the σ value can be set for a coarse scale of convolution. On the other hand, when the search process is carried out for the limbus boundary, the σ value is set for a finer convolution scale. This is due to the nominal variation of the pixel intensities across the limbus boundary.

In an iris image acquired under near-ideal conditions from a cooperative subject, both the iris boundaries can be easily detected using the integro-differential operator. However, in an image acquired under non-ideal conditions, the limbus boundary may not be completely circular due to the occlusions caused by the eyelids. Therefore, when searching for the limbus boundary, the angular arc of integration, ds , is often restricted to the left and right quadrants (i.e., near the vertical edges of the iris). When searching for the pupillary boundary, this arc can be extended over a wider range, as the eyelid occlusions are relatively small. Once both the iris boundaries are detected, the boundaries of the eyelids can be detected by changing the integration path of the operator from circular to arcuate. Figure 13.14 shows the output of the integro-differential operator when used to detect both the iris boundaries and the eyebrows.

Despite successfully determining the iris boundaries using an integro-differential operator, the localized iris may be further occluded by other noisy regions such as eyelashes, shadows, or specular reflections. Therefore, a noise mask³ that records

³The process of generating a noise mask and the subsequent schemes for iris normalization and matching are very similar in a majority of iris recognition algorithms. However, as this chapter

Fig. 13.14 Output obtained by applying an integro-differential operator to detect both the iris and eyebrow boundaries



the locations of these undesired occlusions is correspondingly generated. This mask is later used during the matching stage to mitigate the effect of noisy pixels.

13.3.2 Hough Transform

Another widely used classical iris segmentation algorithm was proposed by Wildes et al. [26, 27]. To detect each iris boundary, the algorithm relies on the Hough transform: a histogram-based model-fitting approach. First, an edge map of the input image is generated using a gradient-based edge detector. A voting procedure is then applied on the highlighted edge map, to determine the parameter values for a contour that best fits a circle.

Given an iris image $I(x, y)$, the edge map of an input image can be highlighted by thresholding the magnitude of the image intensity gradient. This operation can be mathematically expressed as

$$|\nabla G(x, y) * I(x, y)| \geq \text{th}, \quad (13.3)$$

where $\nabla \equiv (\partial/\partial x, \partial/\partial y)$ and th denotes an empirically chosen intensity threshold. $G(x, y)$ represents a two-dimensional Gaussian with center, (x_0, y_0) , and standard deviation, σ , used for smoothing the image. The purpose of image smoothing is to (a) select the spatial scale of edges under consideration and (b) reduce the effect of noise on the thresholding process. The mathematical expression for the Gaussian is

$$G(x, y) = \frac{1}{2\pi\sigma^2} \exp^{-\frac{(x-x_0)^2 + (y-y_0)^2}{2\sigma^2}}. \quad (13.4)$$

focuses only on iris segmentation, these details are not discussed. The reader is directed to the original publication by Daugman [4] for further information.

The image thresholding operation yields an approximate edge map consisting of the iris boundary (along with other prominent edges). In most cases, such an output would consist of noncontinuous, noncircular contours. Therefore, the edge map is thinned using a morphological operation, and a voting procedure is used to determine the parameters of the iris boundaries. Hough transform [9], a standard machine vision technique for fitting simple contour models to images, is typically used during the voting process.

When searching for the limbus boundary contour, the image intensity derivatives corresponding to vertical edges are weighted more during the voting process. This directional selectivity grants preference to the left and right portions of the limbus boundary over its upper and lower portions. Therefore, even if the upper and lower portions of the limbus boundary are occluded by eyelids, the left and right portions remain clearly visible and oriented (assuming the head is in an upright position).

Consider a set of edge points (x_j, y_j) , $j = \{1, 2, \dots, n\}$ obtained by the image thresholding operation. The goal is to determine if a subset of these points are associated with the contour of a circle. A circle can be parameterized as (x_c, y_c, r) where (x_c, y_c) denotes its center and r denotes its radius. The Hough transform detects circular contours in the edge image by defining an accumulator array, H , whose entries $H(x_c, y_c, r)$ are computed as follows:

$$H(x_c, y_c, r) = \sum_{j=1}^n h(x_j, y_j, x_c, y_c, r) \quad (13.5)$$

where

$$h(x_j, y_j, x_c, y_c, r) = \begin{cases} 1, & \text{if } g(x_j, y_j, x_c, y_c, r) = 0 \\ 0, & \text{otherwise.} \end{cases} \quad (13.6)$$

and

$$g(x_j, y_j, x_c, y_c, r) = (x_j - x_c)^2 + (y_j - y_c)^2 - r^2. \quad (13.7)$$

For each edge point (x_j, y_j) , $g(x_j, y_j, x_c, y_c, r)$ is set to 0 if the parameter triplet (x_c, y_c, r) represents a circle through that point. The parameter triplet that maximizes H is considered to be a reasonable choice to represent the contour of interest. The maximizing parameter triplet is determined by first building $H(x_c, y_c, r)$ as an array (indexed by discretized values for x_c , y_c , and r) and then scanning for the triplet that corresponds to the largest value in the array.

The same process (consisting of thresholding and voting) is used to determine the pupillary boundary but with the following minor modifications:

1. The image is filtered with a gradient-based edge detector that is *not* directionally tuned. This is due to the fact that the pupillary boundary is less prone to occlusion from the eyelids.
2. The permissible parameter values (x_c, y_c, r) are constrained to lie within the circle that describes the limbus boundary.

Once both the iris boundaries are detected, it is necessary to determine the locations of the upper and lower eyelids that may occlude the iris. To perform this operation, a gradient-based edge detector that is tuned to favor horizontal edges is used. This is based on the fact that the contour of the upper and lower eyelids within the limbus boundary would be nearly horizontal (under the assumption that the subject's head is in an upright position). The upper and lower eyelids are modeled as two separate parabolic arcs of the form $x(t) = a_x t^2 + b_x t + c_x$, and $y(t) = a_y t^2 + b_y t + c_y$ with $0 \leq t \leq 1$. The parameters a_x, b_x, c_x, a_y, b_y , and c_y are once again determined using the same histogram-based model-fitting approach.

13.4 Other Prominent Approaches

Both the aforementioned techniques approximate an iris boundary with a circle or an ellipse. However, such an approximation may not be always suitable. Iris recognition performance can be improved by determining the precise boundaries of the iris, rather than their approximations. Determining the precise boundaries of the iris reduces the noise from occlusions, especially those caused by eyelashes. Figure 13.15 shows the difference between an approximated and a precisely determined iris boundary.

Methods that use curve evolution processes (e.g., geodesic active contours, level sets) could be extremely beneficial in determining the precise boundaries of an iris. In this section, some such techniques are discussed.

13.4.1 Geodesic Active Contours

This approach, proposed by Shah and Ross, is based on the relation between active contours and the computation of geodesics (minimal length curves). The strategy is to evolve an arbitrarily initialized curve from within the iris under the influence of




Fig. 13.15 (a) An eye image in which the limbus boundary is occluded by the eyelashes. (b) Approximate limbus boundary obtained using Daugman's integro-differential operator. (c) Precise limbus boundary obtained by a curve evolution technique. Notice that the precise segmentation helps in avoiding the noise caused by eyelash occlusion

geometric properties of the iris boundary. GACs combine the energy minimization approach of the classical “snakes” and the geometric active contours based on curve evolution.

Let $\gamma(t)$ be the curve, that has to gravitate toward the outer boundary of the iris, at a particular time t . The time t corresponds to the iteration number. Let ψ be a function measuring the signed distance from the curve $\gamma(t)$. That is, $\psi(x, y) =$ distance of point (x, y) to the curve $\gamma(t)$.

$$\psi(x, y) = \begin{cases} 0 & \text{if } (x, y) \text{ is on the curve;} \\ < 0 & \text{if } (x, y) \text{ is inside the curve;} \\ > 0 & \text{if } (x, y) \text{ is outside the curve.} \end{cases} \quad (13.8)$$

Here, ψ is of the same dimension as that of the eye image $I(x, y)$. The curve $\gamma(t)$ is called the level set of the function ψ . Level sets are the set of all points in ψ where ψ is some constant. Thus, $\psi = 0$ is the zeroth-level set, $\psi = 1$ is the first level set, and so on. ψ is the implicit representation of the curve $\gamma(t)$ and is called the embedding function since it embeds the evolution of $\gamma(t)$. The embedding function evolves under the influence of image gradients and the region’s characteristics so that the curve $\gamma(t)$ approaches the desired boundary of the iris. The initial curve $\gamma(t)$ is assumed to be a circle of radius r just beyond the pupillary boundary. Let the curve $\gamma(t)$ be the zeroth-level set of the embedding function. This implies that

$$\frac{d\psi}{dt} = 0.$$

By the chain rule,

$$\frac{d\psi}{dt} = \frac{\partial\psi}{\partial x} \frac{dx}{dt} + \frac{\partial\psi}{\partial y} \frac{dy}{dt} + \frac{\partial\psi}{\partial t},$$

i.e.

$$\frac{\partial\psi}{\partial t} = -\nabla\psi \cdot \gamma'(t).$$

Splitting the $\gamma'(t)$ in the normal ($N(t)$) and tangential ($T(t)$) directions,

$$\frac{\partial\psi}{\partial t} = -\nabla\psi \cdot (v_N N(t) + v_T T(t)).$$

Now, since $\nabla\psi$ is perpendicular to the tangent to $\gamma(t)$,

$$\frac{\partial\psi}{\partial t} = -\nabla\psi \cdot (v_N N(t)). \quad (13.9)$$

The normal component is given by

$$N = \frac{\nabla\psi}{\|\nabla\psi\|}.$$

Substituting this in Eq. (13.9),

$$\frac{\partial \psi}{\partial t} = -v_N \|\nabla \psi\|.$$

Let v_N be a function of the curvature of the curve κ , stopping function K (to stop the evolution of the curve) and the inflation force c (to evolve the curve in the outward direction) such that

$$\frac{\partial \psi}{\partial t} = -(\text{div}(K \frac{\nabla \psi}{\|\nabla \psi\|}) + cK) \|\nabla \psi\|.$$

Thus, the evolution equation for ψ_t^4 such that $\gamma(t)$ remains the zeroth-level set is given by

$$\psi_t = -K(c + \varepsilon \kappa) \|\nabla \psi\| + \nabla \psi \cdot \nabla K, \tag{13.10}$$

where K , the stopping term for the evolution, is an image-dependant force and is used to decelerate the evolution near the boundary; c is the velocity of the evolution; ε indicates the degree of smoothness of the level sets; and κ is the curvature of the level sets computed as

$$\kappa = -\frac{\psi_{xx} \psi_y^2 - 2\psi_x \psi_y \psi_{xy} + \psi_{yy} \psi_x^2}{(\psi_x^2 + \psi_y^2)^{\frac{3}{2}}}.$$

Here, ψ_x is the gradient of the image in the x -direction, ψ_y is the gradient in the y -direction. ψ_{xx} is the second order gradient in the x -direction, ψ_{yy} is the second-order gradient in the y -direction, and ψ_{xy} is the second-order gradient, first in the x -direction and then in the y -direction. Equation (13.10) is the level set representation of the geodesic active contour model. This means that the level set C of ψ is evolving according to

$$C_t = K(c + \varepsilon \kappa) \mathbf{N} - (\nabla K \cdot \mathbf{N}) \mathbf{N} \tag{13.11}$$

where N is the normal to the curve. The term $\kappa \mathbf{N}$ provides the smoothing constraints on the level sets by reducing their total curvature. The term $c \mathbf{N}$ acts like a balloon force, and it pushes the curve outward toward the object boundary. The goal of the stopping function is to slow down the evolution when it reaches the boundary. However, the evolution of the curve will terminate only when $K = 0$, i.e., near an ideal edge. In most images, the gradient values will be different along the edge, thus requiring the use of different K values. In order to circumvent this issue, the third geodesic term ($(\nabla K \cdot \mathbf{N})$) is necessary so that the curve is attracted toward the boundary (∇K points toward the middle of the boundary). This term makes it possible to terminate the evolution process even if (a) the stopping function has different values along the edges and (b) gaps are present in the stopping function.

⁴The subscript t denotes the iteration number.




Fig. 13.16 Stopping function for the geodesic active contours. (a) Original iris image, (b) stopping function K , and (c) modified stopping function K'

The stopping term used for the evolution of level sets is given by

$$K(x,y) = \frac{1}{1 + \left(\frac{\|\nabla(G(x,y) * I(x,y))\|}{k} \right)^\alpha} \tag{13.12}$$

where $I(x,y)$ is the image to be segmented, $G(x,y)$ is a Gaussian filter, and k and α are constants. As can be seen, $K(x,y)$ is not a function of t .

Consider an iris image to be segmented as shown in Fig. 13.16a. The stopping function K obtained from this image is shown in Fig. 13.16b (for $k = 2.8$ and $\alpha=8$). Assuming that the inner iris boundary (i.e., the pupillary boundary) has already been detected, the stopping function K is modified by deleting the circular edges corresponding to the pupillary boundary, resulting in a new stopping function K' . This ensures that the evolving level set is not terminated by the edges of the pupillary boundary (Fig. 13.16c).

A contour is first initialized near the pupil (Fig. 13.17a). The embedding function ψ is initialized as a signed distance function to $\gamma(t = 0)$ which looks like a cone (Fig. 13.17b). Discretizing Eq. (13.10) leads to the following equation:

$$\frac{\psi'_{i,j}{}^{t+1} - \psi'_{i,j}{}^t}{\Delta t} = -cK'_{i,j} \|\nabla \psi^t\| - K'_{i,j} (\varepsilon \kappa'_{i,j} \|\nabla \psi^t\|) + \nabla \psi'_{i,j} \cdot \nabla K'_{i,j}{}^t, \tag{13.13}$$

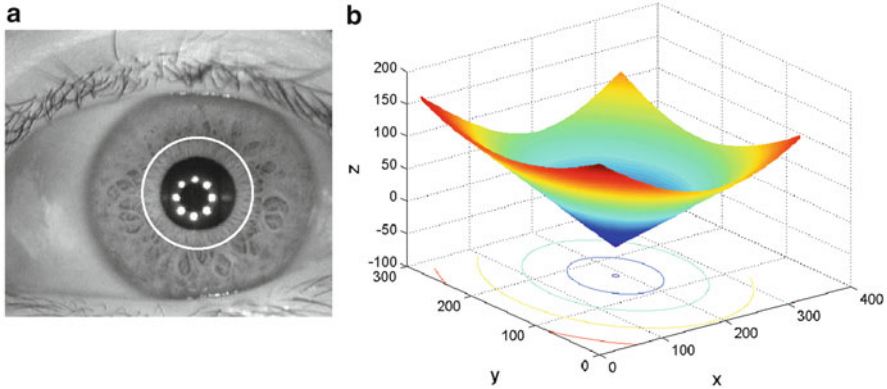


Fig. 13.17 Contour initialization for iris segmentation using GAC. (a) Zeroth-level set (initial contour), (b) mesh plot denoting the signed distance function ψ

where Δt is the time step (e.g., Δt can be set to 0.05). The first term ($cK'_{i,j}\|\nabla\psi^t\|$) on the right-hand side of the above equation is the velocity term (advection term) and, in the case of iris segmentation, acts as an inflation force. This term can lead to singularities and hence is discretized using upwind finite differences. The upwind scheme for approximating $\|\nabla\psi\|$ is given by

$$\begin{aligned} \|\nabla\psi\| &= \sqrt{A}, \\ A &= \min(D_x^- \psi_{i,j}, 0)^2 + \max(D_x^+ \psi_{i,j}, 0)^2 + \\ &\quad \min(D_y^- \psi_{i,j}, 0)^2 + \min(D_y^+ \psi_{i,j}, 0)^2. \end{aligned}$$

where $D_x^- \psi$ is the first-order backward difference of ψ in the x-direction, $D_x^+ \psi$ is the first-order forward difference of ψ in the x-direction, $D_y^- \psi$ is the first-order backward difference of ψ in the y-direction, and $D_y^+ \psi$ is the first-order forward difference of ψ in the y-direction. The second term ($K'_{i,j}(\varepsilon\kappa'_{i,j}\|\nabla\psi^t\|)$) is a curvature-based smoothing term and can be discretized using central differences. In our implementation, $c = 0.65$ and $\varepsilon = 1$ for all iris images. The third geodesic term ($\nabla\psi^t_{i,j} \cdot \nabla K'_{i,j}$) is also discretized using the central differences.

After evolving the embedding function ψ according to Eq. (13.13), the curve begins to grow until it satisfies the stopping criterion defined by the stopping function K' . But at times, the contour continues to evolve in a local region of the image where the stopping criterion is not strong. This leads to over-evolution of the contour. This can be avoided by minimizing the thin plate spline energy of the contour. By computing the difference in energy between two successive contours, the evolution scheme can be regulated. If the difference between the contours is less than a threshold (indicating that the contour evolution has stopped at most places), then the contour evolution process is terminated. The evolution of the curve and the corresponding embedding functions are illustrated in Fig. 13.18.




Fig. 13.18 Evolution of the geodesic active contour during iris segmentation. (a) Iris image with initial contour, (b) embedding function ψ (X and Y axes correspond to the spatial extent of the eye image and the Z axis represents different level sets), (c)–(f) contours after 600 and 1,400 iterations and their corresponding embedding functions, and (g), (h) final contour after 1,800 iterations (contours shown in *white*)

Since the radial fibers may be thick in certain portions of the iris or the crypts present in the ciliary region may be unusually dark, this can lead to prominent edges in the stopping function. If the segmentation technique is based on parametric curves, then the evolution of the curve might terminate at these local minima. However, geodesic active contours are able to split at such local minima and merge again. Thus, they are able to effectively deal with the problems of local minima, thereby ensuring that the final contour corresponds to the true limbus boundary (Fig. 13.19).

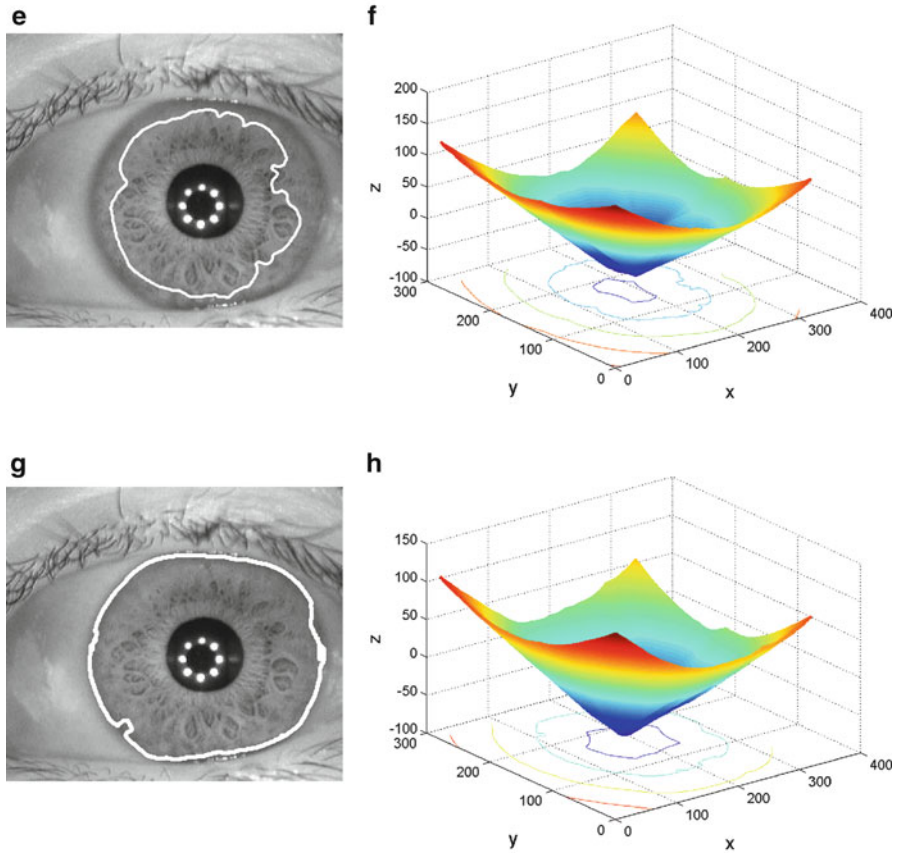


Fig. 13.18 (continued)





Fig. 13.19 The final contour obtained when segmenting the iris using the GAC scheme. (a) Example of a geodesic contour splitting at various local minima, (b) final contour (contours shown in *white*)

Fig. 13.20 Rough contours of the pupillary and limbus boundaries obtained using the elliptical model. Notice that the rough contours do not precisely match the true iris boundaries (Image source: Roy et al. [24] ©Elsevier)



13.4.2 Variational Level Sets

Another approach that can be used to precisely determine the iris boundaries is based on variational level sets [23, 24]. This approach uses partial differential equations (PDE) to numerically solve the evolution of the curves that define the iris boundaries. The iris boundaries are first approximated using elliptical models, which are then refined using geometric active contours with variational formulation.

Given an iris image, first an elliptical model with parameters $(p_1, p_2, r_1, r_2, \varphi)$ is used to roughly determine the pupillary boundary. Here, (p_1, p_2) represents the center of the ellipse; r_1, r_2 denote the semimajor axis and semiminor axis, respectively; and φ denotes the orientation of the ellipse. By limiting the values of the semimajor and semiminor axes to a specified range, the other parameters are iteratively varied with a small step size of three pixels to increase the size of the ellipse. At every iteration, a fixed number of points are randomly chosen on the circumference of the ellipse, and the total intensity difference between the chosen points and the center of the ellipse is computed. The boundary with the maximum intensity variation is chosen as the pupillary boundary. A rough contour of the limbus boundary is also determined in the same manner, with a different set of parameters for the semimajor and semiminor axes. Figure 13.20 shows an iris image with the rough contours of the pupillary and limbus boundaries obtained using the elliptical model.

Once the rough contours of the iris boundaries are obtained, the level set approach is applied to determine the precise contours. In the level set approach, the rough contour determined using the elliptical model is used as an active contour C , which can be represented as the zero level set $C(t) = \{(x, y) \mid \phi(t, x, y) = 0\}$ of a level set function $\phi(t, x, y)$. The objective of the level set approach is to determine a curve within the level set function, such that the total energy of the curve is minimum. The total energy of the curve is defined by the equation

$$\varepsilon(\phi) = \mu\rho(\phi) + \varepsilon_{g,\lambda,v}(\phi), \quad (13.14)$$

where the parameters $\varepsilon_{g,\lambda,v}(\phi)$ and $\rho(\phi)$ denote the external and the internal energies of the curve, respectively, and $\mu > 0$. The external energy parameter depends on the image data and drives the zero level set toward the boundary of the desired contour. On the other hand, the internal energy parameter helps in penalizing the deviation of the level set function, ϕ , from the signed distance function during the evolution of the curve.

The internal energy term in the above equation is further defined as

$$\rho(\phi) = \int_{\Omega} \frac{1}{2} (|\nabla\phi| - 1)^2 dx dy, \tag{13.15}$$

where Ω represents the image domain.

Similarly, the external energy term $\varepsilon_{g,\lambda,v}(\phi)$ can be further defined as

$$\varepsilon_{g,\lambda,v}(\phi) = \lambda L_g(\phi) + v A_g(\phi), \tag{13.16}$$

where $\lambda > 0$ and v are constants. The term g denotes an edge detector function and is defined as

$$g = \frac{1}{1 + |\nabla G_{\sigma} * I|^2}, \tag{13.17}$$

where G_{σ} denotes the Gaussian kernel with a standard deviation of σ and I denotes the image. The term $L_g(\phi)$ is used to measure the length of the zero level set curve of ϕ and is given by

$$L_g(\phi) = \int_{\Omega} g \delta(\phi) |\nabla\phi| dx dy, \tag{13.18}$$

where δ is the univariate Dirac function. The term $A_g(\phi)$ is used to speed up the curve evolution and is defined as

$$A_g(\phi) = \int_{\Omega} g H(-\phi) dx dy, \tag{13.19}$$

where H is the Heaviside function.


The desired evolution equation of the level set function can be obtained by determining the value of $\frac{\partial\phi}{\partial t}$, using the following equation:

$$\frac{\partial\phi}{\partial t} = -\frac{\partial\varepsilon}{\partial\phi}, \tag{13.20}$$

where $\frac{\partial\varepsilon}{\partial\phi}$ represents the Gateaux derivative of ε . The value of ϕ that minimizes the total energy function can be determined by satisfying the Euler-Lagrange equation, $\frac{\partial\varepsilon}{\partial\phi} = 0$. The Gateaux derivative of the functional ε can be written as follows:

$$\frac{\partial\varepsilon}{\partial\phi} = -\mu \left[\Delta\phi - \text{div} \left(\frac{\nabla\phi}{|\nabla\phi|} \right) \right] - \lambda \delta(\phi) \text{div} \left(\left(g \frac{\nabla\phi}{|\nabla\phi|} \right) - v g \delta(\phi) \right). \tag{13.21}$$

Fig. 13.21 Final output obtained using the variational level set approach, where both the iris boundaries are precisely determined (Image source: Roy et al. [24] ©Elsevier)



Thus, the desired equation of the level set function can be defined as

$$\frac{\partial \phi}{\partial t} = \mu \left[\Delta \phi - \operatorname{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) \right] + \lambda \delta(\phi) \operatorname{div} \left(\left(g \frac{\nabla \phi}{|\nabla \phi|} \right) + \nu g \delta(\phi) \right). \quad (13.22)$$

The last two terms on the right-hand side of the above equation represent the gradient flows of the energy functional. These terms help in driving the zero level curve toward the boundaries of the required surface. The Dirac function $\delta(x)$ in the above equation is defined as

$$\delta_\varepsilon(x) = \begin{cases} 0, & |x| > \varepsilon \\ \frac{1}{2\varepsilon} \left[1 + \cos\left(\frac{\pi x}{\varepsilon}\right) \right], & |x| \leq \varepsilon. \end{cases} \quad (13.23)$$

For the active contour ϕ , the curve evolution process for the pupillary boundary is carried out within a small range of ± 10 pixels from the rough contour. For the limbus boundary, this range is increased to ± 20 pixels. The curve evolution process is carried out from the outside of the approximated pupil boundary to avoid the effect of specular reflections. On the other hand, the process is carried out from the inside of the approximated limbus boundary to reduce the effect of eyelids and the eyelashes. Figure 13.21 shows the final output obtained by the variational level set approach, in which both the iris boundaries are precisely detected. Figure 13.22 shows the output of iris segmentation using the variational level set approach on some nonideal iris images.

13.4.3 Fourier-Based Approximation

Daugman [7] suggested the use of the Fourier series approximation in order to deduce the boundaries of the iris. The benefit of such an approximation is that the resulting output satisfies the following expectations:




Fig. 13.22 Output of iris segmentation scheme using variational level set approach on some nonideal iris images from the UBIRIS Version 1 dataset. Notice that the segmentation results are fairly precise (Image source: Roy et al. [24] ©Elsevier)

1. *Completeness*: An iris image may exhibit interruptions in its boundaries (e.g., interruptions caused by the specular reflections and eyelids to the pupillary and limbus boundaries, respectively). Therefore, the boundary detected by the segmentation algorithm must be robust to such interruptions.
2. *Closure*: Both the iris boundaries detected by the segmentation algorithm are expected to continue their trajectory across the interruptions on a principled basis and form closed curves.

In this technique, given an image I , the coarse contour of the iris boundary is determined using active contours. Let the coarse iris contour be represented by N regularly spaced angular samples, given by $\{r_\theta\}$, $\theta = 0$ to $\theta = N - 1$. From this coarse contour, the corresponding iris boundary $\{R_\theta\}$, $\theta = 0$ to $\theta = N - 1$, that satisfies the above two conditions has to be determined. This can be achieved by the Fourier series approximation of the coarse contour data expressed as follows:

$$R_\theta = \frac{1}{N} \sum_{k=0}^{M-1} C_k \exp^{2\pi i k \theta / N}, \quad (13.24)$$


where $\{C_k\}$ represents a set of M discrete Fourier coefficients, for $k = 0$ to $k = M - 1$, determined by the following equation:

$$C_k = \sum_{\theta=0}^{N-1} r_\theta \exp^{-2\pi i k \theta / N}. \quad (13.25)$$

Generally, the zeroth-order Fourier coefficient (or the DC term C_0) describes the average curvature of the obtained boundary. Since this technique is used for determining the iris boundaries, the zeroth-order coefficient determines the approximate radius of the output contour.

The value of M represents the number of active Fourier coefficients that are used in the approximation process. This value represents the number of degrees of freedom for the shape model of the boundary. In the most simplest scenario, i.e., $M = 1$, the model of the boundary will be circular. The value of M also acts as a trade-off between the preciseness of the shape versus the strictness of the constraints (which corresponds to the complexity of the model). A strict set of

Fig. 13.23 Iris segmentation using Daugman's Fourier approximation approach (Image source: Daugman [7] ©IEEE. Image has been edited for clarity)



constraints leads to a complex model, while a weak set of constraints leads to a simple model. Daugman suggests that $M = 17$ is a good choice for the pupillary boundary and $M = 5$ for the limbus boundary. The two different choices for M are supported by the computer vision principle that strong data may be modeled with only weak constraints, while weak data should be modeled with strong constraints. The limbus boundary is considered as weak data because the occlusions caused by the eyelids and eyelashes are generally high. On the other hand, the pupillary boundary is considered as stronger data because the interruptions caused by the specular reflections are relatively minimal.

Figure 13.23 shows the segmentation output obtained using the Fourier-based approximation. The lower left corner of the image shows a *snake* that corresponds to the limbus boundary. The *snake* on the lower right corner of the image corresponds to the pupillary boundary. Both the snakes consist of two components: (a) a fuzzy ribbonlike data distribution that corresponds to the coarse contour, r_θ , and (b) a dotted curve that corresponds to the Fourier series approximation, R_θ . The characteristics of these snakes can be listed as follows:

1. The endpoints for both the snakes meet at the 6 o'clock position.
2. The thickness of each snake roughly corresponds to the sharpness of the corresponding edge.
3. The more uninterrupted an iris boundary is, the flatter and straighter the snake will be.

From the lower right corner of the figure, it can be noticed that the curve r_θ corresponding to the limbus boundary exhibits interruptions caused by the eyelid. However, the curve R_θ continues its trajectory even across the interruptions, proving the effectiveness of the approach.

13.5 Role of Image Acquisition on Iris Segmentation

Image acquisition plays an important role in the performance of an iris recognition system. Early iris recognition systems required significant cooperation from the subjects during image acquisition. This helped in acquiring good quality iris images,



Fig. 13.24 Conventional iris image acquisition system requiring considerable user cooperation (Image source: <http://www.life.com/image/1668585>)

with minimal or no occlusions. As a trade-off, algorithms with low computational complexity [6,26] were sufficient for the task of iris segmentation. With an increased demand for accurate iris recognition under practical scenarios (e.g., from a distance, under covert conditions), the requirements imposed during image acquisition are being relaxed. While this can impart flexibility to the image acquisition process, the quality of the acquired images can reduce drastically. In such cases, complex algorithms may be required to perform segmentation while being robust to the factors mentioned in Sect. 13.2. This section discusses some iris segmentation algorithms, grouped according to the image acquisition conditions they are designed for.

13.5.1 Segmenting Images Acquired Under Highly Constrained Conditions

Images acquired under constrained conditions are often expected to be of high quality. This is because the user is typically still and cooperatively offers the iris images. Figure 13.24 shows an image acquisition system which requires significant user cooperation. In such scenarios, the classical algorithms described in Sect. 13.3 have been observed to provide good segmentation performance.

13.5.2 Segmenting Images Acquired Under Less Constrained Conditions

The public usage iris recognition system proposed by Negin et al. [18] may be considered to be one of the first systems that attempted to relax the image acquisition




Fig. 13.25 Image acquisition setup for the public use iris recognition system (Image source: Negin et al. [18] ©IEEE)

conditions. While other iris recognition systems prior to this work required the user to be in close proximity to the sensor, the proposed system allowed a standoff distance of up to 3 ft. Furthermore, the proposed system allowed for an easy *public setup*, compared to its counterparts. An example of such a public setup could be the use of iris recognition technology to access an automated teller machine (ATM). Even if the user at the ATM does not stand completely still, the system would be able to perform iris recognition during the transaction.

To use the system, the user was required to be reasonably cooperative and focus their gaze toward the system camera. The system would provide feedback to facilitate easy image acquisition of the user's eye. A high-resolution video image of one eye of the user would be captured and used for recognition purposes. The image acquisition setup of the system is shown in Fig. 13.25.


The various steps involved in the working of this system, along with the functioning of the individual components, are provided below:

1. The user stands in front of the system, with a maximum allowable stand-off distance of 3 ft.
2. A wide field of view (WFOV) camera pair is used to capture an image of the user's torso, as shown in Fig. 13.26. The system then applies an image-processing algorithm to locate the eyes of the user.
3. A narrow field of view (NFOV) camera pair is separately used to focus on the eye region. The system controls a gaze director to aid the user to look toward the camera. A pan-tilt mirror is simultaneously used to direct the optical axis

Fig. 13.26 Image of a user's torso acquired by the WFOV camera (Image source: Negin et al. [18] ©IEEE)



Fig. 13.27 Image of the user's eye, acquired by the NFOV camera (Image source: Negin et al. [18] ©IEEE)



of the NFOV camera pair to ensure that the user's eye is focused properly. As infrared illuminators are used during image acquisition, the system could operate even if the subject wore eyeglasses, contact lenses, or in a nighttime environment. Figure 13.27 shows a sample image acquired at this stage.

4. A circular grid was used as a guide by the system to localize the iris region in the image acquired by the NFOV camera. The use of a circular grid simultaneously allowed for the exclusion of noisy regions such as the pupil, sclera, and the eyelids. The region lying within the grid was used for encoding and recognition. A sample image showing the circular grid for localizing the iris region is shown in Fig. 13.28.

While this system relaxed the acquisition conditions only moderately, it is considered to be significant in the field of iris biometrics for the following reasons:

1. This is one of the earliest works related to successful iris recognition from a distance.
2. This work highlights the dependency of iris segmentation on the eye detection scheme. When iris images are acquired from a distance, it has to be noted that eye detection has to be accurately performed to obtain good segmentation performance.




Fig. 13.28 *Left:* Circular grid used for iris localization. *Right:* Iris image with the circular grid overlaid (Image source: Negin et al. [18] ©IEEE)

13.5.3 Segmenting Images Acquired Using IOM (*Iris on the Move*) Systems

The “Iris on the Move” (IOM) system developed by Matey et al. [17] is considered to be another major development in the field of iris recognition from a distance. The IOM system significantly reduces the constraints on the position and the motion of a user during image acquisition. Such a flexibility is made possible by using an improved image acquisition system that uses high-resolution cameras and video-synchronized illumination mounted on a minimally confining portal.

The setup of the IOM system consists of a walk-through portal, similar to a metal detector. Near infrared (NIR) illumination sources and high-resolution image sensors are fixed to the portal. Images are acquired when a user walks through the portal at a normal walking speed (< 1 ms). Standoff distances up to 3 m is possible, with a minimum requirement that the user be moderately cooperative. The system can acquire images even when a user wears eyeglasses or contact lenses, but cannot see through sunglasses. The camera used in the IOM system was Securimetrics PIER 2.3, which can acquire iris images with an approximate diameter of 200 pixels. As the heights of subjects can vary by a large factor, the system uses a set of cameras instead of one single camera. The setup of an IOM system is shown in Fig. 13.29.


The raw data acquired by the IOM system is typically a set of high-resolution facial images ($2,048 \times 2,048$ pixels). This is possible due to the high-resolution and wide field of view of the cameras involved. The key requirement of the IOM system is that it should perform image acquisition, segmentation, and recognition in real time. The IOM system is considered to be an industrial application of the iris recognition technology that is expected to serve a large volume of people in short time. To reduce the computational time and processing speed of the system, the authors suggest a segmentation routine which is significantly different from the other segmentation schemes in the literature.

An iris image acquired by the IOM system typically exhibits a pattern of specular reflections on the iris. These specular reflections are caused by the near infrared (NIR) illumination system used during image acquisition. The pattern of the specular reflections are strongly dependent on the pattern in which the illuminators



Fig. 13.29 Image acquisition setup for the Iris on the Move system (Image source: Matey et al. [17] ©IEEE)


Fig. 13.30 Image acquired using the IOM system, exhibiting a specular reflection pattern that corresponds to the illuminator pattern (Image source: Matey et al. [17] ©IEEE)



are arranged. Figure 13.30 shows a sample NIR image exhibiting the specular reflections that correspond to the illuminator pattern.

Instead of ignoring specular reflections as noise, the segmentation scheme in the IOM system in fact makes effective use of the pattern. A match filter is applied to the captured image producing the highest responses on the in-focus specularities. Once the specularities are highlighted, a thresholding scheme is used to binarize the image. This process yields an image showing only the locations of the specularities. The border, width, height, and center of each specularity is determined in the binary image. As the illuminator pattern is predetermined, specularity patterns that are inconsistent with the illuminator pattern (or those that are oddly shaped) are ignored. For the images that have strong correspondence between the specular reflection and

Fig. 13.31 Iris segmentation using the specular reflection pattern in the image. Once the specular reflection pattern is detected, a specified area around the pattern is used for unwrapping and feature extraction (Image source: Matey et al. [17] ©IEEE)



the illuminator patterns, the iris regions around the specular reflections are extracted. Figure 13.31 shows an iris image in which segmentation is performed using the specular reflection patterns.

The contribution of the IOM system in the field of iris recognition is significant for the following reasons:

1. IOM system is the first image acquisition and recognition system that can work on iris images obtained from users walking at a normal pace.
2. The system allows for real-time recognition while relaxing the constraints imposed on the user, compared to its predecessors.

13.6 Segmenting Visible Wavelength Iris Images

Most iris recognition systems acquire input images using sensors that operate in the infrared wavelength. This is due to the fact that the complex texture of the iris is more easily discernible when imaged under infrared lighting. However, it may not be possible to acquire or operate with infra-red iris images at all times. Some of the reasons that support iris image acquisition in visible wavelength are provided below:

1. The current sensor and illuminator technology makes it difficult to acquire iris images from a distance using sensors that operate in the near infrared wavelength.
2. Surveillance cameras usually work in visible wavelength. When using face and iris recognition systems on the data acquired from surveillance cameras, the iris data will be in the visible wavelength.

A majority of existing segmentation techniques use near infrared images as their input. Such images typically exhibit higher contrast between the pupil and the iris regions and induce the usual option of determining the pupillary border. In contrast, visible wavelength images usually exhibit less contrast between the pupil and the iris. This supports the inversion of the order in which the iris boundaries are segmented

Fig. 13.32 An iris image acquired in the visible wavelength. Notice that the intensity contrast between the pupil and the iris is low



for visible images. An iris image acquired using a sensor that operates under visible wavelength is shown in Fig. 13.32.

Proenca suggests an approach [20, 21] that performs automatic segmentation of the iris images acquired in the visible wavelength. Furthermore, the approach can perform iris segmentation on images that are acquired at a large imaging distance (4–8 m), on the move, and under varying lighting. From the input image, the proposed technique first detects the sclera, as it is the most distinguishable region even under varying lighting conditions. Then, the fact that sclera mandatorily lies next to the iris is taken into account, and the iris regions are detected. A neural pattern recognition approach is later used to perform the segmentation.

The entire process can be broken down into three stages: detection of the regions that correspond to the iris; segmentation of the iris; and detection and elimination of the noisy regions. Given an image I , the sclera region is first detected. This is because in some images, iris detection is rather difficult. However, due to the naturally distinguishable appearance (color) of the sclera, it can be detected first by an approach which analyzes the color spaces of the image. Based on an empirical analysis, the author suggests three color components: hue (h), blue (cb), and red chroma (cr) that can characterize information about the sclera. The contrast between the sclera and the remaining parts of the eye can be maximized using these three color components.

Given an iris image I , a 20-dimensional feature vector is constructed for each pixel, which can be denoted by the following expression:

$$\{x, y, h_{0,3,7}^{\mu, \sigma}(x, y), cb_{0,3,7}^{\mu, \sigma}(x, y), cr_{0,3,7}^{\mu, \sigma}(x, y)\} \quad (13.26)$$

where x and y denote the position of the pixel and $h()$, $cb()$, and $cr()$ denote the hue, blue, and red chroma components of the image at that pixel. The subscript denotes the radii of the circle that is centered at the pixel. The parameters μ and σ denote the mean and the standard deviation, respectively, of the set of pixels which fall within those circular regions. For example, the term $cb_{0,3,7}^{\mu, \sigma}(x, y)$ means that six features were extracted from regions of the blue color component: three averages and three standard deviations, computed locally within regions of circles with radii 0, 3, and 7, centered at the considered pixel (x, y) . Once the feature vectors for all the image pixels are calculated for all the images, a neural network classifier is used to obtain the sclera map. This map indicates the location of sclera in the image.

To detect the iris, the information obtained using the sclera map is used. The author suggests that a pixel which lies within iris boundaries, when frontally imaged, will have similar number of sclera pixels on both sides. On the other hand, if the iris is off-angled, then the number of pixels of sclera on one side will be more than the number of pixels on the other side. To detect the iris pixels, another feature vector is generated. The data obtained from the sclera detection stage is used to obtain a new feature, referred to as “proportion of sclera” $p(x,y)$, for each image pixel. This feature helps in measuring the proportion of pixels that belong to the sclera in a direction d , with respect to the pixel at location (x,y) . The notation used for the directions are \uparrow for north, \downarrow for south, \leftarrow for east, and \rightarrow for west. The feature vectors for each pixel are generated using the proportion of sclera information as follows:

$$p_{\leftarrow}(x,y) = \mu(sc((1,y-1),(x,y))), \quad (13.27)$$

$$p_{\rightarrow}(x,y) = \mu(sc((x,y-1),(w,y))), \quad (13.28)$$

$$p_{\uparrow}(x,y) = \mu(sc((x-1,1),(x,y))), \quad (13.29)$$

$$p_{\downarrow}(x,y) = \mu(sc((x-1,y),(x,h))), \quad (13.30)$$

where $sc((.,.),(.,.))$ denotes the regions of the image cropped from the detected sclera, delimited by the top-left and bottom-right corner coordinates. w and h denote the width and height, respectively. The value of $p()$ is set to 0 for all sclera pixels. The “proportion of sclera” values, pixel positions, local image saturation, and blue chrominance are then used to form a new feature vector represented as

$$\{x,y,s_{0,3,7}^{\mu,\sigma}(x,y),cb_{0,3,7}^{\mu,\sigma}(x,y),p_{\leftarrow,\rightarrow,\uparrow,\downarrow}(x,y)\} \quad (13.31)$$

where $s()$ and $cb()$ denote saturation and blue chrominance. Once again, the choice of color components is based on empirical evaluation.

Multilayered perceptron feed forward neural networks with one hidden layer are used by both classification stages. The neural network is trained using sample images, and its output on a test image is considered to effectively decide the boundaries of the iris. Once the set of image pixels that correspond to a noise-free iris are identified, the goal is to determine the contours of the pupil and sclera of the iris. For this purpose, shape parameterization techniques are used.

13.7 Performance Evaluation of a Few Iris Segmentation Techniques

To provide an understanding of the iris segmentation performance, three techniques were chosen. The three algorithms are (a) integro-differential operator, (b) Hough transformation, and (c) geodesic active contours. The segmentation performance is evaluated on the ICE database [19], which contains a total of 2,953 iris images

Table 13.1 Segmentation accuracies of the three techniques evaluated in this chapter

Segmentation technique	Number of input images	Number of correctly segmented images	Segmentation accuracy (%)
Integro-differential operator	2,953	2,597	87.9
Hough transform	2,953	2,632	89.1
Geodesic active contours (GAC)	2,953	2,699	91.3

corresponding to both left and right eyes. The images are in the TIFF format, with a resolution of 640×480 pixels. The quality of the images is reasonably good, as they were collected from cooperative users. However, some images exhibit eyelid or eyelash occlusions. The performance of an iris segmentation technique was measured by computing the segmentation accuracy as follows:

$$\text{Segmentation accuracy} = \frac{\text{Number of correctly segmented images}}{\text{Number of input images}} \times 100 \quad (13.32)$$

Table 13.1 lists the segmentation accuracies of the three techniques. From the results, it can be observed that geodesic active contours provide better segmentation performance when compared to the classical approaches. However, it has to be noted that the computational expense of the classical approaches is much lower than that of the geodesic active contours. Thus, it is necessary to understand the computational demands and operational requirements of an application prior to engaging a specific iris segmentation algorithm. It may be possible to dynamically determine *which* iris segmentation algorithm is appropriate to be used based on characteristics of the image to be segmented. Further, the outputs of multiple segmentation algorithms may be combined to generate a single hypothesis for the boundaries of the iris.

13.8 Approaches to Refine Iris Segmentation

One of the major concerns in iris segmentation is the over- or under-segmentation of the iris boundaries. Over-segmentation refers to the situation where the radius of the detected iris boundary is larger than that of the actual boundary. On the other hand, under-segmentation refers to the situation where the radius of the detected iris boundary is smaller than the actual boundary. Figure 13.33 shows a sample image for each case.

In both cases, the offset between the actual iris boundary and the detected iris boundary is not large. However, such minor offsets can significantly lower iris recognition performance. This is due to the inclusion of noise or regions that do not contain discriminatory texture information (e.g., sclera, eyelashes). In this section, some approaches are discussed that attempt to refine segmentation by operating on finer details in the vicinity of the iris.




Fig. 13.33 *Left*: Over-segmented limbus boundary. *Right*: Under-segmented limbus boundary




Fig. 13.34 Unwrapped iris images containing eyelash occlusion (Image source: Zhang et al. [28] ©IEEE)

13.8.1 Eyelash Removal in Iris Images

One of the many factors that affect iris recognition performance is the occlusion caused by eyelashes. While some approaches ignore the iris regions occluded by eyelashes [14], others detect the eyelashes and mask them while encoding [5]. On the contrary, the technique proposed by Zhang et al. [28] attempts to restore the iris regions that are occluded by the eyelashes. A nonlinear conditional directional filtering approach is used to perform the restoration. The proposed technique is an iterative approach involving the following steps: (a) detecting the pixels that correspond to eyelash occlusion, (b) detecting the direction of the eyelash that causes the occlusion, (c) local filtering of occlusion region in the direction that is perpendicular to the eyelash, and (d) restoring the pixel intensity by using a 1D median filter.

The proposed technique is more suited for unwrapped or normalized iris images, rather than the original iris images. The unwrapped iris images in this work are of size 512×80 pixels, and the top 48 rows of pixels nearest to the pupil are used by the eyelash removal technique. Figure 13.34 shows eyelash occlusion in an unwrapped iris image.

An eyelash is expected to cause a discontinuity along its edges. Therefore, given an unwrapped iris image I , an eyelash is detected, and its direction is estimated by using an edge filter. For this purpose, a 3×3 Sobel edge filter is applied on the image. A Sobel edge filter is shown in Fig. 13.35.

-1	-2	-1	z ₁	z ₂	z ₃	-1	0	1
0	0	0	z ₄	z ₅	z ₆	-2	0	2
1	2	1	z ₇	z ₈	z ₉	-1	0	1

Fig. 13.35 *Left:* x derivative for a Sobel edge filter. *Center:* image region under consideration. *Right:* y derivative for a Sobel edge filter

For each pixel, the gradients in the x - and y -directions $[G_x, G_y]$ can be determined by the following equations:

$$G_x = (z_7 + 2z_8 + z_9) - (z_1 + 2z_2 + z_3) \tag{13.33}$$

$$G_y = (z_3 + 2z_6 + z_9) - (z_1 + 2z_4 + z_7) \tag{13.34}$$

Similarly, the magnitude of the gradient at the center of the mask, Grad, is computed as

$$\text{Grad} = \sqrt{(G_x^2 + G_y^2)}. \tag{13.35}$$

The local gradient direction that is perpendicular to the edge can then be determined by

$$\theta = \arctan(G_y/G_x). \tag{13.36}$$

To determine if a pixel is occluded or not, a window of size $m \times n$ is centered at a pixel, and the gradient direction variance for the r pixels that lie within the window and have a Grad value above a specific threshold is computed as follows:

$$\text{Var}_{-} \text{Grad} = \frac{1}{r-1} \sum_{i=1}^r (\theta_i - \bar{\theta})^2. \tag{13.37}$$

A strong edge is indicated if the gradient direction has a small variance, and the pixel is classified as being affected by an eyelash. To restore such a pixel, a 1D median filter of length L is applied along the direction θ . This process outputs an estimate of the value of the image with the eyelash removed.

To avoid incorrectly filtering non-eyelash pixels, pixel alteration is carried out only if the change in the pixel exceeds a certain threshold, related to the total variance of the image. For this purpose, a parameter Recover is computed as follows:

$$\text{Recover} = \text{Diff} - k * \text{Var}(\text{Image}) \tag{13.38}$$

where Diff represents the difference in the intensity between the filtered and the unfiltered pixel, Var(Image) represents the intensity variance of the whole unfiltered image, and k denotes parameter used to tune the threshold. The pixel is replaced by the filtered value, only if Recover is positive. By using this approach, the visual appearance is not significantly changed, but the recognition performance is improved. Figure 13.36 shows an unwrapped iris image before and after applying the approach.



Fig. 13.36 *Top*: image affected by eyelash occlusion. *Bottom*: output obtained by applying the eyelash removal technique (Image source: Zhang et al. [28] ©IEEE)

13.8.2 Improving Daugman's Classical Segmentation Algorithm

Libor Masek's MATLAB package [16] is one of the most widely used open source implementation for iris segmentation. For good quality iris images acquired under regular imaging conditions, Libor Masek's implementation results in good segmentation performance. However, Liu et al. [13] showed that the segmentation performance can be further improved by incorporating minor modifications to Libor Masek's implementation. The two most significant modifications that were used to improve the segmentation performance of Libor Masek's implementation are as follows:

1. Reversal of the detection order of the iris boundaries: In Masek's implementation, the limbus boundary is detected first, followed by the pupillary boundary. However, by reversing the detection order, slightly better segmentation performance can be observed. This is based on the fact that the pupillary boundary exhibits strong intensity variation at its boundary, when compared to that of the limbus boundary. By reversing the order, the dependency problem can be minimized.
2. Eyelid detection: Libor Masek's implementation models the eyelids as two horizontal lines. As a result of such an approximation, some of the iris texture can be occluded. The authors suggest splitting the top and bottom eyelid regions into two different portions each as shown in Fig. 13.37 and then performing eyelid detection. As a result of this, the eyelids will no longer be approximated as straight lines but will appear curvy. This modification avoids unnecessary occlusion of the iris during eyelid estimation. Figure 13.38 shows the difference observed in eyelid detection using the proposed modification.

13.9 Predicting Errors in Iris Segmentation

The performance of an iris recognition system is highly dependent on the output of iris segmentation. If an incorrectly segmented region is used for recognition, the lack of rich distinctive texture can reduce recognition performance. Thus, designing an

Fig. 13.37 Splitting of the eyelid into four portions to allow better detection of eyelid (Image source: Liu et al. [13] ©IEEE)





Fig. 13.38 *Left:* Eyelid detection using Libor Masek's approach. *Right:* Improved eyelid detection using the proposed modifications (Image source: Liu et al. [13] ©IEEE)

algorithm that can examine the input image and predict in advance if segmentation is likely to fail or not would be beneficial. Such a scheme can be used to provide feedback during image acquisition, requiring the user to provide a better, more useful image. When acquiring a new image is not a possibility, the image can at least be flagged to avoid being used for recognition purposes. Manual segmentation could then be used to further process such an image. Another possibility is to design an algorithm that can evaluate the output of the segmentation routine and determine if the segmentation has failed or not.

Some algorithms [3, 10] compute the quality of the iris region using local image analysis. However, such algorithms require at least a coarsely segmented iris. If the

segmentation is improper, the quality estimate will be incorrect, thus defeating the motive for such an exercise.

Kalka et al. [11] propose an algorithm which evaluates the output of the pupil and iris segmentation routines. The algorithm is based on combining probabilistic intensity features with geometric features to generate scores that indicate the success of segmentation. A decision tree-based machine learning approach is used to render a binary decision: success or failure. The proposed method, therefore, predicts whether the output of segmentation is good (both the pupil and iris boundaries were correctly estimated) or bad (at least one of the boundaries was incorrectly estimated).

The following measures are taken into account:

1. Pupil segmentation measure: One of the major concerns here is related to the over- or under-segmentation of the pupil. To take into account such problems, a probabilistic model is used to fit the segmentation output for the pupil. The output of the model is then used to generate an over-segmentation or under-segmentation score.
2. Geometric iris measure: This is based on the fact that the limbus and pupillary boundaries are actually concentric (circles or ellipses) when they are frontally imaged. Therefore, a measure based on eccentricity and concentricity of the boundaries is taken into account.

Given an iris image segmentation output, the goal is to assign a score to that output, which indicates the “goodness” of segmentation based on the above factors. First, the pupil boundary is checked to establish whether it is over- or under-segmented. For this, consider an image I , and let its pupillary contour be represented as \bar{x} . The task is to decide whether the pixels lying within the boundary \bar{x} actually belong to the pupil or not. To this end, a probabilistic model is associated with \bar{x} , which formulates a likelihood ratio test, $\Lambda(\bar{x})$, as follows:

$$\Lambda(\bar{x}) = \frac{P(\bar{x} | H_1)}{P(\bar{x} | H_0)} \geq \eta \quad (13.39)$$

where $H_1 : \bar{x}$ corresponds to a pupil pixel and $H_0 : \bar{x}$ corresponds to a non-pupil pixel. Based on empirical evaluation, the authors use a Gamma distribution $\Gamma(k, \Theta)$, for $P(\bar{x} | H_1)$, and a Gaussian distribution $N(\mu, \sigma^2)$ for $P(\bar{x} | H_0)$, respectively. The value of the shape parameter k is fixed to 1. To obtain the scale parameter value Θ , the spatial histogram of the image intensities in the region of interest is computed. From the histogram, the scale parameter is estimated as

$$\hat{\Theta} = \sum_{i=0}^{\text{Bin}(P_t)} x_i w_i \quad (13.40)$$

where P_t denotes the threshold used to constrain the size of the pupil region, x_i is a gray-level bin from the histogram of the region of interest, and w_i is the weight

associated with bin x_i . In other words, Θ is obtained by summing the product of the gray-level bins and the associated weights until the bin corresponding to P_t is reached. The parameters of the Gaussian are estimated using the following equations:

$$\hat{\mu} = \sum_{i=\text{Bin}(P_t)+1}^{\text{Bin}(I_t)} x_i w_i \quad (13.41)$$

$$\hat{\sigma}^2 = \sum_{i=\text{Bin}(P_t)+1}^{\text{Bin}(I_t)} w_i (x_i - \hat{\mu})^2 \quad (13.42)$$

where I_t is a threshold used to constrain the size of the iris region, x_i is the gray-level bin in the histogram corresponding to the region of interest, and w_i is the weight associated with bin x_i . Thresholds P_t and I_t are determined experimentally. Every pixel within the pupil boundary is assigned a 0 or 1, based on the likelihood ratio test. Once values are assigned to all pixels in the region of interest, the over-segmentation measure P_{over} is computed as the ratio of the number of 0s to that of 1s.

To measure the under-segmentation of the pupil boundary, an iterative approach is employed where the estimated pupil radius (or major and minor axes in case of an ellipse) is increased, and it is determined whether the pixels inside the expanded region are pupil pixels by the same approach as above. This process is continued until the pupil radius reaches the size of the iris radius or the ratio of pupil to non-pupil pixels is less than 20%. The use of a threshold prevents the influence of heterogenous factors such as dark eyelashes/eyelids and reduces unnecessary computations. The final under-segmentation score is defined as follows:

$$P_{\text{under}} = \frac{P_{\text{over}}}{P_{\text{over}} + P_{\text{est_under}}} \quad (13.43)$$

$P_{\text{est_under}}$ is the total number of estimated pupil pixels over all iterations.

To obtain yet another score value for the accuracy of segmentation, the concentricity and eccentricity of the iris and pupil boundaries are used. These values are measured using the following equations:

$$I_C = \sqrt{(p_x - i_x)^2 + (p_y - i_y)^2} \quad (13.44)$$

$$I_E = \sqrt{(p_x - i_x)^2 + (p_y - i_y)^2} + \arccos\left(\frac{b_i}{a_i}\right) * 100 + \arccos\left(\frac{b_p}{a_p}\right) * 100 \quad (13.45)$$

where (p_x, p_y) are the pupil center coordinates, (i_x, i_y) are the iris center coordinates, b_i and a_i are the semiminor and semimajor axes for the iris ellipse, respectively, and b_p and a_p are the semiminor and semimajor axes for the pupil ellipse, respectively.

Once both the pupil segmentation score and the iris segmentation scores are available, a Naive-Bayes tree classifier is used to generate a final score. This final score is binarized to indicate the performance of iris segmentation.

13.10 Conclusion

Segmentation plays a significant role in the recognition performance of an iris biometric system. A majority of algorithms proposed in the past assume that the input image of the eye is of good quality. However, with an increasing need for faster and accurate recognition systems, the conditions imposed during image acquisition have been relaxed. As a result, the quality of images may be reduced, forcing the need for robust segmentation algorithms that can tolerate a wide variety of image degradations (see Table 13.2). A number of iris segmentation algorithms were discussed in this chapter. From this discussion, it is evident that (a) iris segmentation is an important problem, (b) substantial effort has been invested by researchers in solving the problem of iris segmentation under different scenarios, (c) the performance of an iris segmentation algorithm depends upon a number of

Table 13.2 Examples of a few iris segmentation techniques proposed in the literature

Authors	Segmentation
Daugman	Integro-differential operator
Wildes [26]	Edge detection and Hough transform
Bole and Boashash	Edge and contour detection
Masek and Kovesi [16]	Edge detection and Hough transform
Ma et al. [15]	Edge detection and Hough transform
Lim et al.	Edge detection and Hough transform
Ma et al. [14]	Edge detection and Hough transform
Huang et al. [8]	Edge detection and Hough transform
Huang et al. [8]	Phase congruency and Hough transform
Yuan et al.	Edge detection and Hough transform
Dorairaj et al.	Integro-differential operator
Kumar et al. [12]	Hough transform
Schuckers et al. [25]	Active shape models
Du et al.	Scale invariant feature transform
Pundlik et al. [22]	Graph cuts
Shah and Ross	Geodesic active contours
Daugman [7]	Fourier based approximation
Roy et al. [23, 24]	Variational level sets
Zuo and Schmid [29]	Ellipse fitting
He et al.	Pulling and pushing model

Notice that a number of them are based on edge detection and the Hough transform. Some other examples can be found in [2]

intrinsic image characteristics including spatial resolution of the iris and degree of occlusion, (d) the computational demands of different iris segmentation algorithms can vary considerably, and (e) evaluating the output of an iris segmentation routine and combining the outputs of multiple iris segmentation methods are ongoing activities in the iris biometric community.

Acknowledgements This work was partially funded by US National Science Foundation CAREER Grant No. IIS 0642554.

References

1. Baker, S., et al.: Degradation of iris recognition performance due to non-cosmetic prescription contact lenses. *Comput. Vis. Image Underst.* **114**, 1030–1044 (2010)
2. Bowyer, K., Hollingsworth, K., Flynn, P.: Image understanding for iris biometrics: a survey. *Comput. Vis. Image Underst.* **110**(2), 281–307 (2008)
3. Chen, Y., Dass, S., Jain, A.: Localized iris image quality using 2-d wavelets. In: Tistarelli, M., Bigün, J., Grosso, E. (eds.) *Advances in Biometrics. Lecture Notes in Computer Science*, vol. 3832, pp. 373–381. Springer, Berlin/Heidelberg (2005)
4. Daugman, J.: How iris recognition works. *IEEE Trans. Circuits Syst. Video Technol.* **14**(1), 21–30 (2004)
5. Daugman, J.: The importance of being random: statistical principles of iris recognition. *Pattern Recognit.* **36**(2), 279–291 (2003)
6. Daugman, J.: How iris recognition works. *IEEE Trans. Circuits Syst. Video Technol.* **14**(1), 21–30 (2004). doi:<http://dx.doi.org/10.1109/TCSVT.2003.818350>.
http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=1262028
7. Daugman, J.: New methods in iris recognition. *IEEE Trans. Syst. Man Cybern. Part B Cybern.* **37**(5), 1167–1175 (2007)
8. Huang, J., et al.: Learning based resolution enhancement of iris images. In: *British Machine Vision Conference*, pp. 153–162. British Machine Vision Association, Great Britain (2003)
9. Illingworth, J., Kittler, J.: A survey of the Hough transform. *Computer Vis. Graph. Image Process.* **44**(1), 87–116 (1988)
10. Kalka, N.D., Zuo, J., Schmid, N.A., Cukic, B.: Image quality assessment for iris biometric. In: *Proceedings of the SPIE Conference on Biometric Technologies for Human Identification III*, pp. 62020D-1–62020D-11. SPIE, Bellingham (2006)
11. Kalka, N., Bartlow, N., Cukic, B.: An automated method for predicting iris segmentation failures. In: *IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems*, pp. 1–8. IEEE, Piscataway (2009)
12. Kumar, B., et al.: Biometric verification with correlation filters. *Appl. Opt.* **43**(2), 391–402 (2004)
13. Liu, X., Bowyer, K., Flynn, P.: Experiments with an improved iris segmentation algorithm. In: *Fourth IEEE Workshop on Advanced Technologies for Automatic Identification*, pp. 118–123. IEEE, Los Alamitos (2005)
14. Ma, L., et al.: Efficient iris recognition by characterizing key local variations. *IEEE Trans. Image Process.* **13**(6), 739–750 (2004)
15. Ma, L., et al.: Efficient iris recognition by characterizing key local variations. *IEEE Trans. Image Process.* **13**, 739–750 (2004)
16. Masek, L., Kovesi, P.: MATLAB source code for a biometric identification system based on iris patterns. Technical Report the School of Computer Science and Software Engineering, The University of Western Australia (2003)

17. Matey, J., et al.: Iris on the move: acquisition of images for iris recognition in less constrained environments. *Proc. IEEE* **94**(11), 1936–1947 (2006)
18. Negin, M., et al.: An iris biometric system for public and personal use. *Computer* **33**(2), 70–75 (2000)
19. Newton, E.M., Phillips, P.J.: Meta-analysis of third-party evaluations of iris recognition. *IEEE Trans. Syst. Man Cybern.* **39**(1), 4–11 (2009)
20. Proenca, H.: Iris recognition: a method to segment visible wavelength iris images acquired on-the-move and at-a-distance. In: *ISVC 2008: 4th International Symposium on Visual Computing 1*. Springer Lecture Notes in Computer Science #5358, pp. 731–742. Springer, Berlin/New York (2008)
21. Proenca, H.: Iris recognition: on the segmentation of degraded images acquired in the visible wavelength. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(8), 1502–1516 (2010)
22. Pundlik, S.J., Woodard, D.L., Birchfield, S.T.: Non-ideal iris segmentation using graph cuts. In: *IEEE Computer Society Workshop on Biometrics*, pp. 1–6. IEEE, Piscataway (2008)
23. Roy, K., Bhattacharya, P.: Variational level set method and game theory applied for nonideal iris recognition. In: *Proceedings of the International Conference on Image Processing (ICIP)*, pp. 2721–2724. IEEE, Piscataway (2009)
24. Roy, K., Bhattacharya, P., Suen, C.Y.: Iris segmentation using variational level set method. *Opt. Lasers Eng.* **49**(4), 578–588 (2011)
25. Schuckers, S., et al.: On techniques for angle compensation in nonideal iris recognition. *IEEE Trans. Syst. Man Cybern. Part B Cybern.* **37**(5), 1176–1190 (2007)
26. Wildes, R.: Iris recognition: an emerging biometric technology. *Proc. IEEE* **85**(9), 1348–1363 (1997)
27. Wildes, R., et al.: A system for automated iris recognition. In: *Proceedings of the Second IEEE Workshop on Applications of Computer Vision*, pp. 121–128. IEEE, Los Alamitos (1994)
28. Zhang, D., Monro, D., Rakshit, S.: Eyelash removal method for human Iris recognition. In: *ICIP is the International Conference on Image Processing 06*, pp. 285–288. IEEE, Piscataway (2006)
29. Zuo, J., Schmid, N.: On a methodology for robust segmentation of nonideal iris images. *IEEE Trans. Syst. Man Cybern. Part B Cybern.* **40**(3), 703–718 (2010)
30. Zuo, J., Kalka, N., Schmid, N.: A robust iris segmentation procedure for unconstrained subject presentation. In: *2006 Biometrics Symposium: Special Session on Research at the Biometric Consortium Conference (BCC)*, pp. 1–6. IEEE, Piscataway (2006)

Chapter 14

Iris Segmentation for Challenging Periocular Images

Raghavender Jillela, Arun A. Ross, Vishnu Naresh Boddeti, B.V.K. Vijaya Kumar, Xiaofei Hu, Robert Plemmons, and Paúl Pauca

Abstract This chapter discusses the performance of five different iris segmentation algorithms on challenging periocular images. The goal is to convey some of the difficulties in localizing the iris structure in images of the eye characterized by variations in illumination, eyelid and eyelash occlusion, defocus blur, motion blur, and low resolution. The five algorithms considered in this regard are based on the (a) integrodifferential operator, (b) Hough transform, (c) geodesic active contours, (d) active contours without edges, and (e) directional ray detection method. Experiments on the Face and Ocular Challenge Series (FOCS) database highlight the pros and cons of the individual segmentation algorithms.

14.1 Introduction

The iris is a moving object with a small surface area that is located within the independently movable eyeball. The eyeball itself is located within another moving object – the head. Therefore, reliably localizing the iris in eye images obtained at a distance from unconstrained human subjects can be difficult. Furthermore, since the iris is typically imaged in the near-infrared portion (700–900nm) of the electromagnetic (EM) spectrum, appropriate invisible lighting is required to illuminate it prior to image acquisition. This poses a problem when the illuminator is at a considerable distance from the subject.

R. Jillela (✉) • A.A. Ross
West Virginia University, Morgantown, WV, USA
e-mail: Raghavener.Jillela@mail.wvu.edu

V.N. Boddeti • B.V.K. Vijaya Kumar
Carnegie Mellon University, Pittsburgh, PA, USA

X. Hu • R. Plemmons • P. Pauca
Wake Forest University, Winston-Salem, NC, USA



Fig. 14.1 Image acquisition setup used by the Iris On the Move (IOM) system (Image source: Matey et al. [8]. ©IEEE)

In order to mitigate this concern and to extend the depth of field of ocular-based biometric systems, the use of a small region around the eye has been proposed as an additional biometric cue. This use of this region – referred to as the *perioocular* region – has several benefits:

1. In images where the iris cannot be reliably obtained (or used), the surrounding skin region may be used to either confirm or refute an identity.
2. The use of the perioocular region represents a good trade-off between using the entire face region and using only the iris for recognition. When the entire face is imaged from a distance, the iris information is typically of low resolution; this means the matching performance due to the iris modality will be poor. On the other hand, when the iris is imaged at close quarters, the entire face may not be available, thereby forcing the recognition system to rely only on the iris.
3. The perioocular region can offer information about eye shape that may be useful as a soft biometric.
4. The depth of field of iris systems can be increased if the surrounding ocular region was to be included as well.

The use of the perioocular region is especially significant in the context of IOM (Iris on the Move) systems where the eye of a moving subject can be imaged when the individual passes through a choke point (e.g., a portal). Figure 14.1 shows the image acquisition setup used by an Iris on the Move (IOM) system.

Fig. 14.2 A sample periocular image



A sample of periocular image is shown in Fig. 14.2. Periocular image acquisition depends on the following two factors:

1. Sensor parameters (e.g., field of view, zoom factor, resolution, and view angle),
2. Standoff distance (distance between the sensor and the subject)

To perform iris recognition using periocular images, the iris region has to be segmented successfully. However, performing iris segmentation in periocular images can be very challenging. This is due to the fact that images acquired from moving subjects in a relatively unconstrained environment can be of poor quality. Such images often exhibit nonideal attributes such as off-angle iris, occlusions, blur, and poor illumination. Some challenging periocular images are shown in Fig. 14.3.

There are several benefits in determining the spatial extent of the iris in nonideal periocular images:

1. **Defining the area of the periocular region:** In images that contain poor-quality iris, periocular information can be used to improve the recognition performance [16]. However, it is very important to define a rectangular region of fixed size, from which the periocular features can be extracted. The width and height of this region are often expressed as multiples of the iris radius [10], which in turn can be determined by iris segmentation.
2. **Selective quality enhancement:** If the iris region is of poor quality in a given periocular image, image enhancement techniques can be applied exclusively within the iris region. This operation can lead to improved iris recognition performance. The appropriate region for applying selective enhancement can be determined by first segmenting the iris.
3. **Periocular image alignment:** In some cases, it is possible to encounter rotated periocular images, caused by head tilt. The center of the iris, determined by iris segmentation, can be used as an anchor point to perform rotation and to register the images appropriately.

Toward this end, this chapter discusses various techniques that can be used to perform iris segmentation in challenging periocular images.




Fig. 14.3 Periocular images exhibiting some of the nonideal attributes referred to in the narrative. (a) Poor illumination. (b) Blur. (c) Occlusion. (d) Off-angle iris

14.1.1 Periocular Region

The word periocular is a combination of *peri* (meaning the vicinity) and *ocular* (meaning related to the eye). In general, the periocular region refers to the skin and the anatomical features (e.g., eyebrows and birthmarks) contained within a specified region surrounding the eye. Periocular recognition describes the process of using the discriminative information contained in the periocular region to perform human recognition.

The region defining the periocular biometric has not been precisely defined in the biometric literature. However, Park et al. [11] suggest that a rectangular region centered at the iris center and containing the eyebrows will be the most beneficial size for recognition. Active research is being carried out to study the performance of periocular recognition under various conditions (e.g., distance [1] and, types of features that can be used [9]). It has to be noted that the periocular region can be considered to be a soft biometric trait. This chapter focuses exclusively on performing iris segmentation in periocular images and does not deal with periocular recognition.

Fig. 14.4 A sample of periocular image with the pupillary (*inner*) and limbus (*outer*) boundaries highlighted



14.1.2 Iris Segmentation

Iris segmentation refers to the process of automatically detecting the pupillary (inner) and limbus (outer) boundaries of an iris in a given image. This process helps in extracting features from the discriminative texture of the iris while excluding the surrounding regions. A periocular image showing the pupillary and limbus boundaries is seen in Fig. 14.4. Iris segmentation plays a key role in the performance of an iris recognition system. This is because improper segmentation can lead to incorrect feature extraction from less discriminative regions (e.g., sclera, eyelids, eyelashes, and pupil), thereby reducing the recognition performance.

A significant number of iris segmentation techniques have been proposed in the literature. Two most popular techniques are based on using an integrodifferential operator [3] and the Hough transform [15], respectively. The performance of an iris segmentation technique is greatly dependent on its ability to precisely isolate the iris from the other parts of the eye. Both the above-listed techniques rely on curve-fitting approach on the edges in the image. Such an approach works well with good quality, sharply focused iris images. However, under challenging conditions (e.g., nonuniform illumination, motion blur, and off-angle), the edge information may not be reliable.

In this chapter, the following techniques are considered for performing iris segmentation in challenging periocular images:

1. Integrodifferential operator
2. Hough transform
3. Geodesic active contours
4. Active contours without edges
5. Directional ray detection

The first two curve-fitting techniques are classical approaches that are computationally inexpensive. The other three techniques present relatively newer approaches for iris segmentation and are more suited for nonideal periocular images. The above selection of techniques ensures a good combination between contour-fitting and curve evolution–based approaches for performing iris segmentation in a challenging database.

14.2 Face and Ocular Challenge Series (FOCS) Database

14.2.1 Database

The Face and Ocular Challenge Series (FOCS) database was collected primarily to study the possibility of performing iris and periocular recognition in images obtained under nonideal conditions. Periocular images of resolution 750×600 pixels were captured from subjects walking through a portal in an unconstrained environment. The image acquisition system contained a set of near-infrared (NIR) sensors and illuminators. The amount of illumination observed in an image varied drastically across images due to the variation in the standoff distance, which in turn was caused by subject movement. Figure 14.5 shows some images exhibiting this effect. Some of the other challenges observed in the images include:

1. Out-of-focus blur
2. Specular reflections
3. Partially or completely occluded iris
4. Off-angled iris
5. Small size of the iris region, compared to the size of the image
6. Smudged iris boundaries
7. Sensor noise

In some cases, the size of the periocular images was smaller than 750×600 pixels. Such images were padded with pixels of zero intensity (by the distributors of the database) in order to maintain a constant image resolution (Fig. 14.6). All these factors render FOCS to be a challenging database.




Fig. 14.5 A set of images showing the significant variations in illumination caused by varying standoff distance




Fig. 14.6 Sample images in FOCS database showing the padding with zero pixels (by the distributors of the database) to maintain a fixed resolution of 750×600 pixels

14.2.2 Preprocessing

As the image quality of the FOCS database was very poor, two image preprocessing schemes were used to improve the iris segmentation performance: (i) illumination normalization, and (ii) eye center detection.

14.2.2.1 Illumination Normalization

For a majority of the images in the FOCS database, it is very difficult, even for a human expert, to determine the exact location of the iris boundaries. This is caused by low or almost no illumination in the images. In such cases, the image contrast is very low, and the iris boundaries are obscured. To alleviate this problem, illumination normalization was performed prior to iris segmentation. This was performed by adjusting the histogram of the image using the *imadjust* command in MATLAB. This step helps in increasing the contrast of the image and highlights the intensity variation across the iris boundaries. Figure 14.7 shows sample images before and after illumination normalization.

14.2.2.2 Eye Center Detection

Many iris segmentation techniques determine the rough location of the pixels lying within the pupil region by a simple image thresholding process. This is based on the fact that the pixels corresponding to the pupil area, in a uniformly illuminated image, are usually of the lowest intensity. In nonideal images, however, such an approach may not work due to the presence of nonuniform illumination. In this work, an eye center detector [5, 6] was used to output the 2D location of the center of the iris in a given image. The eye center detector is based on the shift-invariance property of the correlation filters.

The correlation filter for the eye center detector was trained on 1,000 images, in which the eye centers were manually labeled. When the correlation filter is applied to a given periocular image, a peak is observed, whose location corresponds to the




Fig. 14.7 (a) and (b): Images before illumination normalization. (c) and (d): Images after illumination normalization. Notice that the iris boundaries are better distinguishable only after performing illumination normalization




Fig. 14.8 Results of the eye center detector on sample images in FOCS data set (shown by a *white dot*). (a) Correctly detected eye center. (b) and (c) Incorrect output

center of the eye. The output of the detector was observed to be reliable in a majority of images and can be used as an anchor point (a) to perform geometric normalization of the images and (b) to initialize contours in the curve evolution–based techniques.

The eye center detector can be of significant use in nonideal images containing off-centered or occluded eyes. Figure 14.8a shows a sample of periocular image,

with the eye center detector output marked by a white dot. Figure 14.8b, c shows examples of some rare cases where the eye center was not accurately determined.

14.3 Integrodifferential Operator

Given a preprocessed image, $I(x, y)$, Daugman's integrodifferential operator [3] can be used to first determine the limbic boundary of the iris. This operation can be mathematically expressed as:

$$\max(r, x_0, y_0) \left| G_\sigma(r) * \frac{\partial}{\partial r} \oint_{r, x_0, y_0} \frac{I(x, y)}{2\pi r} ds \right|, \quad (14.1)$$

where

$$G_\sigma(r) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(r-r_0)^2}{2\sigma^2}\right) \quad (14.2)$$

represents a radial Gaussian with a center r_0 , standard deviation σ , and the symbol $*$ denotes the convolution operation. Convoluting the image with a Gaussian operator helps in smoothing the image, thereby highlighting the edge information. The term r denotes the radius of the circular arc ds , centered at (x_0, y_0) . To normalize the circular integral with respect to its perimeter, it is divided by a factor of $2\pi r$. In short, Daugman's integrodifferential operator performs circular edge detection, which can be controlled by the parameter set $\{x_0, y_0, r\}$. The computational expense associated with an iris boundary search process can be minimized by providing a range of estimates for the parameter r that are close to the actual boundary radius.

Once the limbus boundary is detected, the search process for the pupillary boundary is carried out only within the predetermined limbus region. Daugman suggests that the radius of the pupillary boundary can range from 0.1 to 0.8 of the limbus boundary radius. Figure 14.9 shows some images in which the iris boundaries are successfully segmented using an integrodifferential operator. Some examples of poorly segmented boundaries using this technique are shown in Fig. 14.10.

14.4 Hough Transform

Given a preprocessed image, $I(x, y)$, a Canny edge detector is first used to determine the edges contained in the image. Consider the set of edge points obtained by the Canny edge detector to be (x_i, y_i) , where $i = 1, 2, \dots, n$. Since this set of edge points could represent a noncontinuous or noncircular contour, a voting procedure is used to fit a circle to the boundary. For this purpose, Hough transform [4], a standard contour-fitting algorithm, is used. The voting procedure in the Hough transform technique is carried out in a parameter space, from which object candidates (in this




Fig. 14.9 Images showing successful iris segmentation output obtained using the integrodifferential operator technique




Fig. 14.10 Images showing unsuccessful iris segmentation output obtained using the integrodifferential operator technique

case, circular contours) are obtained as local maxima in an accumulator space constructed by the algorithm. In the field of iris recognition, Wildes et al. [15] demonstrated the use of Hough transform to determine the iris boundaries.

For a given set of edge points, (x_i, y_i) , $i = 1, 2, \dots, n$, Hough transform can be used to fit a circle with center (x_c, y_c) , and radius r , as follows:

$$H(x_c, y_c, r) = \sum_{i=1}^n h(x_i, y_i, x_c, y_c, r) \quad (14.3)$$




Fig. 14.11 Images showing successful iris segmentation output obtained using the Hough transform technique

where

$$h(x_i, y_i, x_c, y_c, r) = \begin{cases} 1, & \text{if } g(x_i, y_i, x_c, y_c, r) = 0 \\ 0, & \text{otherwise} \end{cases} \quad (14.4)$$

and

$$g(x_i, y_i, x_c, y_c, r) = (x_i - x_c)^2 + (y_i - y_c)^2 - r^2. \quad (14.5)$$

For each edge point contained in the set (x_i, y_i) , $g(x_i, y_i, x_c, y_c, r)$ is considered to be 0, if the parameter triplet (x_c, y_c, r) represents a circle through that point. $H(x_c, y_c, r)$ is an accumulator array, and its values (indexed by discretized values for x_c , y_c , and r) are incremented as per the equations above. The parameter triplet that corresponds to the largest value in the array is considered to be the most suitable parameter set for the circle that fits the given contour. Equation 14.5 can be modified to accommodate various contours such as circle, parabola, or ellipse. However, the computational complexity associated with a parabola or an ellipse is much higher than for a circle.

Similar to the integrodifferential operator, Hough transform–based segmentation first detects the limbus boundary of the iris. To detect the pupillary boundary, the region within the limbus boundary is used for localization purposes. Figure 14.11 shows some sample images in which the segmentation was successful using Hough transform. On the other hand, unsuccessful segmentation outputs are shown in Fig. 14.12.




Fig. 14.12 Images showing unsuccessful iris segmentation output obtained using the Hough transform technique

14.5 Geodesic Active Contours (GAC)

This approach is based on the relation between active contours and the computation of geodesics (minimal length curves). The strategy is to evolve an arbitrarily initialized curve from within the iris under the influence of geometric properties of the iris boundary. GACs combine the energy minimization approach of the classical “snakes” and the geometric active contours based on curve evolution [13].

Let $\gamma(t)$ be the curve that has to gravitate toward the outer boundary of the iris at a particular time t . The time t corresponds to the iteration number. Let ψ be a function measuring the signed distance from the curve $\gamma(t)$. That is, $\psi(x, y) =$ distance of point (x, y) to the curve $\gamma(t)$.

$$\psi(x, y) = \begin{cases} 0 & \text{if } (x, y) \text{ is on the curve;} \\ < 0 & \text{if } (x, y) \text{ is inside the curve;} \\ > 0 & \text{if } (x, y) \text{ is outside the curve.} \end{cases} \quad (14.6)$$

Here, ψ is of the same dimension as that of the eye image $I(x, y)$. The curve $\gamma(t)$ is called the level set of the function ψ . Level sets are the set of all points in ψ where ψ is some constant. Thus, $\psi = 0$ is the zeroth level set, $\psi = 1$ is the first level set, and so on. ψ is the implicit representation of the curve $\gamma(t)$ and is called the embedding function since it embeds the evolution of $\gamma(t)$. The embedding function evolves under the influence of image gradients and the region’s characteristics so that the curve $\gamma(t)$ approaches the desired boundary of the iris. The initial curve $\gamma(t)$

is assumed to be a circle of radius r just beyond the pupillary boundary. Let the curve $\gamma(t)$ be the zeroth level set of the embedding function. This implies that

$$\frac{d\psi}{dt} = 0.$$

By the chain rule,

$$\frac{d\psi}{dt} = \frac{\partial\psi}{\partial x} \frac{dx}{dt} + \frac{\partial\psi}{\partial y} \frac{dy}{dt} + \frac{\partial\psi}{\partial t},$$

that is,

$$\frac{\partial\psi}{\partial t} = -\nabla\psi \cdot \gamma'(t).$$

Splitting the $\gamma'(t)$ in the normal ($N(t)$) and tangential ($T(t)$) directions,

$$\frac{\partial\psi}{\partial t} = -\nabla\psi \cdot (v_N N(t) + v_T T(t)).$$

Now, since $\nabla\psi$ is perpendicular to the tangent to $\gamma(t)$,

$$\frac{\partial\psi}{\partial t} = -\nabla\psi \cdot (v_N N(t)). \quad (14.7)$$

The normal component is given by

$$N = \frac{\nabla\psi}{\|\nabla\psi\|}.$$

Substituting this in Eq. (13.9),

$$\frac{\partial\psi}{\partial t} = -v_N \|\nabla\psi\|.$$

Let v_N be a function of the curvature of the curve κ , stopping function K (to stop the evolution of the curve) and the inflation force c (to evolve the curve in the outward direction) such that

$$\frac{\partial\psi}{\partial t} = -(\text{div}(K \frac{\nabla\psi}{\|\nabla\psi\|}) + cK) \|\nabla\psi\|.$$

Thus, the evolution equation for ψ_t ¹ such that $\gamma(t)$ remains the zeroth level set is given by

$$\psi_t = -K(c + \varepsilon\kappa) \|\nabla\psi\| + \nabla\psi \cdot \nabla K, \quad (14.8)$$

¹The subscript t denotes the iteration number.

where K , the stopping term for the evolution, is an image-dependant force and is used to decelerate the evolution near the boundary; c is the velocity of the evolution; ε indicates the degree of smoothness of the level sets; and κ is the curvature of the level sets computed as

$$\kappa = -\frac{\psi_{xx}\psi_y^2 - 2\psi_x\psi_y\psi_{xy} + \psi_{yy}\psi_x^2}{(\psi_x^2 + \psi_y^2)^{\frac{3}{2}}}.$$

Here, ψ_x is the gradient of the image in the x direction; ψ_y is the gradient in the y direction; ψ_{xx} is the second-order gradient in the x direction; ψ_{yy} is the second-order gradient in the y direction; and ψ_{xy} is the second-order gradient, first, in the x direction and then in the y direction. Equation (13.10) is the level set representation of the geodesic active contour model. This means that the level set C of ψ is evolving according to

$$C_t = K(c + \varepsilon\kappa)\mathbf{N} - (\nabla K \cdot \mathbf{N})\mathbf{N}, \quad (14.9)$$

where N is the normal to the curve. The term $\kappa\mathbf{N}$ provides the smoothing constraints on the level sets by reducing their total curvature. The term $c\mathbf{N}$ acts like a balloon force, and it pushes the curve outward toward the object boundary. The goal of the stopping function is to slow down the evolution when it reaches the boundary. However, the evolution of the curve will terminate only when $K = 0$, that is, near an ideal edge. In most images, the gradient values will be different along the edge, thus requiring the use of different K values. In order to circumvent this issue, the third geodesic term $(\nabla K \cdot \mathbf{N})$ is necessary so that the curve is attracted toward the boundary (∇K points toward the middle of the boundary). This term makes it possible to terminate the evolution process even if (a) the stopping function has different values along the edges and (b) gaps are present in the stopping function.

The stopping term used for the evolution of level sets is given by

$$K(x, y) = \frac{1}{1 + \left(\frac{\|\nabla(G(x, y) * I(x, y))\|}{k}\right)^\alpha}, \quad (14.10)$$

where $I(x, y)$ is the image to be segmented, $G(x, y)$ is a Gaussian filter, and k and α are constants. As can be seen, $K(x, y)$ is not a function of t .

Consider an iris image to be segmented as shown in Fig. 14.13a. The stopping function K obtained from this image is shown in Fig. 14.13b (for $k = 2.8$ and $\alpha = 8$). Assuming that the inner iris boundary (i.e., the pupillary boundary) has already been detected, the stopping function K is modified by deleting the circular edges corresponding to the pupillary boundary, resulting in a new stopping function K' . This ensures that the evolving level set is not terminated by the edges of the pupillary boundary (Figs. 14.13c).




Fig. 14.13 Stopping function for the geodesic active contours. (a) Original iris image, (b) stopping function K , and (c) modified stopping function K'

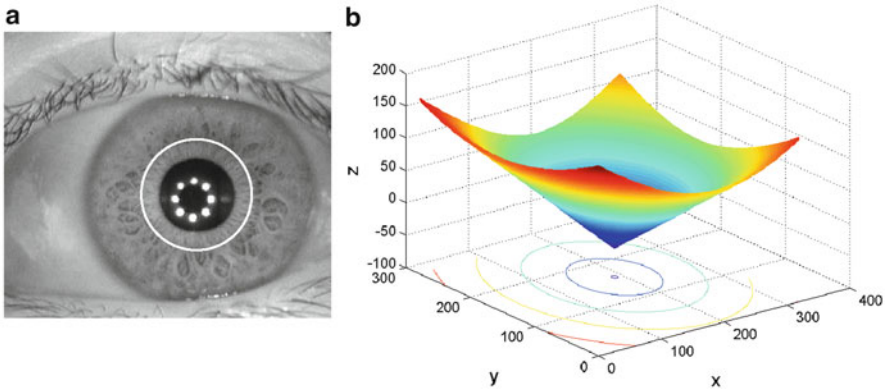


Fig. 14.14 Contour initialization for iris segmentation using GAC. (a) Zeroth level set (initial contour) and (b) mesh plot denoting the signed distance function ψ

A contour is first initialized near the pupil (Fig. 14.14a). The embedding function ψ is initialized as a signed distance function to $\gamma(t=0)$ which looks like a cone (Fig. 14.14b). Discretizing equation (14.10) leads to the following equation:

$$\frac{\psi_{i,j}^{t+1} - \psi_{i,j}^t}{\Delta t} = -cK'_{i,j}\|\nabla\psi^t\| - K'_{i,j}(\varepsilon\kappa'_{i,j}\|\nabla\psi^t\|) + \nabla\psi_{i,j}^t \cdot \nabla K'_{i,j}, \quad (14.11)$$

where Δt is the time step (e.g., Δt can be set to 0.05). The first term ($cK'_{i,j}\|\nabla\psi^t\|$) on the right-hand side of the above equation is the velocity term (advection term) and, in the case of iris segmentation, acts as an inflation force. This term can lead to singularities and hence is discretized using upwind finite differences. The upwind scheme for approximating $\|\nabla\psi\|$ is given by

$$\begin{aligned} \|\nabla\psi\| &= \sqrt{A}, \\ A &= \min(D_x^- \psi_{i,j}, 0)^2 + \max(D_x^+ \psi_{i,j}, 0)^2 + \\ &\quad \min(D_y^- \psi_{i,j}, 0)^2 + \min(D_y^+ \psi_{i,j}, 0)^2. \end{aligned}$$

$D_x^- \psi$ is the first-order backward difference of ψ in the x direction; $D_x^+ \psi$ is the first-order forward difference of ψ in the x direction; $D_y^- \psi$ is the first-order backward difference of ψ in the y direction; and $D_y^+ \psi$ is the first-order forward difference of ψ in the y direction. The second term ($K'_{i,j}(\varepsilon\kappa'_{i,j}\|\nabla\psi^t\|)$) is a curvature-based smoothing term and can be discretized using central differences. In our implementation, $c = 0.65$ and $\varepsilon = 1$ for all iris images. The third geodesic term ($\nabla\psi_{i,j}^t \cdot \nabla K'_{i,j}$) is also discretized using the central differences.

After evolving the embedding function ψ according to Eq. (13.13), the curve begins to grow until it satisfies the stopping criterion defined by the stopping function K' . But at times, the contour continues to evolve in a local region of the image where the stopping criterion is not strong. This leads to over-evolution of the contour. This can be avoided by minimizing the thin plate spline energy of the contour. By computing the difference in energy between two successive contours, the evolution scheme can be regulated. If the difference between the contours is less than a threshold (indicating that the contour evolution has stopped at most places), then the contour evolution process is terminated. The evolution of the curve and the corresponding embedding functions are illustrated in Fig. 14.15.

Since the radial fibers may be thick in certain portions of the iris, or the crypts present in the ciliary region may be unusually dark, this can lead to prominent edges in the stopping function. If the segmentation technique is based on parametric curves, then the evolution of the curve might terminate at these local minima. However, geodesic active contours are able to split at such local minima and merge again. Thus, they are able to effectively deal with the problems of local minima, thereby ensuring that the final contour corresponds to the true limbus boundary (Fig. 14.16).

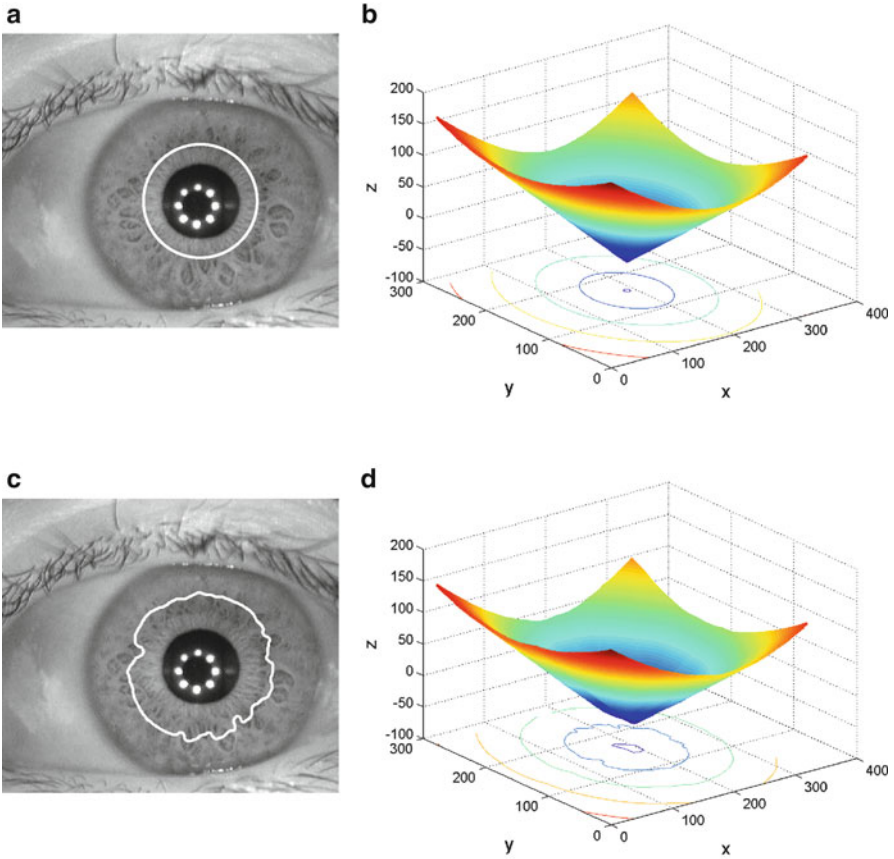


Fig. 14.15 Evolution of the geodesic active contour during iris segmentation. (a) Iris image with initial contour, (b) embedding function ψ (X and Y axes correspond to the spatial extent of the eye image, and the Z axis represents different level sets), (c–f) contours after 600 and 1,400 iterations, and their corresponding embedding functions, and (g, h) final contour after 1,800 iterations (contours shown in *white*)

Figures 14.17 and 14.18 show sample images corresponding to successful and failed segmentation outputs, respectively, obtained using the geodesic active contours technique.

14.6 Active Contours Without Edges

As explained in Sect. 14.1.2, the classical iris segmentation algorithms depend on edge information to perform boundary detection. However, in poor-quality images of FOCS database, the sharp edges required for iris segmentation are smudged to the

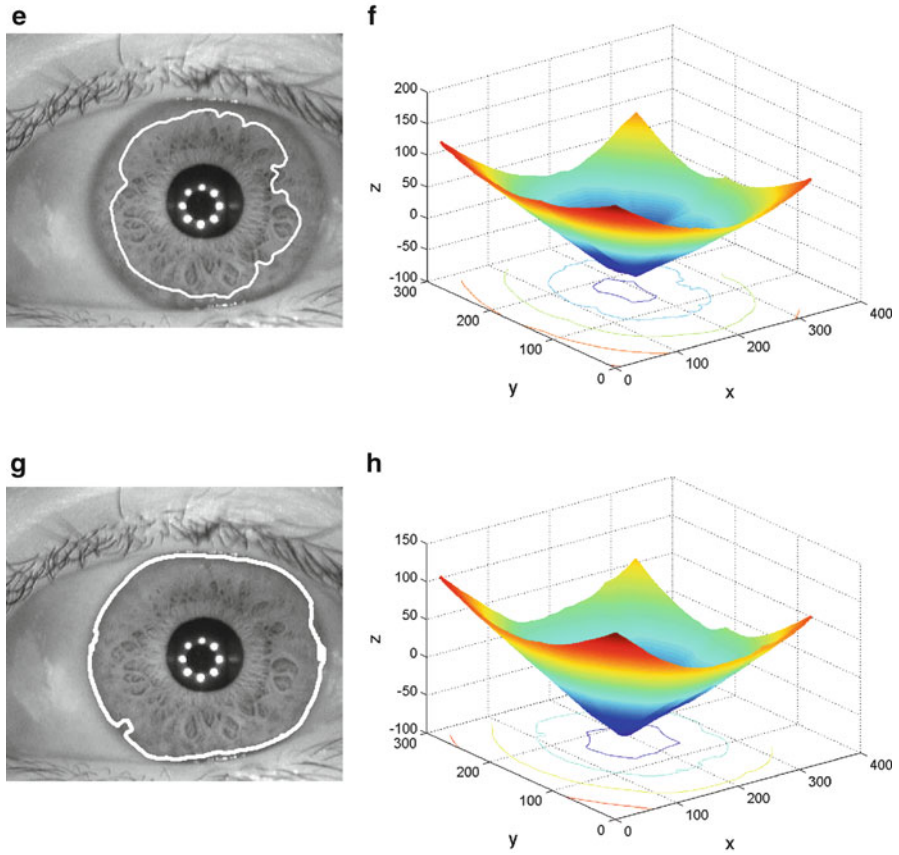


Fig. 14.15 (continued)




Fig. 14.16 The final contour obtained when segmenting the iris using the GAC scheme. (a) Example of a geodesic contour splitting at various local minima and (b) final contour (contours shown in white)




Fig. 14.17 Sample images showing successful segmentation output obtained using the geodesic active contours technique




Fig. 14.18 Sample images showing failed segmentation output obtained using the geodesic active contours technique

point that there is no discernible edge information. This problem can be alleviated to an extent by basing the segmentation algorithm on the statistics of different regions of the eye (e.g., pupil and iris), instead of the sparse edge information. To this end, a region-based active contour segmentation algorithm is developed, which is inspired by the seminal work of Chan and Vese [2]. In this technique, the contour evolution is governed by the image statistics of the *foreground* and *background* regions of a considered contour. Such an approach has been shown to work very well in cases

where both the *foreground* (the region that has to be segmented) and the *background* are homogeneous, but is known to fail when the regions are not homogeneous.

For iris images, the foreground inside the contour (pupil or iris) is homogeneous, while the background is not, due to the presence of skin, eyelashes, and eyelids. Recently, Sundaramoorthi et al. [14] addressed this issue by defining a *lookout region*. This lookout region is typically a region just outside the region of interest, from which the background statistics are computed as a function of the foreground. The basic idea behind this is that for a Gaussian distributed data (assumption for the foreground), the region outside the foreground (i.e., required to detect a transition from the foreground to background) is dependent on the image statistics of the foreground region. More precisely, the lookout region for the quickest change detection for Gaussian distributed data is given by $\Delta_{\sigma(I|\Omega)}\Omega \setminus \Omega$, where Δ_{σ} denotes dilation by $\sigma(I|\Omega)$.

14.6.1 Description of the Technique/Contour Formulation

The proposed technique segments the image based on the distribution of pixel intensities or features extracted from the eye image from a region both inside and outside the contour rather than looking for sharp edges, making it more robust to blurring and illumination variations than an edge-based active contour. The segmentation involves two steps: pupil segmentation, followed by iris segmentation. The pupil segmentation algorithm is posed as an energy minimization problem, with the objective to be minimized defined as follows:

$$E(\Omega, \mu, \bar{\mu}, \lambda_1, \lambda_2) = \int_{\Delta_{\sigma(I|\Omega)}\Omega \setminus \Omega} |I(x) - \bar{\mu}|^2 dx + \lambda_1 \int_{\Omega} |I(x) - \mu|^2 dx + \lambda_2 \Gamma(\Omega), \quad (14.12)$$

where $I(x)$ is the eye image. For simplicity reason, $I(x)$ is used instead of $I(x, y)$. Ω is the current contour, $\sigma(I|\Omega) = \frac{\int_{\Omega} |I(x) - \mu|^2 dx}{\int_{\Omega} dx}$ computes the statistics of the region within a contour which is then used to define the lookout region $\Delta_{\sigma(I|\Omega)}\Omega \setminus \Omega$ outside the current contour, $\Gamma(\Omega)$ is the regularization term, μ is the mean pixel intensity within the contour, $\bar{\mu}$ is the mean pixel intensity in the lookout region, and λ_1 and λ_2 are scalars weighting the different criteria defining the contour energy. The output of the eye center detection is used to initialize a contour for pupil segmentation. Once the pupil is segmented, a contour is initialized just outside the pupil. However, Eq. 14.12 may not be used directly because the pupil region needs to be excluded from within the current contour. This leads to the following energy formulation for detecting the outer boundary:

$$E(\Omega, \mu, \bar{\mu}, \lambda_1, \lambda_2) = \int_{\Delta_{\sigma(I|\bar{\Omega})}\Omega \setminus \Omega} |I(x) - \bar{\mu}|^2 dx + \lambda_1 \int_{\bar{\Omega}} |I(x) - \mu|^2 dx + \lambda_2 \Gamma(\Omega). \quad (14.13)$$




Fig. 14.19 Some examples showing successful segmentation of iris regions under challenging conditions (poor illumination and blur) using the active contours without edges technique on the FOCS database

where $\bar{\Omega} = \Omega \setminus \Omega_{\text{pupil}}$ defines the region within the current contour excluding the pupil. Due to occlusions like specular reflections and eye lashes as part of the contour evolution, many small contours are also formed along with the pupil or iris contours. These extraneous contours are pruned based on their regional statistics like size of region and eccentricity of region. Further, the final contour is smoothed by applying a simple moving average filter to the contour. Figures 14.19 and 14.20 show some examples of successful and unsuccessful iris segmentation outputs, respectively.

14.7 Directional Ray Detection Method

This section presents an iris segmentation technique for low quality and off-angle iris images that are based on a novel directional ray detection segmentation scheme. This method can employ calculus of variations or directional gradients to better approximate the boundaries of the pupillary and limbic boundaries of the iris. Variational segmentation methods are known to be robust in the presence of noise and can be combined with shape-fitting schemes when some information about the object shape is known a priori. Quite commonly, circle fitting is used to approximate the boundaries of the iris, but this assumption may not necessarily hold for noncircular boundaries or off-axis iris data. For computational purposes, this technique uses directional gradients and circle-fitting schemes, but other shapes can also be easily considered.




Fig. 14.20 Some examples showing unsuccessful segmentation of iris regions under challenging conditions (poor illumination and blur) using the active contours without edges technique on the FOCS database. Some of the mistakes are due to eye-detection failure and closed eye

This technique extends the work by Ryan et al. [12], who approach the iris segmentation problem by adapting the starburst algorithm to locate pupillary and limbic feature pixels used to fit a pair of ellipses. The starburst algorithm was introduced by Li et al. [7], for the purpose of eye tracking. The proposed method involves multiple stages of an iterative ray detection scheme, initialized at points radiating out from multiple positions around eye region, to detect the pupil, iris, and eyelid boundaries. The scheme also involves the use of a correlation filter method for eye center detection.

14.7.1 Description of the Technique

The proposed segmentation approach includes the following four sequential steps:

1. The location of the eye center is first estimated using a specially designed correlation filter.
2. The circular boundary of the pupil is obtained by uniform key point extraction and Hough transformation.
3. The iris boundary is determined by the directional ray detection scheme.
4. Finally, the eyelid boundary is also determined by directional ray detection scheme, applied at multiple points.




Fig. 14.21 One iteration of iterative directional ray segmentation for pupil detection. (a) $I(x,y)$. (b) $I_p(x,y)$. (c) Directional rays. (d) Segmented pupil

14.7.1.1 Eye Center Detection

Given a periocular image, the location of the eye center is obtained using the correlation-based eye center detector described in Sect. 14.2.2.2. It is noticed that the accuracy of the current iris segmentation technique is crucially related to the correctness of the eye center detector output.

14.7.1.2 Pupil Segmentation

In this step, it is assumed that the eye center (x_c, y_c) of image $I(x,y)$ is located within the pupil. Iterative directional ray detection is then applied within a square region $I_p(x,y)$ of size of $2(r_p + R_p)$ centered at (x_c, y_c) , where $0 < r_p < R_p$ are pupil radius bounds determined experimentally from the data set. A sample square region $I_p(x,y)$ containing a pupil is shown in Fig. 14.21b.

Each iteration includes ray detection to extract structural key points and the Hough transformation to estimate the circular pupil boundary. As shown in Fig. 14.21c, a number of rays are chosen starting at the estimated eye center (x_c, y_c) along m directions $\theta \in \Theta$, where Θ is a subset of $[0, 2\pi]$. For the FOCS data set, the value of Θ is set as $\Theta = [\frac{\pi}{4}, \frac{3\pi}{4}] \cup [\frac{5\pi}{4}, \frac{7\pi}{4}]$, to select pixel information contained along the vertical direction. The directional gradient difference along these rays is then calculated. Locations for which the gradient difference reaches an absolute maximum are then selected as key points. The Hough transform is performed to find the circular pupil boundary, the estimated pupil center (x_p, y_p) , and the radius of the pupil \hat{r}_p as shown in Fig. 14.21d.

Gradient calculation is notoriously sensitive to image noise. To compensate for the presence of noise, an adaptive approach is implemented that:

1. Enhances the edges at the location of selected key points
2. Uses the estimated pupil center (x_p, y_p) instead of the eye center location (x_c, y_c)
3. Properly sets up the length of test rays to vary toward \hat{r}_p

In this implementation, at most three iterations are needed to obtain a very reliable and accurate pupil segmentation, as shown in Fig. 14.21d.




Fig. 14.22 Iterative directional ray detection results for iris and eyelids. (a) Rays ($\theta = \frac{\pi}{6}$). (b) Segmented iris. (c) Rays (multiple pos). (d) Eyelids detection

14.7.1.3 Iris Segmentation

In this step, the estimate pupil center (x_p, y_p) and pupil radius \hat{r}_p are used to select a square region of size $2(\hat{r}_p + R_I)$ centered at (x_p, y_p) , where $R_I > \hat{r}_p$ is a bound parameter determined experimentally from the data set. Iterative directional ray detection is then applied within this region. Specifically, a set of rays are chosen to emanate from a set of points Γ along n directions $\theta \in \Theta$, where $\Gamma = \{(x, y) | x = x_p + (1 + \alpha)\hat{r}_p, y = y_p + (1 + \alpha)\hat{r}_p\}$, $\Theta = [-\frac{\pi}{6}, \frac{\pi}{6}] \cup [\frac{5\pi}{6}, \frac{7\pi}{6}]$, and $\alpha = 0.5$. As shown in Fig. 14.22a, the directions of test rays in this stage are close to the horizontal axis to avoid the upper and lower eyelid regions. Similarly, to perform pupil segmentation, the key points are determined using the adaptive gradient different scheme, and Hough transform is applied to find the iris, as shown in Fig. 14.22b. The main advantage of this approach is a reduced sensitivity to light reflection and other impurities in the iris region as well as increased computational efficiency of the Hough transform.

14.7.1.4 Eyelid Boundary Detection

Detecting the eyelid boundaries is an important step in the accurate segmentation of the visible iris region as well as for the determination of iris quality metrics. For this, a two-step directional ray detection approach from multiple starting points is used. As shown in Fig. 14.22c, the testing rays are set to emanate from the boundaries of the pupil and iris regions. Two groups of rays are chosen roughly along the vertical and horizontal directions. Gradient differences are computed along these rays, and key points are appropriately selected. A least squares best fitting model is then applied to the binary key point maps, and a first-step fitting estimation to the eyelid boundary is produced. To improve accuracy of the estimation in the previous step, the noisy key points are removed outside the estimated eyelid boundary, and the edges are enhanced around the remaining key points. Applying ray detection from multiple positions, a new set of key points is obtained. The least squares fitting model is then applied again, resulting in an improved estimation of eyelids, shown in Fig. 14.22d.



Fig. 14.23 Directional ray detection method – some successful examples from the FOCS data set




Fig. 14.24 Directional ray detection method – some unsuccessful examples from the FOCS data set

The proposed iterative directional ray detection method is a promising segmentation approach as it has been shown to be robust in the segmentation of poor-quality data. An adaptive gradient difference scheme is used as a core procedure to select key points in specific image regions, simplifying the overall problem into simpler more manageable pieces. The computational cost for this technique is comparable to other existing techniques in the literature. Figure 14.23 provides examples of successful segmentation obtained by the directional ray detection technique. Figure 14.24 provides images of unsuccessful examples.

14.8 Experimental Evaluation

The total number of images contained in the FOCS database is 9,581. Performing an evaluation of iris segmentation algorithms using the entire database is prohibitive due to the high number of images and the computational cost aspects involved. Therefore, a subset of 404 images was chosen from the FOCS database and used to determine the best performing iris segmentation algorithm. It was ensured that the sample data set was representative of the full FOCS database, in terms of the challenges posed (e.g., nonuniform illumination, blur, and occlusion). The resolution of the images was unaltered (750×600 pixels). The number of unique subjects contained in this 404 image data sets was 108, with the number of samples per subject ranging from 1 to 11.

14.8.1 Segmentation Accuracies

The correlation filter approach for detecting the eye centers, when applied on the full FOCS data set, yielded a success rate of over 95%. The output of the eye center detector was used only for three techniques: geodesic active contours, active contours without edges, and directional ray detection techniques. The performance of an iris segmentation technique was measured by computing the *segmentation accuracy*, defined as follows:

$$\text{Segmentation accuracy} = \frac{\text{number of correctly segmented images}}{\text{number of input images provided}} \times 100 \quad (14.14)$$

The types of iris segmentation technique used and their corresponding segmentation accuracies are summarized in the table below.

Segmentation technique	Number of input images provided	Number of correctly segmented images	Segmentation accuracy (%)
Integrodifferential operator	404	207	51.2
Hough transform	404	210	52
Geodesic active contours (GAC)	404	358	88.6
Active contours without edges	404	365	90.3
Directional ray detection	404	343	84.9

From the results obtained on the 404-image data set, it was observed that the techniques based on active contours resulted in the best performance among the considered set of techniques.

14.8.2 Analysis

The preprocessing schemes appear to have a significant role in the segmentation performance for all the techniques. Illumination normalization helps in increasing the contrast of the image, thereby highlighting the iris boundaries. Similarly, eye center detector helps in localizing a region for iris boundary search process.

The considered set of iris segmentation techniques ensures a balance between the classical approaches and the relatively newer approaches to handle challenging iris images. Both the integrodifferential operator and Hough transform require relatively less computations, when compared to the other three techniques. However, their performance was observed to be low, due to the poor-quality input data. On the other hand, geodesic active contours, active contours without edges, and directional ray detection algorithm provide better performance, at the expense of higher computational complexity.

Geodesic active contours can be effectively used to evolve a contour that can fit to a noncircular iris boundary (typically caused by eyelids or eyelash occlusions). However, edge information is required to control the evolution and stopping of the contour. The performance of geodesic active contours for this database was limited to 88.6% due to the lack of edge information, which is caused by poor illumination levels.

14.9 Summary

Performing iris recognition at a distance is a challenging task. Significant amount of research is being conducted toward improving the recognition performance in iris images acquired under unconstrained conditions. Employing better image acquisition systems can significantly improve the quality of input images and, thereby, the recognition performance. However, for practical purposes, it is necessary to develop algorithms that can handle poor-quality data. In this regard, this chapter discusses the problem of iris segmentation in challenging periocular images. A set of 5 iris segmentation techniques were evaluated on a periocular database containing various nonideal factors such as occlusions, blur, and drastic illumination variations. This work helps serve two main purposes: (a) it describes a real-world problem of uncontrolled iris image acquisition and the associated challenges, and (b) it highlights the need for robust segmentation techniques to handle the poor-quality iris images.

Acknowledgments This work was sponsored under IARPA BAA 09-02 through the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-10-2-0013. The views and conclusions contained in this chapter are those of the authors and should not be interpreted as representing of official policies, either expressed or implied, of IARPA, the Army Research Laboratory, or the US government. The US government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation herein.

References

1. Bharadwaj, S., Bhatt, H., Vatsa, M., Singh, R.: Periocular biometrics: when iris recognition fails. In: Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS), BTAS'10. IEEE, Piscataway (2010)
2. Chan, T., Vese, L.: Active contours without edges. *IEEE Trans. Image Process.* **10**(2), 266–277 (2001)
3. Daugman, J.: How iris recognition works. *IEEE Trans. Circuits Syst. Video Technol.* **14**(1), 21–30 (2004)
4. Illingworth, J., Kittler, J.: A survey of the Hough transform. *Comput. Vis. Graph. Image Process.* **44**(1), 87–116 (1988)
5. Kumar, B.V.K.V., Hassebrook, L.: Performance measures for correlation filters. *Appl. Opt.* **29**(20), 2997–3006 (1990)

6. Kumar, B., et al.: Biometric verification with correlation filters. *Appl. Opt.* **43**(2), pp. 391–402 (2004)
7. Li, D., Babcock, J.S., Parkhurst, D.J.: Openeyes: a low-cost head-mounted eye-tracking solution. In: *ETRA*, San Diego, California, USA, pp. 95–100 (2006)
8. Matey, J., Naroditsky, O., Hanna, K., Kolczynski, R., LoIacono, D., Mangru, S., Tinker, M., Zappia, T., Zhao, W.: Iris On the Move: acquisition of images for iris recognition in less constrained environments. *Proc. IEEE* **94**(11), 1936–1947 (2006)
9. Miller, P., Lyle, J., Pundlik, S., Woodard, D.: Performance evaluation of local appearance based periocular recognition. In: *Fourth IEEE International Conference on Biometrics: Theory Applications and Systems (BTAS), BTAS'10*. IEEE, Piscataway (2010)
10. Park, U., Ross, A., Jain, A.K.: Periocular biometrics in the visible spectrum: a feasibility study. In: *Proceedings of the 3rd IEEE International Conference on Biometrics: Theory, applications and Systems, BTAS'09*, pp. 153–158. IEEE, Piscataway (2009)
11. Park, U., Jillela, R.R., Ross, A., Jain, A.K.: Periocular biometrics in the visible spectrum. *IEEE Trans. Inf. Forensics Secur.* **6**(1), 96–106 (2011)
12. Ryan, W., Woodard, D., Duchowski, A., Birchfield, S.: Adapting starburst for elliptical iris segmentation. In: *Proceedings of the 2nd IEEE International Conference on Biometrics: Theory, applications and Systems, BTAS'08*. IEEE, Piscataway (2008)
13. Shah, S., Ross, A.: Iris segmentation using geodesic active contours. *IEEE Trans. Inf. Forensics Secur.* **4**(4), 824–836 (2009)
14. Sundaramoorthi, G., Soatto, S., Yezzi, A.J.: Curious snakes: a minimum latency solution to the cluttered background problem in active contours. In: *CVPR*, San Francisco, pp. 2855–2862 (2010)
15. Wildes, R., Asmuth, J., Green, G., Hsu, S., Kolczynski, R., Matey, J., McBride, S.: A system for automated iris recognition. In: *Proceedings of the Second IEEE Workshop on Applications of Computer Vision*, Sarasota, pp. 121–128 (1994)
16. Woodard, D.L., Pundlik, S.J., Miller, P.E., Lyle, J.R.: Appearance-based periocular features in the context of face and non-ideal iris recognition. *Signal Image Video Process.* **5**(4), 443–455 (2011)

Chapter 15

Periocular Recognition from Low-Quality Iris Images

Josh Klontz and Mark J. Burge

Abstract Definitions of the periocular region vary, but typically encompass the skin covering the orbit of the eye. Especially in cases where the iris has not been acquired with sufficient quality to reliably compute an IrisCode, the periocular region can provide additional discriminative information for biometric identification. The NIR periocular images which form NIST's Face and Ocular Challenge Series (FOCS) are characterized by large variations in illumination, eye-lid and eye-lash occlusion, de-focus blur, motion blur and low resolution. We investigate periocular recognition on the FOCS dataset using three distinct classes of features: photometric, keypoint, and frequency-based. We examine the performance of these features alone, in combination, and when fused with classic IrisCodes.

15.1 Introduction

Section 15.2 introduces NIST's Face Ocular Challenge Series (FOCS) dataset, which is characterized by large variations in illumination, eyelid and eyelash occlusion, defocus blur, motion blur, and low resolution. Section 15.3 identifies some of the properties which make the FOCS dataset difficult for IrisCode systems. Section 15.4 investigates three alternatives to IrisCodes which take advantage of the periocular information available in the FOCS images: Gabor wavelets, local binary patterns (LBP), and scale-invariant feature transform (SIFT). Section 15.5 demonstrates how applying a sequence of principle component analysis (PCA) and linear discriminant analysis (LDA) to the alternative feature encodings can improve match performance. Finally, Sect. 15.6 explores the result of fusing our alternative encodings with traditional IrisCodes.

J. Klontz (✉) • M.J. Burge
The MITRE Corporation, McLean, VA, USA
e-mail: jklontz@mitre.org; mburge@mitre.org

15.2 Face Ocular Challenge Series (FOCS)

The FOCS dataset is an example of low-quality iris images which exhibit a high rate of failure to acquire (i.e., failure to enroll) on commercial iris systems. A NIST document (<http://www.nist.gov/itl/iad/ig/focs.cfm>) describes the rationale for the FOCS dataset as “A significant amount research has gone into iris recognition from ‘high’ quality iris images captured from cooperative subjects. The results of this research are commercially available products for matching ‘high’ quality iris images. The cutting edge of research is recognition from images that contain the eye region of the face and the iris. The iris in these images can be of variable quality” [15]. We investigate using the periocular area of the FOCS images to improve match performance.

NIST’s FOCS dataset was commissioned by IARPA to serve as a challenge set for performers in the ocular track of the IARPA BEST program. It consists of NIR frames extracted from the Iris on the Move [12] (IoM) videos (see Fig. 15.1) of the MBGC Portal Track V2. It contains 4,792 left and 4,789 right eyes of 136 subjects. Images are of a single iris (see Fig. 15.2) and the surrounding periocular region. They are encoded as 750×600 , 8-bit grayscale, JPEG-quality images.

The COIR [17] ocular collection was a precursor to the FOCS dataset which was developed “to compensate for the lack of a dedicated ocular dataset and to leverage the large amount of existing face and iris image and video data already available” [17]. In contrast to the FOCS dataset, COIR used visible light images extracted from NIST’s MBGC [19] series.

15.3 FOCS: A Difficult Dataset for IrisCodes

FOCS images pose a number of challenges (e.g., large variations in illumination, eyelid and eyelash occlusion, defocus blur, motion blur, and low resolution) which make them difficult for traditional IrisCode systems. Two traditional IrisCode algorithms were used in the study, NIST’s VASIR [9] system and a commercial



Fig. 15.1 The FOCS dataset consists of frames extracted from NIST’s Multiple Biometrics Grand Challenge (MBGC) IoM (Iris on the Move) NIR video sequences




Fig. 15.2 Example of FOCS images. The image on the *left* 05301d128_06_r was enrollable by commercial algorithms, while the image on the *right* 05301d116_09_r was not

algorithm. The iris-pupil and -scelera boundary were manually located, and the coordinates were provided to both algorithms. Without this information, both systems failed to enroll over 90% of the FOCS images.

After localizing the iris-sclera boundary (Fig. 15.3a, b), the first step in a traditional IrisCode algorithms is to segment and unwrap the iris texture. Figure 15.3d gives an example of the quality and resolution of a typical unwrapped FOCS images and can be compared to that of a typical “high-”quality iris images such as that of CASIA V1, Fig. 15.3c.

15.3.1 *IrisCode Performance*

While both VASIR [9] and the commercial algorithm were easily able to separate match and non-match scores,¹ on the higher-quality CASIA images, they failed to enroll nearly all FOCS images.

15.3.2 *Dynamic Range*

The dynamic range of an image can be defined as the difference between the image’s highest and lowest pixel values. In cases where a small number of outlier values

¹As can be expected for one of the earlier standard test sets like CASIA [18], the commercial algorithm was able to perfectly separate match and non-match pairs.

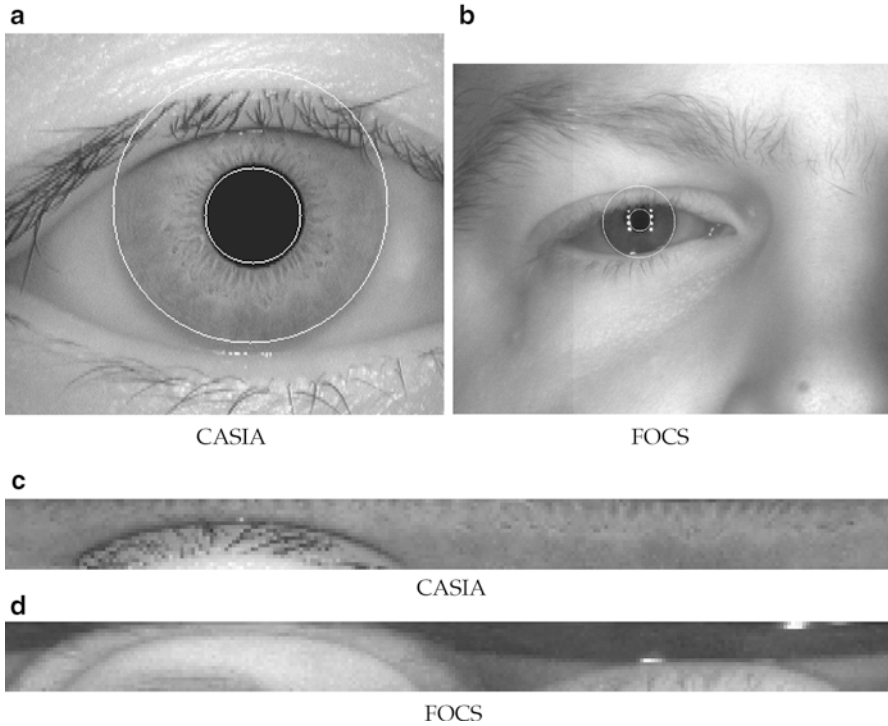


Fig. 15.3 Example of the segmentation of (a) “high-”quality NIR iris image CASIA, (b) a “lower-”quality FOCS image, (c) unwrapped CASIA, and (d) unwrapped FOCS

skew the distribution, a number of more robust variations, such as discarding 0.5% of both tails, can be used. A low dynamic range can indicate acquisition problems such as insufficient illumination on target or nonoptimal sensor bias settings.

Figure 15.4 graphs the mean \bar{x} and standard deviation σ of all the pixel values (i.e., each graph point corresponds to the mean and sigma of a single image) in each of 2,000 NIR iris images sampled from the ND-IRIS [20] and FOCS datasets. Figure 15.4 illustrates how the distribution of the dynamic range of the FOCS dataset differs from that of a typical NIR iris dataset like ND-IRIS. The low dynamic range of the FOCS images makes both segmentation and feature extraction challenging. Figure 15.5 shows the increasing dynamic range of the FOCS image, note that one-fourth of the images have a σ less than 23 and a \bar{x} of less than 76.

15.4 Alternatives to IrisCodes

Using a Bayesian graphical model to find the correlations between patches on the FOCS images, Boddeti [3] reported an equal error rate (EER) of 26.81% for the left ocular region, 23.83% for the right ocular region, and 30.8% for a customized

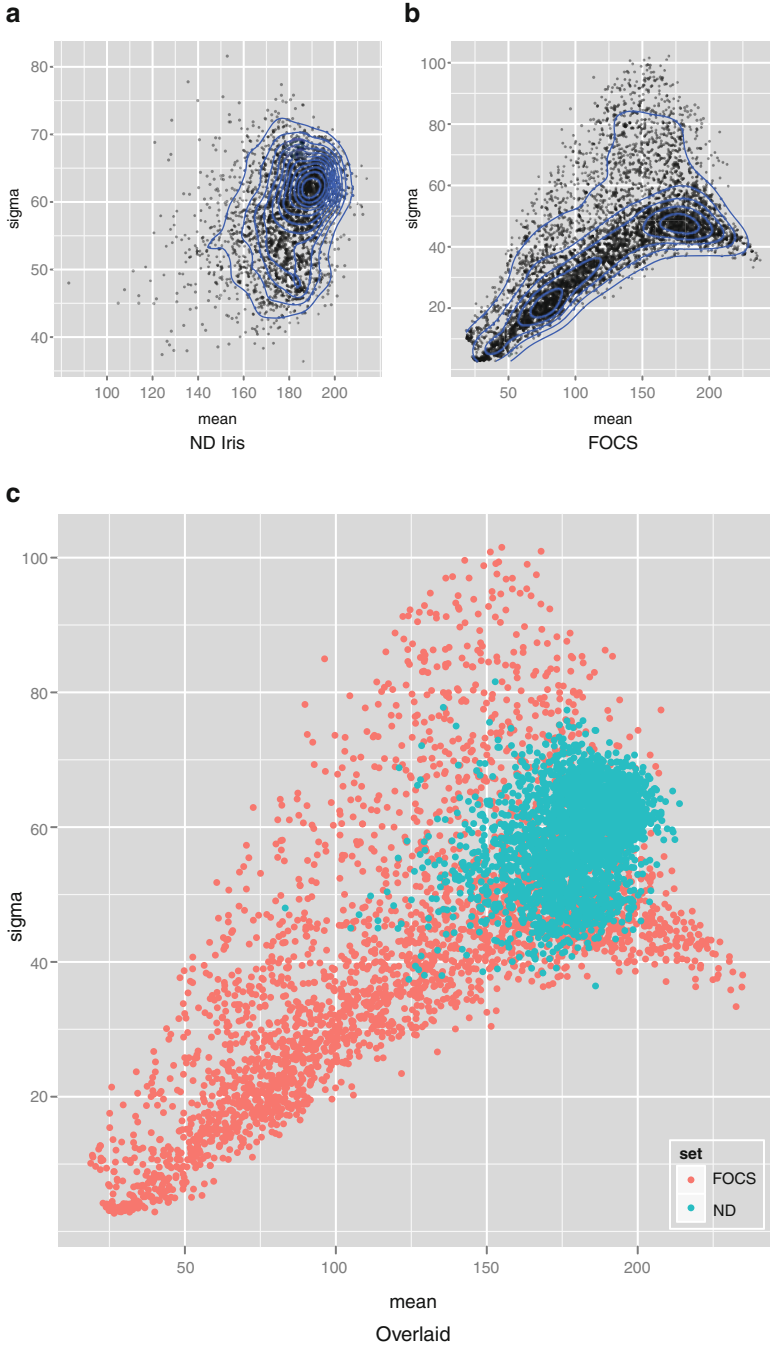


Fig. 15.4 Mean \bar{x} and standard deviation σ of all the pixel values in each of 2,000 NIR iris images from (a) ND-IRIS, (b) FOCS, and (c) overlay of ND-IRIS (teal) and FOCS (red). Each point represents the \bar{x} and σ of a single image




Fig. 15.5 FOCS images sorted (increasing) by the standard deviation of their dynamic range, equally sampled (i.e., 4/12 of the images have a σ less than the *upper-right* image)

IrisCode implementation. This established that the ocular region could be used to improve match performance over IrisCodes alone. To further assess the usefulness of periocular information for recognition, we investigated the use of the three alternatives to IrisCodes: Gabor wavelets, local binary patterns (LBP), and scale-invariant feature transform (SIFT) features for ocular recognition.

15.4.1 Keypoint Descriptors

The skin texture in periocular images contains discriminating marks arising from variations in the skin surface. In general, keypoint descriptors are those which are able to identify edge-like features which persist over scale variations. In the following experiments, the Speeded-Up Robust Feature (SURF) [24] variation of the SIFT [11] descriptor was used. Compared to SIFT, SURF uses integral images to decrease computation time and Haar wavelets to increase robustness. Feature vectors were constructed by densely sampling an 11×11 grid of 24 pixel radius SURF features.

Table 15.1 Performance of LBP on FOCS. (a) Impact of LBP window size on match performance. (b) Highest scoring LBP variation (i.e., $LBP_{48,8}^{u2}$) vs. VASIR. Note LBP's TAR of .61 vs. VASIR's TAR of .06 at the same FAR of 10^{-2}

	$LBP_{256,8}^{u2}$	$LBP_{128,8}^{u2}$	$LBP_{64,8}^{u2}$	$LBP_{48,8}^{u2}$	$LBP_{32,8}^{u2}$
FAR = 10^{-3}	0.04	0.26	0.52	0.53	0.43
FAR = 10^{-2}	0.16	0.49	0.65	0.66	0.58

15.4.2 Local Binary Patterns

A seminal method in face recognition is the use of local binary patterns [16] (LBP) to represent the face [1]. Local binary patterns are level 2 features that represent small patches across the face with histograms of binary patterns that encode the structure and texture of the face. In this experiment, we encode the periocular region using LBPs.

Local binary patterns describe each pixel using a p -bit binary number. Each bit is determined by sampling p pixel values at uniformly spaced locations along a circle of radius r , centered at the pixel being described. For each sampling location, the corresponding bit receives the value 1 if it is greater than or equal to the center pixel, and 0 otherwise.

A special case of LBP, called the uniform LBP [16], is generally used in face recognition. Uniform LBP (LBP U2) assigns any nonuniform binary number to the same value, where uniformity is defined by whether more than u transitions between the values 0 and 1 occur in the binary number. In the case of $p = 8$ and $u = 2$, the uniform LBP has 58 uniform binary numbers, and the 59th value is reserved for the remaining $256 - 58 = 198$ nonuniform binary numbers. Thus, each pixel will take on a value ranging from 1 to 59. Overlapping histograms of LBP U2 patterns were extracted from window sizes ranging from 32 to 256 pixels (see Table 15.1), and the resulting histograms were compared using a Chi-squared distance metric. Table 15.1 gives the match results from five variations.

In the context of periocular recognition, LBP values are first computed at each pixel in the (normalized) periocular image. The image is tessellated into patches with a height and width of 12 pixels. For each patch i , a histogram of the LBP values $S'_i \in \mathbb{Z}^{d_s}$ is computed (where $d_s = 59$). This feature vector is then normalized to the feature vector $S_i \in \mathbb{R}^{d_s}$ by

$$S_i = \frac{S'_i}{\sum_i^{d_s} S'_i}.$$

Finally, we concatenate the N vectors into a single vector x of dimensionality $d_s \cdot N$. Note that, in our implementation, the illumination filter proposed by Tan and Triggs [23] is used prior to computing the LBP codes in order to suppress nonuniform illumination variations.

15.4.3 Gabor Wavelets

Gabor features are one of the first level 2 facial features [8] to have been used with wide success in representing facial images [10, 22, 25]. One reason Gabor features are popular for representing both facial and natural images is their similarity with human neurological receptor fields [13, 14].

A Gabor image representation is computed by convolving a set of Gabor filters with an image (in this case, a periocular image). The Gabor filters are defined as

$$G(x, y, \theta, \eta, \gamma, f) = \frac{f^2}{\pi\gamma\eta} e^{-\left(\frac{f^2}{\gamma^2}x'^2 + \frac{f^2}{\eta^2}y'^2\right)} e^{(j2\pi f x')} \quad (15.1)$$

$$x' = x \cos \theta + y \sin \theta \quad (15.2)$$

$$y' = -x \sin \theta + y \cos \theta, \quad (15.3)$$

where f sets the filter scale (or frequency), θ is the filter orientation along the major axis, γ controls the filter sharpness along the major axis, and η controls the sharpness along the minor axis. Typically, combinations across the following values for the scale f and orientation θ are used: $f = \{0, 1, \dots, 4\}$ and $\theta = \{\pi/8, \pi/4, 3\pi/8, \dots, \pi\}$. This creates a set (or bank) of filters with different scales and orientations. Given the bank of Gabor filters, the input image is convolved with each filter, which results in a Gabor image for each filter. The combination of these scale and orientation values results in 40 different Gabor filters, which in turn results in 40 Gabor images, for example.

In this chapter, the recognition experiments using a Gabor image representation operate by (i) performing illumination correction using the method proposed by Tan and Triggs [23], (ii) computing the phase response of the Gabor images with $f = \{8, 16, 32\}$ and $\theta = 0, \pi/4, \pi/2, 3\pi/4$, (iii) tessellating the Gabor image(s) into patches of size 12×12 , (iv) quantizing the phase response (which ranges from 0 to 2π) into 24 values and computing the histogram within each patch, and (v) concatenating the histogram vectors into a single feature vector. Given two (aligned) periocular images, the distance between their corresponding Gabor feature vectors is used to measure the dissimilarity between the two periocular images.

15.5 Improving Performance Using Learning

The trainable algorithm used in this study is the Spectrally Sampled Structural Subspace Features algorithm [7], which is abbreviated as 4SF@. This algorithm uses multiple discriminative subspaces to perform recognition. After geometric normalization of a periocular image, illumination correction is performed using the illumination correction filter presented by Tan and Triggs [23]. Periocular images

are then represented using histograms of local binary patterns at densely sampled patches [1] (to this point, 4SF is the same as the non-trainable LBP algorithm described in Sect. 15.4.2). For each patch, principal component analysis (PCA) is performed so that 95% of the variance is retained.

Given a training set of subjects, multiple stages of weighted random sampling are performed, where the spectral densities (i.e., the eigenvalues) from each face patch are used for weighting. The randomly sampled subspaces are based on Ho’s original method [5]; however, the proposed approach is unique in that the sampling is weighted based on the spectral densities. For each stage of random sampling, LDA [2] is performed on the randomly sampled components. The LDA subspaces are learned using subjects randomly sampled from the training set (i.e., bagging [4]). Finally, distance-based recognition is performed by projecting the LBP representation of periocular images into the per-patch PCA subspaces and then into each of the learned LDA subspaces. The sum of the Euclidean distance in each subspace is the dissimilarity between two periocular images.

15.6 Fusing Periocular Features with IrisCodes

Ross [21] reported recognition rates on the FOCS dataset based on fusing results from Gradient Orientation Histogram (GOH), SIFT, and probabilistic deformation model (PDM). In Table 15.2, we present the results of Z-score fusing [6] combinations of LBP, SURF, Gabor wavelets, and traditional IrisCodes on the FOCS dataset (see Fig. 15.6).

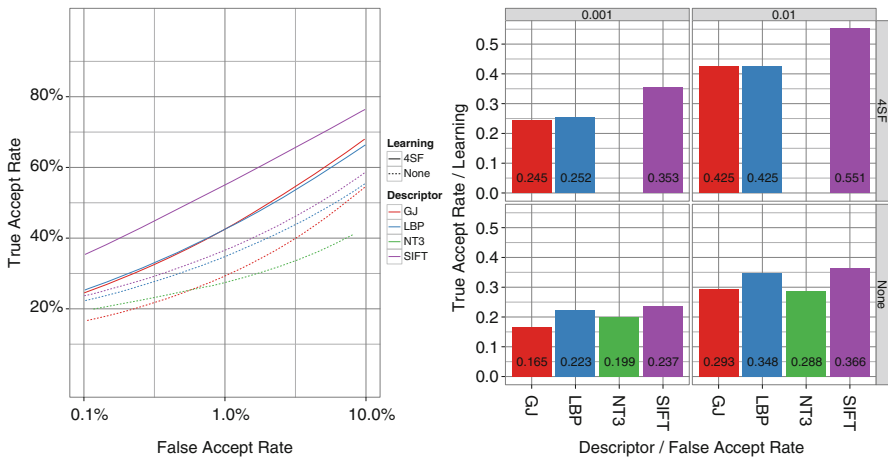


Fig. 15.6 Performance of a commercial iris code system (NT), Gabor wavelets (GJ), local binary patterns (LBP), and keypoint descriptors (SIFT), with (solid line) and without (dashed line) learning

Table 15.2 Fusion method for the following was Z-score score normalization then weighing the higher performing descriptor 2:1

	without learning				with learning			
	GJ	LBP	SIFT	NT	GJ	LBP	SIFT	NT
GJ	0.293	0.357	0.357	0.398	0.398	0.546	0.485	0.506
LBP	0.357	0.348	0.381	0.430	0.546	0.430	0.556	0.488
SIFT	0.357	0.381	0.366	0.445	0.485	0.556	0.445	0.591
NT	0.398	0.430	0.445	0.288	0.506	0.488	0.591	na

15.7 Conclusions

The NIR periocular images which form NIST's Face and Ocular Challenge Series (FOCS) are characterized by large variations in illumination, eyelid and eyelash occlusion, defocus blur, motion blur, and low resolution. We investigated periocular recognition on the FOCS dataset using three distinct classes of features: photometric, keypoint, and frequency-based. We examined the performance of these features alone, in combination, and when fused with classic IrisCodes. We demonstrated that in cases where the iris has not been acquired with sufficient quality to reliably compute an IrisCode (as in NIST's FOCS dataset), the periocular region can provide additional discriminative information for biometric identification.

References

1. Ahonen, T., Hadid, A., Pietikainen, M.: Face description with local binary patterns: application to face recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **28**(12), 2037–2041 (2006)
2. Belhumeur, P., Hespanda, J., Kriegman, D.: Eigenfaces vs. fisherfaces: recognition using class specific linear projection. *IEEE Trans. Pattern Anal. Mach. Intell.* **19**(7), 711–720 (1997)
3. Boddeti, V., Smereka, J., Kumar, B.: A comparative evaluation of iris and ocular recognition methods on challenging ocular images. *IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–8. IEEE, Piscataway, Washington, DC (2011)
4. Breiman, L.: Bagging predictors. *Mach. Learn.* **24**, 123–140 (1996)
5. Ho, T.K.: The random subspace method for constructing decision forests. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**(8), 832–844 (1998)
6. Jain, A., Nandakumar, K., Ross, A.: Score normalization in multimodal biometric systems. *Pattern Recogn.* **38**(12), 2270–2285 (2005)
7. Klare, B.: Spectrally sampled structural subspace features (4SF). In: Michigan State University Technical Report, MSU-CSE-11-16 (2011)
8. Klare, B., Jain, A.: On a taxonomy of facial features. In: *Proceedings of the IEEE Conference on Biometrics: Theory, Applications and Systems*, Washington, DC (2010)
9. Lee, Y., Micheals, R.J., Phillips, P.J., Filliben, J.J.: Robust iris recognition baseline for the grand challenge. *Nat. Inst. Stand. Technol. NISTIR 7777*, 1–8 (2011)
10. Liu, C., Wechsler, H.: Gabor feature based classification using the enhanced fisher linear discriminant model for face recognition. *IEEE Trans. Image Process.* **11**(4), 467–476 (2002)
11. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision* **60**(2), 91–110 (2004)

12. Matey, J.R., Naroditsky, O., Hanna, K., Kolczynski, R., Loiacono, D.J., Mangru, S., Tinker, M., Zappia, T.M., Zhao, W.Y.: Iris on the move: acquisition of images for iris recognition in less constrained environments. *Proc. IEEE* **94**(11), 1936–1947 (2006)
13. Meyers, E., Wolf, L.: Using biologically inspired features for face processing. *Int. J. Comput. Vis.* **76**(1), 93–104 (2008)
14. M, R., Poggio, T.: Hierarchical models of object recognition in cortex. *Nat. Neurosci.* **2**(11), 1019–1025 (1999)
15. NIST: Face and ocular challenge series. <http://www.nist.gov/itl/iad/ig/focs.cfm>
16. Ojala, T., Pietikäinen, M., Mäenpää, T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**(7), 971–987 (2002)
17. Pauca, V.P., Forkin, M., Xu, X., Plemmons, R., Ross, A.A.: Challenging ocular image recognition. *Proc. SPIE 8029 Biometric Technology for Human Identification VIII*, pp. 8029IV-1-8029IV-13 (2011)
18. Phillips, P.J., Member, S., Bowyer, K.W., Flynn, P.J., Member, S.: Comments on the casia version 1.0 iris dataset. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**, 1869–1870 (2007)
19. Phillips, P.J., Flynn, P.J., Beveridge, J.R., Scruggs, W.T., O’Toole, A.J., Bolme, D., Bowyer, K.W., Draper, B.A., Givens, G.H., Lui, Y.M., Sahibzada, H., Scallan Iii, J.A., Weimer, S.: Overview of the multiple biometrics grand challenge. In: *Proceedings of the Third International Conference on Advances in Biometrics, ICB ’09*, pp. 705–714. Springer, Berlin/Heidelberg (2009)
20. Phillips, P., Scruggs, W., O’Toole, A., Flynn, P., Bowyer, K., Schott, C., Sharpe, M.: FRVT 2006 and ICE 2006 large-scale experimental results. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(5), 831–846 (2010)
21. Ross, A., Jillela, R., Smereka, J., Boddeti, V.N., VijayaKumar, B.V.K., Barnard, R., Hu, X., Pauca, P., Plemmons, R.: Matching highly non-ideal ocular images: An information fusion approach. *Proceedings of the 5th IAPR International Conference on Biometrics (ICB)*, New Delhi (2012)
22. Shen, L., Bai, L.: A review on gabor wavelets for face recognition. *Pattern Anal. Appl.* **9**, 273–292 (2006)
23. Tan, X., Triggs, B.: Enhanced local texture feature sets for face recognition under difficult lighting conditions. *IEEE Trans. Image Process.* **19**(6), 1635–1650 (2010)
24. Thomee, B., Bakker, E.M., Lew, M.S.: TOP-SURF: a visual words toolkit. In: *ACM Multimedia’10*, pp. 1473–1476, Firenze, Italy (2010)
25. Wiskott, L., Fellous, J.M., Kuiger, N., von der Malsburg, C.: Face recognition by elastic bunch graph matching. *IEEE Trans. Pattern Anal. Mach. Intell.* **19**(7), 775–779 (1997)

Chapter 16

An Introduction to the IrisCode Theory

Adams Wai Kin Kong, David Zhang, and Mohamed Kamel

Abstract IrisCode is the most successful iris recognition method. Developed for over 18 years, IrisCode still dominates the market even though numerous iris recognition algorithms have been proposed in the academics. Currently, more than 60 million people have been mathematically enrolled by this algorithm. Its computational advantages, including high matching speed, predictable false acceptance rates, and robustness against local brightness and contrast variations, play a significant role in its commercial success. To further these computational advantages, researchers have modified this algorithm to enhance iris recognition performance and recognize other biometric traits (e.g., palm print). Many scientific papers on iris recognition have been published, but its theory is almost completely ignored. In this chapter, we will report our most recent theoretical work on the IrisCode.

A.W.K. Kong (✉)

Forensics and Security Laboratory, School of Computer Engineering, Nanyang Technological University, Nanyang Avenue, Singapore, Singapore 639798
e-mail: AdamsKong@ntu.edu.sg

D. Zhang

Biometrics Research Centre Department of Computing, The Hong Kong Polytechnic University, Kowloon, Hong Kong
e-mail: csdzhang@comp.polyu.edu.hk

M. Kamel

Pattern Analysis and Machine Intelligence Research Group, University of Waterloo, 200 University Avenue West, Waterloo, ON, Canada
e-mail: mkamel@pami.uwaterloo.ca

16.1 Introduction

IrisCode¹ has drawn considerable attention in the last 18 years [1–3] because of its commercial success, and outstanding performance in terms of speed and accuracy. The success of IrisCode is due to the characteristics of the human iris (e.g., it contains rich texture information and is stable for a long period of time) and also its algorithmic design. IrisCode, which is an algorithm based on coarse phase information in irises, has a number of desirable properties, including rapid matching, binomial impostor distribution, robustness against contrast and brightness variations, and a predictable false acceptance rate. Therefore, researchers have modified this algorithm to enhance iris recognition performance and recognize other biometric traits. It is extremely important to completely understand IrisCode because over 60 million people are using this algorithm, and many other biometric algorithms are extended from IrisCode.

In the past two decades, numerous papers on iris recognition have been published, but our understanding of IrisCode is still very incomplete because very limited work has been devoted to its theory. In the original paper on the IrisCode, Daugman reported that the bits of “0” and “1” in IrisCodes are equally probable and the importer hamming distance of IrisCode follows a binomial distribution with a high degree of freedom [1]. In addition to Daugman, Yao et al. investigated the relationship between the distributions of bits and the bandwidth of the Gabor filter [4]. In contrast to Daugman’s conclusion, they stated that the bits of “0” and “1” in IrisCodes are not equally probable. Their study was based on a nonzero DC Gabor filter, while Daugman’s analysis was based on zero DC Gabor filters. Hollingsworth et al. attempted to look for the most stable bits in their codes, which were not generated from 2D Gabor filters [5].

Many people believe that the inequality generating bit pairs are the core of the IrisCode because they quantize the phase information in bitwise format for high-speed matching. Thus, some researchers have replaced the Gabor filters with other linear filters or functions to design their new coding methods. Due to the operators “ \leq ” and “ $>$,” some researchers claim that IrisCode is an ordinal feature. This claim is controversial because the bit pairs in IrisCodes are generated from a periodic feature of the Gabor phase (see Sect. 16.2 in this chapter). Furthermore, some researchers believe that the impostor distributions of their coding methods also follow binomial distributions [6]. One clear weakness of the current understanding is that only one bit of information can be extracted from each filtering. It is an inflexible representation. In this chapter, we will summarize our recent theoretical work for a better understanding of the IrisCode [7–9].

Before reporting on our theoretical work, we will offer a brief computational summary of the IrisCode [1] for the sake of presentation convenience and consistency of notations. Two-dimensional Gabor filters with zero DC are applied to an

¹In this paper, IrisCode is used interchangeably for both the method and features of iris recognition developed by Daugman. Recently, this method has also been dubbed the Daugman algorithm.

iris image in a dimensionless polar coordinate system, $I(\rho, \phi)$. The complex Gabor response is quantized into two bits by using the following inequalities:

$$h_{\text{Re}} = 1 \quad \text{if} \quad \text{Re} \left(\int_{\rho} \int_{\phi} I(\rho, \phi) e^{\frac{-(r_0-\rho)^2}{\alpha^2}} e^{\frac{-(\theta_0-\phi)^2}{\beta^2}} e^{-i\omega(\theta_0-\phi)} \rho d\rho d\phi \right) \geq 0, \quad (16.1)$$

$$h_{\text{Re}} = 0 \quad \text{if} \quad \text{Re} \left(\int_{\rho} \int_{\phi} I(\rho, \phi) e^{\frac{-(r_0-\rho)^2}{\alpha^2}} e^{\frac{-(\theta_0-\phi)^2}{\beta^2}} e^{-i\omega(\theta_0-\phi)} \rho d\rho d\phi \right) < 0, \quad (16.2)$$

$$h_{\text{Im}} = 1 \quad \text{if} \quad \text{Im} \left(\int_{\rho} \int_{\phi} I(\rho, \phi) e^{\frac{-(r_0-\rho)^2}{\alpha^2}} e^{\frac{-(\theta_0-\phi)^2}{\beta^2}} e^{-i\omega(\theta_0-\phi)} \rho d\rho d\phi \right) \geq 0, \quad (16.3)$$

$$h_{\text{Im}} = 0 \quad \text{if} \quad \text{Im} \left(\int_{\rho} \int_{\phi} I(\rho, \phi) e^{\frac{-(r_0-\rho)^2}{\alpha^2}} e^{\frac{-(\theta_0-\phi)^2}{\beta^2}} e^{-i\omega(\theta_0-\phi)} \rho d\rho d\phi \right) < 0, \quad (16.4)$$

where r_0 , θ_0 , ω , α , and β are the parameters of the Gabor filters [2]. Initially, Daugman used the bitwise hamming distance defined as $HD = \sum_{i=1}^{2048} (A_i \otimes B_i) / 2048$, where A_i and B_i are the bits in two IrisCodes and \otimes represents the bitwise operator, XOR. Currently, masks are employed to exclude the corrupted bits from eyelashes, reflections, eyelids, and low signal-to-noise ratios [2]. The hamming distance between two IrisCodes is redefined as

$$HD = \frac{\sum_{i=1}^{2048} ((A_i \otimes B_i) \cap (A_i^M \cap B_i^M))}{\sum_{i=1}^{2048} (A_i^M \cap B_i^M)}, \quad (16.5)$$

where A^M and B^M are the masks of IrisCodes A and B , respectively, and \cap represents bitwise operator AND.

The rest of this chapter is organized as follows: Sect. 16.2 will report a recent theory on the IrisCode. Section 16.3 will demonstrate an application of the theory. Section 16.4 gives some concluding remarks.

16.2 Theoretical Properties of IrisCode

In this section, we will prove that the IrisCode is a clustering algorithm with four prototypes; the locus of a Gabor function is a two-dimensional ellipse with respect to a phase parameter and can be approximated by a circle in many cases; the Gabor

function can be considered as a phase-steerable filter, and the bitwise hamming distance can be regarded as a bitwise phase distance. Moreover, we will present a unified framework to link different coding methods.

16.2.1 *IrisCode Is a Clustering Process*

Let us define a filter-generating function, $Z(\varphi, \rho, \phi) = (\cos(\varphi)M_R(\rho, \phi) + \sin(\varphi)M_I(\rho, \phi))$, where

$$M_R(\rho, \phi) = e^{-\frac{-(r_0-\rho)^2}{\alpha^2}} e^{-\frac{-(\theta_0-\phi)^2}{\beta^2}} (\cos(-\omega(\theta_0 - \phi))), \quad (16.6)$$

$$M_I(\rho, \phi) = e^{-\frac{-(r_0-\rho)^2}{\alpha^2}} e^{-\frac{-(\theta_0-\phi)^2}{\beta^2}} (\sin(-\omega(\theta_0 - \phi))), \quad (16.7)$$

represent the real and imaginary parts of a Gabor filter, respectively. For the sake of convenience, we use $Z(\varphi)$, M_R , and M_I to denote $Z(\varphi, \rho, \phi)$, $M_R(\rho, \phi)$, and $M_I(\rho, \phi)$, respectively. The same notations for other symbols will be used in the rest of this chapter. It should be noted that $Z(\varphi)$ is a periodic function with respect to the phase parameter φ , i.e., $Z(\varphi) = Z(\varphi + 2\pi)$ and M_R are not a zero DC filter. In this chapter, we will suppose that the DC of iris patches for filtering has been removed, and therefore, we can retain the DC components in Eq. 16.6.

By substituting $5\pi/4$, $7\pi/4$, $\pi/4$, and $3\pi/4$ to φ , we can obtain four filters:

$$Z_0 = Z\left(\frac{5\pi}{4}\right) = -\frac{(M_R + M_I)}{\sqrt{2}}, \quad (16.8)$$

$$Z_1 = Z\left(\frac{7\pi}{4}\right) = \frac{(M_R - M_I)}{\sqrt{2}}, \quad (16.9)$$

$$Z_2 = Z\left(\frac{\pi}{4}\right) = \frac{(M_R + M_I)}{\sqrt{2}}, \quad (16.10)$$

$$Z_3 = Z\left(\frac{3\pi}{4}\right) = -\frac{(M_R - M_I)}{\sqrt{2}}. \quad (16.11)$$

Note that $Z_0 = -Z_2$ and $Z_1 = -Z_3$. These four filters shown in Fig. 16.1 will be considered as four cluster centers. To cluster an input iris image, we define $j = \arg \max_i \int \int \rho I Z_i d\rho d\varphi / \|\rho I\| \|Z_i\|$ as a clustering condition, where j is called the winning index and I is the iris image in the dimensionless polar coordinate system. The clustering condition is in fact the cosine measure between Z_i and ρI . Because $\|Z_u\| = \|Z_v\|$, for all u and v , the clustering condition can be simplified as:




Fig. 16.1 The four filters used in the IrisCode, (a) Z_0 , (b) Z_1 , (c) Z_2 and (d) Z_3

Table 16.1 Comparison of IrisCode, winning index and coded winning index

Winning index	Coded winning indexes		IrisCode	
	Bit 2	Bit 1	h_{Im}	h_{Re}
1	0	0	0	0
2	0	1	0	1
3	1	1	1	1
4	1	0	1	0

$$j = \arg \max_i \left(\int_{\rho} \int_{\varphi} \rho I Z_i d\rho d\varphi \right). \tag{16.12}$$

To prove that $\|Z_u\| = \|Z_v\|$ for all u and v , the orthogonal property between M_R and M_I , i.e., $\int_{\rho} \int_{\varphi} M_I M_R d\rho d\varphi = 0$, should be used. By using the coding table under heading “Coded winning indexes” in Table 16.1 to encode the winning indexes and comparing the resultant bits with the bits of IrisCodes, we can easily note that these two sets of bits are equivalent. In other words, IrisCode is a clustering process, in which the cosine measure is used as a clustering criterion and Z_0, Z_1, Z_2 , and Z_3 are the four prototypes.

16.2.2 Properties of the Filter-Generating Function

In this subsection, we will provide the properties of the filter-generating function. First of all, we will discuss the physical meaning of Z and then the locus of φ . Next,

we will demonstrate that filters generated from the filter-generating function can be regarded as phase-steerable filters. Finally, we will pinpoint the relationship between φ and $\tan^{-1} \left(\frac{\int_{\rho} \int_{\varphi} \rho I M_I d\rho d\varphi}{\int_{\rho} \int_{\varphi} \rho I M_R d\rho d\varphi} \right)$ which is commonly employed as a feature in image-based applications.

16.2.2.1 The Physical Meaning of Z and the Locus of φ

By substituting Eqs. 16.6 and 16.7 into Z and using a compound angle formula, we can derive that

$$Z(\phi) = e^{-(r_0-\rho)^2/\sigma^2} e^{-(\theta_0-\varphi)^2/\beta^2} \cos(-\omega(\theta_0 - \varphi) - \phi). \tag{16.13}$$

Clearly, Z is a Gabor function, and φ is its phase.

Now, we will prove that the locus of φ is an ellipse. First, we discretize M_R , M_I , and Z to obtain three vectors, \mathbf{M}_R , \mathbf{M}_I , and $\mathbf{Z}(\phi)$, respectively. The unit vectors of \mathbf{M}_R and \mathbf{M}_I are defined as $\mathbf{v}_R = \mathbf{M}_R / \|\mathbf{M}_R\|$ and $\mathbf{v}_I = \mathbf{M}_I / \|\mathbf{M}_I\|$. By using them to decompose $\mathbf{Z}(\phi)$, we can obtain

$$\vec{Z}(\phi) = \left(\cos(\phi) \|\mathbf{M}_R\| \mathbf{v}_R + \sin(\phi) \|\mathbf{M}_I\| \vec{v}_I \right). \tag{16.14}$$

Clearly, $\mathbf{Z}(\phi)$ is in a two-dimensional space spanned by \mathbf{v}_R and \mathbf{v}_I , and its coordinate in this space is $(\|\mathbf{M}_R\| \cos(\phi), \|\mathbf{M}_I\| \sin(\phi))$, which fulfills the equality

$$\frac{(\|\mathbf{M}_R\| \cos(\phi))^2}{\|\mathbf{M}_R\|^2} + \frac{(\|\mathbf{M}_I\| \sin(\phi))^2}{\|\mathbf{M}_I\|^2} = 1. \tag{16.15}$$

Equation 16.15 pinpoints that the locus of φ is an ellipse. It can be proven that $\lim_{k \rightarrow \infty} (\|\mathbf{M}_R\|^2 - \|\mathbf{M}_I\|^2) = 0$, where $k = \omega\beta$, which is to say that the locus of φ is a circle because Eq. 16.15 can be simplified as $(\|\mathbf{M}_R\| \cos(\phi))^2 + (\|\mathbf{M}_I\| \sin(\phi))^2 = \|\mathbf{M}_R\|^2$. The proof is given in Appendix A. Figure 16.2 illustrates the loci and the filters.

16.2.2.2 Gabor Function and Phase-Steerable Filter

By substituting $Z(\varphi, \rho, \phi) = (\cos(\varphi)M_R(\rho, \phi) + \sin(\varphi)M_I(\rho, \phi))$ into $\int_{\rho} \int_{\varphi} \rho I Z(\varphi) d\rho d\varphi$, we can obtain




Fig. 16.2 Illustration of the loci of φ and filters generated from Z

$$\cos(\varphi) \int \int_{\rho \varphi} \rho IM_R d\rho d\varphi + \sin(\varphi) \int \int_{\rho \varphi} \rho IM_I d\rho d\varphi. \tag{16.16}$$

Equation 16.16 pinpoints that no matter how many filters are generated from Z , two filtering operations $\int \int_{\rho \varphi} \rho IM_R d\rho d\varphi$ and $\int \int_{\rho \varphi} \rho IM_I d\rho d\varphi$ are enough to perform the clustering calculation based on $j = \arg \max_i \int \int_{\rho \varphi} \rho IZ(\varphi_i) d\rho d\varphi / \|\rho I\| \|Z(\varphi_i)\|$ because $\|Z(\varphi_i)\|$ can be precomputed. Clearly, Z is a phase-steerable filter.

16.2.2.3 The Difference Between Gabor Phase and φ

Let $\phi_c = \tan^{-1} \left(\frac{\int \int_{\rho \varphi} \rho IM_I d\rho d\varphi}{\int \int_{\rho \varphi} \rho IM_R d\rho d\varphi} \right)$ be a Gabor phase. Researchers generally consider ϕ_c as a phase of a local patch, instead of φ . Their relationship is not clear. It can be proven that $|\varphi - \phi_c| \leq \tan^{-1}(1/\sqrt{K}) - \tan^{-1}(\sqrt{K})$, where $K = \|M_I\|^2 / \|M_R\|^2$. This proof is given in [8], which is lengthy. Note that when $K = 1$, $\phi_c = \varphi$ and the locus of φ is a circle.




Fig. 16.3 Illustration of phase distance

16.2.3 Phase Distance and Its Relationship with Bitwise Hamming Distance

In this subsection, we will demonstrate that the phase distance can be represented as the bitwise hamming distance. Assume that we generate a number of filters from Z for clustering based on the cosine measure. If the number of filters are more than four, we need a new distance measure to compare the prototypes, i.e., $Z(\omega)$ and $Z(\gamma)$. In Sect. 16.2.2.1, we have already shown that Z is on a two-dimensional space and the locus of φ is an ellipse, and therefore, $Z(\omega)$ and $Z(\gamma)$ can be measured by the phase distance between ω and γ defined as $\min(|\omega - \gamma|, 2\pi - |\omega - \gamma|)$. If ω and γ are generated from uniform sampling, i.e., $\omega = 2\pi p/2n$ and $\gamma = 2\pi q/2n$, where p and q are two integers between 0 and $2n - 1$, the phase distance can be rewritten as $\min(\frac{\pi}{n}|p - q|, \frac{\pi}{n}(2n - |p - q|))$. n is called the order of the coding scheme. This formula can be further simplified as $\min(|p - q|, (2n - |p - q|))$ if we consider π/n as one unit distance. Note that n is equal to 2 for the IrisCode because it uses four prototypes in the clustering process. Figure 16.3 illustrates the phase distance and winning indexes.

The relationship between the bitwise hamming distance employed in the IrisCode and phase distance based on the integer representations of ω and γ is still not clear. By using the coding scheme given in Fig. 16.4 to encode p and q , the phase distance can be computed through bitwise hamming distance, i.e., $\sum_{i=1}^n b_{i,q} \otimes b_{i,p} = \min(|p - q|, 2n - |p - q|)$. The proof of this equation is given in Appendix B. Table 16.2 shows the coding tables for $n = 3$ and $n = 4$. It clearly indicates that the difference between two adjacent winning indexes is only one bit, and most importantly, this coding scheme retains the phase distance. It should be emphasized that in rotating between any adjacent phase quadrant, only a single bit in the IrisCode changes, while in rotating between one phase quadrant to the opposite phase quadrant, both two bits in the IrisCode change, which is to say that the distances between phase quadrants in the IrisCode are also retained. By exploiting the relationship between the bitwise hamming distance and the phase distance, we can extend the IrisCode from using the four prototypes in the clustering process to using $2n$ prototypes and high-speed matching.

Fig. 16.4 Pseudocode of the coding table

```

if  $j \leq n$  and  $1 \leq i < j$ ,
     $b_{i,j} = 1$ 
elseif  $j > n$  and  $j - n \leq i \leq n$ ,
     $b_{i,j} = 1$ 
else
     $b_{i,j} = 0$ 
    
```

Table 16.2 The coding tables for (a) $n = 3$ and (b) $n = 4$

(a)			
Winning indexes	Bit 2	Bit 1	Bit 0
0	0	0	0
1	0	0	1
2	0	1	1
3	1	1	1
4	1	1	0
5	1	0	0

(b)				
Winning indexes	Bit 3	Bit 2	Bit 1	Bit 0
0	0	0	0	0
1	0	0	0	1
2	0	0	1	1
3	0	1	1	1
4	1	1	1	1
5	1	1	1	0
6	1	1	0	0
7	1	0	0	0

16.2.4 A Generalized Framework from the IrisCode

Many coding methods that are very similar to the IrisCode [13–20] have been proposed for iris and palm print recognition [10, 11]. A common approach is to substitute other linear transforms or filters for the Gabor filters in the IrisCode. We have already proven that the IrisCode is a clustering with four prototypes. It is not surprising that these coding methods can also be considered as clustering algorithms, but with two prototypes. Mathematically, these methods can be represented by the following two equations:

$$h = 1 \quad \text{if} \quad \int_{\rho} \int_{\varphi} FId \rho d\varphi \geq 0, \tag{16.17}$$

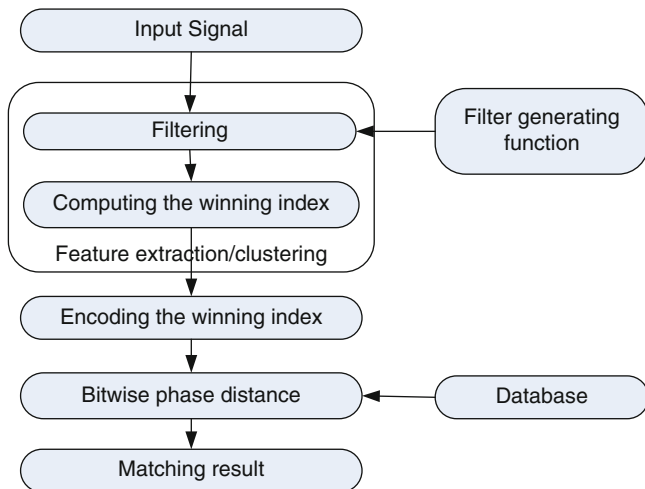


Fig. 16.5 A common framework employed by most of the existing coding methods

$$h = 0 \quad \text{if} \quad \int_{\rho} \int_{\varphi} FId \rho d\varphi < 0, \tag{16.18}$$

where h is a resultant bit and F is a linear filter employed in these methods. Let $(-1)^{v+1}F$ be a filtering-generating function, where $v \in \{0, 1\}$. It can only generate two filters, F and $-F$. Note that these two filters have the same power, i.e., $\|F\| = \|-F\|$; the clustering criterion based on the cosine measure can be simplified as

$$j = \arg \max_i \left(\int_{\rho} \int_{\varphi} (-1)^{i+1} FId \rho d\varphi \right). \tag{16.19}$$

If $\int_{\rho} \int_{\varphi} FId \rho d\varphi > \int_{\rho} \int_{\varphi} -FId \rho d\varphi, j = 1$ and $\int_{\rho} \int_{\varphi} FId \rho d\varphi > 0$. If $\int_{\rho} \int_{\varphi} -FId \rho d\varphi > \int_{\rho} \int_{\varphi} FId \rho d\varphi, j = 0$ and $\int_{\rho} \int_{\varphi} FId \rho d\varphi < 0$. By using the first-order coding scheme defined in Fig. 16.4 to encode the winning index j , we get back Eqs. 16.17 and 16.18.

In addition to these coding schemes, other coding methods that use log Gabor filters and Gabor filters employ the order 2 coding scheme, which is to say that in their clustering processes, they have four prototypes, as with IrisCode [17, 20]. Our competitive code for palm print identification [12] uses the negative real part of a Gabor function as a filter-generating function to produce six filters for the clustering process. It is the first coding method that uses an order higher than 2. All of these coding methods can be put into a unified framework, which is composed of a filter-generating function, clustering, coding scheme, and bitwise phase matching as illustrated in Fig. 16.5.

16.3 An Application Based on the Theoretical Results

To demonstrate the theoretical results that have an impact on applications, an algorithm that uses higher precision to present the Gabor phase was designed to improve matching accuracy for iris recognition. Its computational steps are given below:

1. Compute $\int \int_{\rho \varphi} \rho IM_{Rd} \rho d\varphi$ and $\int \int_{\rho \varphi} \rho IM_{Id} \rho d\varphi$.
2. Use the results in Sect. 16.2.2, i.e., $\cos(i\pi/n) \int \int_{\rho \varphi} \rho IM_{Rd} \rho d\varphi + \sin(i\pi/n) \int \int_{\rho \varphi} \rho IM_{Id} \rho d\varphi$ to compute the winning index, $j = \arg \max_i \int \int_{\rho \varphi} \rho IZ(i\pi/n) d\rho d\varphi / \|\rho I\| \|Z(i\pi/n)\|$.
3. Use the results in Sect. 16.2.3 to encode the winning indexes and bitwise hamming distance to perform high-speed matching.

The algorithm is called precise phase representation, which inherits the properties of the IrisCode (e.g., high-speed matching). The precision of the phase is controlled by the parameter n . In total, n bits are allocated for a phase value. Note that the IrisCode uses two bits.

16.3.1 Database

The CASIA-1 iris database that contains 756 images from 108 irises is used for this evaluation. One hundred and forty images were used to train the parameters of the Gabor filters. The remaining was used for the following tests. In these experiments, we used d' index and receiver operating characteristic (ROC) curves as a plot of the genuine acceptance rates against the false acceptance rates for all possible operating points as performance indexes. The experimental results based on the West Virginia University (WVU) iris database are reported in [9].

16.3.2 Experimental Results

We compared the original IrisCode and precise phase representations for $n = 3, 4,$ and 5 to demonstrate the effectiveness of the latter. Irises from the same and different individuals were matched to respectively obtain 1,848 genuine hamming distances and 187,572 impostor hamming distances for each representation. These distances were used to estimate the genuine and impostor distributions. The genuine and impostor distributions of the IrisCode and precise phase representations for $n = 3, 4,$ and 5 are given in Fig. 16.6d, respectively. Figure 16.7 depicts the corresponding ROC curves for comparison. It shows that the precise phase representation of

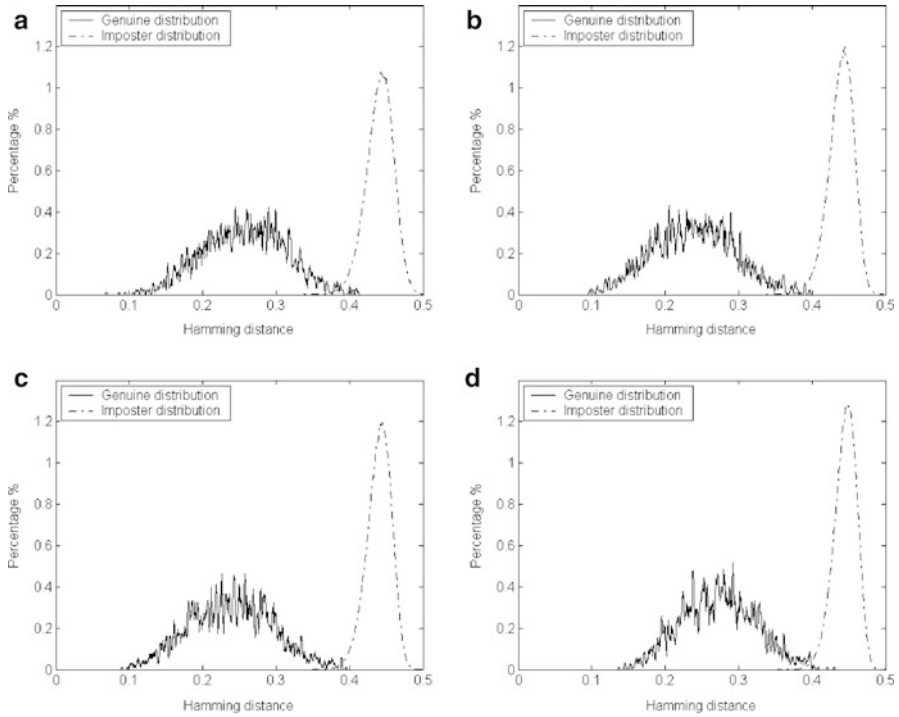


Fig. 16.6 Genuine and impostor distributions for (a) the IrisCode, (b)–(d) precise phase representations for $n=3-5$, respectively




Fig. 16.7 ROC curves of the IrisCode and precise phase representations for $n=3, 4$, and 5

the order 4 always provides the best results. Its equal error rate is 0.32%. In a comparison of the IrisCode and the precise phase representation of the order 4, the latter has a 3.1% improvement in the genuine acceptance rate when the false acceptance rate is $5.3 \times 10^{-4}\%$. However, an increase in the precision of the phase does not always improve accuracy, because some encoded bits are in fact not stable. It is also the case in eigenface in which increasing numbers of principle components do not always increase accuracy. Note that the matching speed of the precise phase representation of the order 4 is slower than the IrisCode; roughly speaking, it is half of that of the IrisCode since the latter only uses two bits to represent a phase value while the former requires four bits. For some applications, such as identifying a person in a midsized company for access control, one million comparisons per second is not necessary. We can use the precise phase representation to obtain high accuracy for these applications. It should be emphasized that the precise phase representation is a flexible representation to balance speed and accuracy.

16.4 Conclusion

The IrisCode has been developed for nearly two decades, and more than 60 million people are using this algorithm. Although numerous papers on iris recognition have been published, our theoretical understanding of the IrisCode is not yet complete because only limited research efforts have been put into analyzing the IrisCode. In this chapter, we have summarized our recent work that fills this knowledge gap [7–9]. If readers are interested in this theoretical development, please keep track of the author’s publications.

Acknowledgements We would like to thank CASIA for sharing their database. This work is partially supported by the Ministry of Education, Singapore, through AcRF Tier 1.

A.1 Appendices

A.1.1 Appendix A

This appendix proves that $\|M_R\| = \|M_I\|$ when $k \rightarrow \infty$, where $k = \omega\beta$.

Considering $\|M_R\|^2 - \|M_I\|^2$

$$\begin{aligned} &= \iint \left(e^{-\frac{\rho^2}{\alpha^2}} e^{-\frac{\varphi^2}{\beta^2}} \cos(\omega\varphi) \right)^2 d\rho d\varphi - \iint \left(e^{-\frac{\rho^2}{\alpha^2}} e^{-\frac{\varphi^2}{\beta^2}} \sin(\omega\varphi) \right)^2 d\rho d\varphi \\ &= \iint e^{-\frac{2\rho^2}{\alpha^2}} e^{-\frac{2\varphi^2}{\beta^2}} (\cos^2(\omega\varphi) - \sin^2(\omega\varphi)) d\rho d\varphi \end{aligned}$$

$$\begin{aligned}
&= \int e^{-\frac{2\rho^2}{\alpha^2}} d\rho \int e^{-\frac{2\varphi^2}{\beta^2}} (\cos^2(\omega\varphi) - \sin^2(\omega\varphi)) d\varphi \\
&= \frac{\alpha\sqrt{2\pi}}{2} \int e^{-\frac{2\varphi^2}{\beta^2}} (\cos^2(\omega\varphi) - \sin^2(\omega\varphi)) d\varphi
\end{aligned}$$

Let $\gamma = \frac{k}{\beta}\varphi$. Thus,

$$\begin{aligned}
&= \frac{\alpha\beta\sqrt{2\pi}}{2k} \int e^{-\frac{2\gamma^2}{k^2}} (\cos^2(\gamma) - \sin^2(\gamma)) d\gamma \\
&= \frac{\alpha\beta\sqrt{2\pi}}{2k} \int e^{-\frac{2\gamma^2}{k^2}} \cos(2\gamma) d\gamma
\end{aligned}$$

Let $2\gamma = \tau$

$$\begin{aligned}
&= \frac{\alpha\beta\sqrt{2\pi}}{4k} \int e^{-\frac{\tau^2}{2k^2}} \cos(\tau) d\tau \\
&= \frac{\alpha\beta\sqrt{2\pi}}{4k} \sqrt{2\pi} k e^{-\frac{k^2}{2}} \\
&= \frac{1}{2} \alpha\beta\pi e^{-\frac{k^2}{2}}.
\end{aligned}$$

$\|M_R\|^2 - \|M_I\|^2$ is always greater than zero because α , β , and k are greater than zero. Note that $\lim_{k \rightarrow \infty} (\|M_R\|^2 - \|M_I\|^2) = \lim_{k \rightarrow \infty} \frac{1}{2} \alpha\beta\pi e^{-\frac{k^2}{2}} = 0$.

A.1.2 Appendix B

This appendix shows that the phase distance can be calculated through bitwise hamming distance.

Let two winning indexes be $j - l$ and $j - l + k$, where $1 \leq j \leq j + k < 2n$. Their phase distance is $\min(k, 2n - k)$. By using the coding scheme given in Fig. 16.4, they are represented by the j th and $j + k$ th column vectors of matrix B . We would like to prove that

$$\sum_{i=1}^n b_{i,j} \otimes b_{i,j+k} = \min(k, 2n - k).$$

Because all $b_{i,j}$ are either one or zero, $\sum_{i=1}^n b_{i,j} \otimes b_{i,j+h} = \sum_{i=1}^n b_{i,j} - b_{i,j+k}$.

Case 1: If $j \leq n$ and $j+k \leq n$:

From the definition of A , we know $\sum_{i=1}^n |b_{i,j} - b_{i,j+k}| = k$.

Case 2: If $j > n$ and $j+k > n$:

As in Case 1, we know $\sum_{i=1}^n |b_{i,j} - b_{i,j+k}| = k$.

Case 3: If $j \leq n$ and $j+k > n$ and $k \leq n$:

$$\text{Consider } b_{i,j} = 1 \text{ and } b_{i,j+k} = 1. \quad (16.20)$$

From the definition of A , we have $1 \leq i < j$ and $j+k-n \leq i \leq n$.

Then, $j+k-n \leq i < j$.

The number of i that satisfies condition (16.20) is

$$\max(0, j - (j+k-n)). \quad (16.21)$$

Since $k \leq n$, $\max(0, j - (j+k-n)) = n - k$.

$$\text{Consider } b_{i,j} = 0 \text{ and } b_{i,j+k} = 0. \quad (16.22)$$

From the definition of A , we have $i \geq j$ and $i < j+k-n$.

Then $j \leq i < j+k-n$.

The number of i that satisfies condition (16.22) is

$$\max(0, j+k-n-j). \quad (16.23)$$

Since $k \leq n$, $\max(0, k-n) = 0$.

Thus, $\sum_{i=1}^n |b_{i,j} - b_{i,j+k}| = n - (n-k) = k$.

Case 4: If $j \leq n$ and $j+k > n$ and $k > n$:

$$\text{Consider } b_{i,j} = 1 \text{ and } b_{i,j+k} = 1. \quad (16.24)$$

From (16.21), the number of i that satisfies condition (16.24) is $\max(0, j - (j+k-n))$.

Since $k > n$, $\max(0, n-k) = 0$.

$$\text{Consider } b_{i,j} = 0 \text{ and } b_{i,j+k} = 0. \quad (16.25)$$

From (16.23), the number of i that satisfies condition (16.25) is $\max(0, j+k-n-j)$.

Since $k > n$, $\max(0, k-n) = k-n$.

Thus, $\sum_{i=1}^n |b_{i,j} - b_{i,j+k}| = n - (k-n) = 2n - k$.

Thus, $\sum_{i=1}^n b_{i,j} \otimes b_{i,j+k} = k$ for Cases 1–3, and $\sum_{i=1}^n b_{i,j} \otimes b_{i,j+k} = 2n - k$ for Case 4. Since $2n - k \geq k$ for Cases 1–3 and $2n - k < k$ for Case 4, $\sum_{i=1}^n b_{i,j} \otimes b_{i,j+k} = \min(k, 2n - k)$.

References

1. Daugman, J.G.: High confidence visual recognition of persons by a test of statistical independence. *IEEE Trans. Pattern Anal. Mach. Intell.* **15**(11), 1148–1161 (1993)
2. Daugman, J.: How iris recognition works. *IEEE Trans. Circuit. Syst. Video Technol.* **14**(1), 21–30 (2004)
3. Daugman, J.: Probing the uniqueness and randomness of IrisCode: results from 200 billion iris cross-comparisons. *Proc. IEEE* **94**, 11 (2006)
4. Yao, P., Li, J., Ye, X., Zhuang, Z., Li, B.: An analysis and improvement of an iris identification algorithm. *Proceeding of the 18th ICPR vol. 4*, pp. 362–365. IEEE, Piscataway, NJ (2006)
5. Hollingsworth, K.P., Bowyer, K.W., Flynn, P.J.: The best bits in an iris code. *TPAMI* **31**(6), 964–973 (2009)
6. Masek, L.: Recognition of human iris patterns for biometric identification. Bachelor thesis, The University of Western Australia
7. Kong, A.W.K.: Palmprint identification based on generalization of IrisCode. PhD thesis, University of Waterloo (2007)
8. Kong, A.: An analysis of Gabor detection. *International Conference on Image Analysis and Recognition (ICIAR)*, Halifax, Canada, 6–8 July 2009
9. Kong, A.W.K., Zhang, D., Kamel, M.: An analysis of IrisCode. *IEEE Trans. Image Process.* **19**(2), 522–532 (2010)
10. Kong, A.W.K., Zhang, D., Kamel, M.: Palmprint identification using feature-level fusion. *Pattern Recognit.* **39**, 478–487 (2006)
11. Zhang, D., Kong, W.K., You, J., Wong, M.: On-line palmprint identification. *IEEE Trans. Pattern Anal. Mach. Intell.* **25**(9), 1041–1050 (2003)
12. Kong, A.W.K., Zhang, D., Kamel, M.: An analysis of brute-force break-ins of a palmprint authentication system. *IEEE Trans. Syst. Man and Cybern. Part B* **36**(5), 1201–1205 (2006)
13. Ma, L., Tan, T., Wang, Y., Zhang, D.: Efficient iris recognition by characterizing key local variations. *IEEE Trans. Image Process.* **13**(6), 739–750 (2004)
14. Krichen, E., Mellakh, M.A., Garcia-Salicetti, S., Dorizzi, B.: Iris identification using wavelet packets. *Proc. Int. Conf. Pattern Recognit.* **4**, 226–338 (2004)
15. Noh, S.I., Bae, K., Park, Y., Kim, J.: A Novel Method to Extract Features for Iris Recognition System. *Lecture Notes in Computer Science*, vol. 2688, pp. 861–868. Springer, Berlin (2003)
16. Bea, K., Noh, S., Kim, J.: Iris Feature Extraction Using Independent Component Analysis. *Lecture Notes in Computer Science*, vol. 2688, pp. 838–844. Springer, Berlin (2003)
17. Zhang, P.F., Li, D.S., Wang, Q.: A novel iris recognition method based on feature fusion. In: *Proceedings of the Third International Conference on Machine Learning and Cybernetics*, pp. 26–29. IEEE, Piscataway, NJ (2004)
18. Ea, T., Valentian, A., Rossant, F., Amiel, F. Amara, A.: Algorithm implementation for iris identification. In: *Proceeding of 48th Midwest Symposium on Circuits and Systems*, pp. 1207–1210. IEEE, Piscataway, NJ (2005)
19. Park, C.H., Lee, J.J., Oh, S.K., Song, Y.C., Choi, D.H., Park, K.H.: Iris Feature Extraction and Matching Based on Multiscale and Directional Image Representation. *LNCS*, vol. 2695, pp. 576–583. Springer, Berlin (2004)
20. Rydgren, E., Amiel, T.E.A.F., Rossant, F., Amara, A.: Iris features extraction using wavelet packets. *Proc. Int. Conf. Image Process.* **2**, 861–864 (2004)

Chapter 17

Application of Correlation Filters for Iris Recognition

B.V.K. Kumar, Jason Thornton, Marios Savvides, Vishnu Naresh Boddeti, and Jonathon M. Smereka

Abstract Excellent recognition accuracies have been reported when using iris images, particularly when high-quality iris images can be acquired. The best-known strategy for matching iris images requires segmenting the iris from the background, converting the segmented iris image from Cartesian coordinates to polar coordinates, using Gabor wavelets to obtain a binary code to represent that iris, and using the Hamming distances between such binary representations to determine whether two iris images match or do not match. However, some of the component operations may not work well when the iris images are of poor quality, perhaps as a result of the long distance between the camera and the subject. One approach to matching images with appearance variations is the use of correlation filters (CF). In this chapter, we discuss the use of CFs for iris recognition. CFs exhibit important benefits such as shift-invariance and graceful degradation and have proven worthy of consideration in other pattern recognition applications such as automatic target recognition. In this chapter, we will discuss the basics of CF design and show how CFs can be used for iris segmentation and matching.


17.1 Introduction

Among the many biometric modalities proposed to recognize individuals, iris recognition [1] has become a major research direction because of the excellent accuracy iris recognition methods seem to offer [2], particularly when the iris

B.V.K. Kumar (✉) • M. Savvides • V.N. Boddeti • J.M. Smereka
Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA
e-mail: kumar@ece.cmu.edu; marios@andrew.cmu.edu; naresh@cmu.edu;
jsmereka@andrew.cmu.edu

J. Thornton
MIT Lincoln Laboratory, Lexington, MA, USA
e-mail: jthornton@ll.mit.edu

Fig. 17.1 Iris is the texture-rich region between the *black* pupil and the *white* sclera



images are of high quality. As shown in Fig. 17.1, iris refers to the texture-rich region of the eye surrounded by the black pupil (inside) and the white sclera region (outside). It is believed that iris patterns do not change substantially over a lifetime and that an iris pattern is unique to an eye, i.e., iris patterns from the left eye and the right eye of the same individual are different.

The original iris recognition method pioneered by Daugman [3] consists of the following major components:

- Segmenting the iris from the pupil and the sclera by determining its inner and outer boundaries
- Mapping the segmented iris pattern from Cartesian coordinates to polar coordinates to normalize for iris size variations caused by the pupil dilation and other factors
- Producing a binary code from the phases of the inner products of Gabor wavelets (of different widths, orientations, locations, and frequencies) with the mapped (also called unwrapped) iris pattern
- Using the normalized Hamming distance (NHD) between the binary code of an enrolled iris image and the binary code of a query (also called probe) iris image as an indicator of the quality of the match between the two iris patterns

The above iris recognition method based on comparing the resulting binary codes is attractive for its high-speed matching and its excellent accuracy when the training and testing iris images are of good quality. However, iris images require near-infrared (NIR) illumination to produce images with good contrast, and NIR illumination levels cannot be too high because of safety considerations. As a result of the limitations on illumination levels, the iris images acquired from distance are not expected to be of sufficiently good contrast. To add to this challenge, iris images can also exhibit degradations such as occlusions caused by eyelids, non-frontal gazes, and specular reflections. In such challenging iris images, traditional binary code matching methods may not work well because iris regions cannot be easily segmented from the pupil and the sclera regions.

One attractive method to deal with degraded query images is based on correlating the query images with a template carefully designed from training images [4]. This method, known as correlation filter (CF) [5], has proven useful in other challenging image matching applications such as automatic target recognition [6] and face recognition [7]. In this chapter, we will discuss the basics of CF design and show how CFs can be used for iris segmentation and matching.

The rest of this chapter is organized as follows. Section 17.2 provides a brief review of CF design, and Sect. 17.3 shows how CFs can be used for iris segmentation. CFs can be used for matching segmented iris images as well as unsegmented images, and this is discussed in Sect. 17.4. Section 17.5 discusses using CF outputs in a Bayesian graphical model to improve the recognition performance, and Sect. 17.6 provides a summary.

17.2 Correlation Filter Background

A straightforward measure of the similarity between a probe image $p[m, n]$ and a reference image $r[m, n]$ is the inner product between the two two-dimensional arrays or the inner product between column vectors \mathbf{p} and \mathbf{r} obtained by lexicographically scanning $p[m, n]$ and $r[m, n]$. If this inner product (after appropriate normalization) is large, it indicates a high degree of similarity, whereas if this value is small, it suggests that the probe and the reference images are not a good match. However, in practice, the probe image $p[m, n]$ may be a shifted version of the reference $r[m, n]$ necessitating that the inner product be evaluated for all possible shifts between the two images. This leads to the following correlation output $c[x, y]$ which measures the similarity between $r[m - x, n - y]$ and $p[m, n]$ for all possible shifts x and y .

$$c[x, y] = \sum_m \sum_n p[m, n] r[m - x, n - y] \quad (17.1)$$

It can be shown that if the probe image $p[m, n]$ is exactly equal to $r[m - x_0, n - y_0]$, then $c[x, y]$ will have its highest value (called *correlation peak*) at $x = x_0, y = y_0$. Thus, the relative shift between the probe image and the reference image can be determined by locating the peak in the correlation output. If the probe image is from a different class than the class of the reference image, then the resulting $c[x, y]$ will not exhibit a dominant correlation peak, indicating that the two images may come from different classes. The cross-correlation operation above is more efficiently implemented in the frequency domain, i.e., $c[x, y]$ is the 2-D inverse discrete Fourier Transform (DFT) of the product of $P[u, v]R[u, v]$ where $P[u, v]$ and $R[u, v]$ are the 2-D DFTs of $p[m, n]$ and $r[m, n]$, respectively. Here, u and v denote the spatial frequencies corresponding to m and n , respectively, and we use upper case italics to denote frequency-domain functions and lowercase italics to denote image domain functions. The 2-D DFTs are efficiently implemented using fast Fourier transform

(FFT) algorithm. This implementation of the cross-correlation operation via the frequency domain is the main reason for this operation to be termed *correlation filtering*.

In practice, the probe image $p[m,n]$ will differ from the reference image $r[m,n]$ in multiple ways including additive noise, shifts, rotations, scale differences, illumination changes, partial occlusions, and other differences. All these differences cause the correlation peaks to become lower and broader, making it harder to determine the similarity between the two images and to determine the location of the reference image in the probe scene. One method developed to deal with the appearance changes due to such distortions is to design a correlation filter $H[u,v]$ (or equivalently the template $h[m,n]$ in the image domain) that exhibits the following properties of the correlation output $c[x,y]$:

- Correlation output $c[x,y]$ should exhibit large and consistent values at the center in response to centered training images from the authentic class.
- Correlation output $c[x,y]$ should exhibit small values throughout the correlation plane in response to training images from the impostor class.
- Correlation output $c[x,y]$ should take on small values in the rest of the correlation plane so that the controlled values stand out (i.e., yield peaks) in response to centered images from the authentic class.
- Correlation output $c[x,y]$ should exhibit low sensitivity to noise in the probe or query image input.

Over the years, many correlation filter designs have been developed to achieve the above desired properties. For reasons of space, we will focus on the design of one type of correlation filter known as *optimal trade-off synthetic discriminant function* (OTSDF) filter [8]. More details about CF designs can be found elsewhere [5].

Let $r_1[m,n], r_2[m,n], \dots, r_L[m,n]$ denote L training images (assumed to be centered) from the authentic class with each image of size $M \times N$. The 2-D DFTs of these images are lexicographically scanned to yield column vectors $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_L$, each containing $d = MN$ elements. Matrix \mathbf{R} of size $d \times L$ contains $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_L$ as its column vectors. Similarly, let \mathbf{h} denote a d -dimensional column vector, containing as its elements the scanned version of the correlation filter $H[u,v]$. The goal is to find CF vector \mathbf{h} to meet the above objectives.

The OTSDF design requires that the filter vector \mathbf{h} satisfy the following inner product constraints:

$$\mathbf{h}^T \mathbf{r}_i = c_i, \quad i = 1, 2, \dots, L \quad (17.2)$$

Typically c_i is set to 1 for all training images from the authentic class and to 0 for training images from the other classes. It is expected that the resulting CF will yield correlation output values (at the origin) close to 1 for centered non-training (i.e., test) images from the authentic class and values close to 0 for other images. The linear constraints in Eq. (17.2) are under-determined in that there are L constraints and d

unknowns, where d (the number of pixels in the training images) is usually much larger than L (the number of training images). Thus, there are infinite solutions to Eq. (17.2). Among these infinite solutions, the OTSDF design tries to find the vector \mathbf{h} that leads to sharp correlation peaks (sharp peaks make it easy to detect and locate the objects) for authentic class images and that have good tolerance to input noise.

One way to obtain sharp correlation peaks is to force the correlation output values to be small while constraining the value of the peak (i.e., the correlation output at origin given that the input image is a centered training image from the authentic class) to be 1. Mahalanobis et al. [9] tackled this problem by minimizing the average correlation energy (ACE) defined as follows:

$$\begin{aligned} \text{ACE} &= \frac{1}{L} \sum_{i=1}^L \sum_x \sum_y |c_i(x,y)|^2 \propto \frac{1}{L} \sum_{i=1}^L \sum_u \sum_v |C_i(u,v)|^2 \\ &= \frac{1}{L} \sum_{i=1}^L \sum_u \sum_v |H(u,v) R_i(u,v)|^2 \\ &= \sum_u \sum_v |H(u,v)|^2 \left\{ \frac{1}{L} \sum_{i=1}^L |R_i(u,v)|^2 \right\} \\ &= \sum_u \sum_v |H(u,v)|^2 D(u,v) \end{aligned}$$

$$\text{where } D(u,v) = \frac{1}{L} \sum_{i=1}^L |R_i(u,v)|^2 \quad (17.3)$$

The ACE term in Eq. (17.3) can be more compactly expressed as $\mathbf{h}^T \mathbf{D} \mathbf{h}$ where \mathbf{h} is a d -dimensional column vector containing the filter $H(u,v)$ and \mathbf{D} is a $d \times d$ diagonal matrix whose diagonal entries are $D(u,v)$. Minimizing ACE in Eq. (17.3) subject to the peak constraints in Eq. (17.2) leads to the following CF known as the *minimum average correlation energy* (MACE) filter [9]:

$$\mathbf{h} = \mathbf{D}^{-1} \mathbf{R} (\mathbf{R}^+ \mathbf{D}^{-1} \mathbf{R})^{-1} \mathbf{c} \quad (17.4)$$

where superscript $+$ denotes conjugate transpose and where the L -dimensional column vector \mathbf{c} is defined as $\mathbf{c} = [c_1 \ c_2 \ \dots \ c_L]^T$. Since \mathbf{D} is diagonal, forming \mathbf{D}^{-1} is easy and the computation of $(\mathbf{R}^+ \mathbf{D}^{-1} \mathbf{R})^{-1}$ involves an $L \times L$ matrix.

While the MACE filter in Eq. (17.4) produces sharp correlation peaks in response to training images from the authentic class, it exhibits high sensitivity to input noise and other appearance variations. This is mainly because MACE filters tend to be high-frequency-emphasis filters (in order to produce sharp correlation peaks) which end up amplifying noise. One way to reduce the noise sensitivity of a CF is to reduce the output noise variance (ONV) defined as follows. If the input

image is corrupted by a wide-sense stationary noise $w[m, n]$ with power spectral density $P_w[u, v]$, then the variance of the correlation output is given as

$$\begin{aligned} \text{ONV} &= \text{var}\{c(x, y)\} = \sum_u \sum_v P_c(u, v) \\ &= \sum_u \sum_v P_w(u, v) |H(u, v)|^2 \end{aligned} \quad (17.5)$$

where $P_c[u, v]$ is the power spectral density of the noise in the correlation output. Once again, the ONV can be expressed as $\mathbf{h}^T \mathbf{P} \mathbf{h}$ where \mathbf{P} is a $d \times d$ diagonal matrix whose diagonal entries are $P_w[u, v]$. Minimizing ONV in Eq. (17.5) subject to the constraints in Eq. (17.2) leads to the following CF known as the minimum variance synthetic discriminant function (MVSDF) filter [10]:

$$\mathbf{h} = \mathbf{P}^{-1} \mathbf{R} (\mathbf{R} + \mathbf{P}^{-1} \mathbf{R})^{-1} \mathbf{c} \quad (17.6)$$

While the MACE filter in Eq. (17.4) exhibits sharp correlation peaks and high noise sensitivity, the MVSDF filter in Eq. (17.6) typically exhibits broad correlation peaks and good noise tolerance. Refregier [8] introduced the following optimal trade-off synthetic discriminant function (OTSDF) filter formulation that trades off peak sharpness for noise tolerance:

$$\mathbf{h} = \mathbf{T}^{-1} \mathbf{R} (\mathbf{R} + \mathbf{T}^{-1} \mathbf{R})^{-1} \mathbf{c} \quad (17.7)$$

where $\mathbf{T} = \alpha \mathbf{D} + \sqrt{1 - \alpha^2} \mathbf{P}$ with $0 \leq \alpha \leq 1$ is a scalar that controls the trade-off between peak sharpness and noise tolerance. For $\alpha = 0$, the OTSDF filter is same as the MVSDF filter, and for $\alpha = 1$, the OTSDF filter is same as the MACE filter. For other values of α , we achieve a compromise between the two extremes. In practice, we find that α values close to but not equal to 1 (e.g., 0.999) usually produce the best results.

Over the past two decades, CF designs have advanced in many ways [5]. Some examples of these advances are the following:

- Relaxing the hard constraints in Eq. (17.2) by maximizing the average correlation output value (at the origin) rather than requiring that the correlation output take on specific values
- Designing the CF based on the entire correlation output rather than just the value at the origin
- Applying nonlinear operations to input images in the form of point nonlinearities and in the form of quadratic correlation filters
- Combining the shift-invariance properties of the CFs with the good generalization properties of support vector machine (SVM) classifiers in the form of maximum margin correlation filters (MMCFs)

In the next section, we discuss how CFs can be used for iris segmentation.

17.3 Iris Segmentation

Before two iris patterns can be compared, they need to be segmented from the rest of the image. Most iris segmentation approaches rely on the fact that, in gray-scale images, the pupil (i.e., interior) to the iris is usually darker than the iris region and the sclera (on the outside) is brighter than the iris. So, iris boundaries can be detected by looking for regions with large gradient magnitudes (e.g., from pupil to iris and from iris to sclera) as was proposed originally in using integrodifferential operators for iris boundary detection. Another useful feature of iris boundaries is that they may be nearly circular, suggesting the use of circular Hough transforms to identify iris boundaries. More recently, improved iris segmentation results have been obtained using active contour techniques. In this section, we discuss a cross-correlation-based method for iris segmentation.

As discussed above, one way to locate the iris boundaries is to determine regions of high radial gradients of the circular Hough transform [11]. For eye image $E[m, n]$, the circular Hough transform is defined as

$$Z[m, n, r] = \sum_{\theta \in I_\theta} E[m + r \cos \theta, n + r \sin \theta] \quad (17.8)$$

where I_θ is a sub-interval of $\theta \in [0^\circ, 360^\circ]$. In practice, we do not integrate across the entire circle because the upper and lower regions are unreliable due to eyelid interference. Instead, we integrate over symmetric left and right lateral regions extending from 45° to 150° from the top of the circle. This exclusion of the eyelid region leads to a more robust edge detector. Since the iris may not be centered in the image, we need to compute $Z[m, n, r]$ for every possible triplet $[m, n, r]$. This can be computationally prohibitive, as its complexity is order $O(M^4)$ for an $M \times M$ image. As an example, consider the following naive approach to computing the discrete circular Hough transform: for every possible center location $[m, n]$, the polar transform (requiring $O(M^2)$ operations) of the image is obtained using $[m, n]$ as the origin, followed by summing over the angle within the specified angular interval I_θ . This requires M^2 repetitions of the polar transform leading to $O(M^4)$ complexity. This complexity can be reduced by using cross-correlation operation to produce an approximation to the Hough transform, as described below.

For a given triplet $[m_0, n_0, r_0]$, the value of the discrete Hough transform is a sum of pixel intensities that fall along the circular arc of radius r_0 centered at $[m_0, n_0]$. This summation can be approximated as the inner product of the image $E[m, n]$ with a binary template $C_r[m, n]$ that equals 1 along this contour and 0 everywhere else, as shown in Fig. 17.2. The inner product of the eye image with the binary template in Fig. 17.2 yields one value $Z[m_0, n_0, r_0]$.

To determine $Z[m, n, r_0]$ for all $[m, n]$ values, we must compute inner products with all shifted versions of the binary contour, which is equivalent to spatial cross-correlation. So one 2-D cross section of the discrete Hough transform can be obtained as

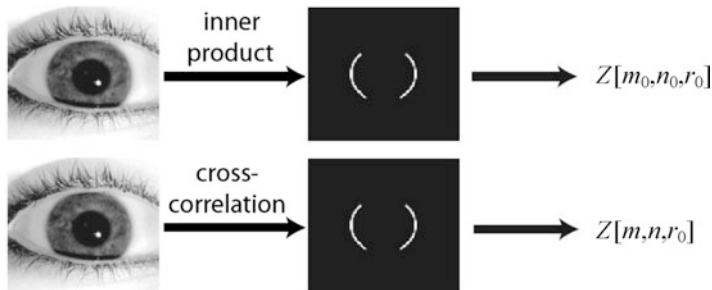


Fig. 17.2 Circular Hough transform approximations using contour filters. An inner product yields a single value, and a cross-correlation yields a 2-D cross section of values




Fig. 17.3 Cross-correlation-based circular boundary detection

$$Z[m, n, r_0] = E[m, n] \otimes C_{r_0}[m, n] \tag{17.9}$$

where cross-correlation is denoted by the \otimes symbol. The cross-correlation in Eq. (17.9) can be efficiently implemented using fast Fourier transform (FFT) algorithm. The computational complexity of the 2-D FFT is $O(M^2 \log_2 M)$. It should be noted that this technique only produces an approximation to the actual Hough transform values, since the use of discrete Fourier transform results in a *circular* cross-correlation. This means that the contour of integration will wrap around when shifted beyond the edge of the image. Fortunately, this does not present a problem in practice. We can either assume that the iris is completely contained within the eye image, or we can zero-pad the eye image if we expect part of the iris to be cut off.

To determine the iris boundaries, the starting point is to build a coarse approximation to the discrete Hough transform. Using a coarse approximation reduces computation time but allows for a reasonable initial estimate of the boundary locations. First, the input eye image is down-sampled to a low-resolution image, e.g., 100 by 100 pixels. Then it is passed through a correlation filter bank where each filter is a binary contour filter of different radius. The output from each filter yields the plane of all Hough transform values at a fixed radius, so concatenating all outputs together gives us the entire discrete Hough space. This is illustrated in Fig. 17.3.

Because the image dimension is low, we need only use 50–100 radial values. The resulting Hough transform is fairly coarse but is usually sufficient for detecting the approximate boundary locations. Each contour filter’s frequency response is computed before hand and stored. Also, the eye image is converted to the frequency domain once, at the beginning of the filtering process. As a result, we only have to compute one inverse FFT for each contour filter. The entire computation has complexity $O(KL^2\log_2L)$ when using K contour filters applied in the low-dimensional space of size $L \times L$. This is significantly better than the $O(M^4)$ complexity of straightforward computation of the circular Hough transform.

The approximate Hough transform is multiplied by $1/r$ to normalize for the circumference. Because we are detecting circular edges, we need the radial gradients of the normalized Hough space. A smoothed difference operator is applied along the radial direction to get these radial gradients. After this step, we have a 3-D set of values which indicate the presence of a circular iris boundary at a range of possible positions. First, a search is conducted for the inner iris boundary, which, depending upon the darkness of the pupil, typically produces the highest radial gradient. We locate the maxima of the gradient, with the minor constraint that the boundary cannot be very near the edge of the image since it is the inner boundary. Once established, the location of the first boundary places a prior on the location of the second boundary. This affects the search in two ways: (1) any potential boundary locations which do not completely surround the inner boundary by some minimal margin are ruled out, and (2) the gradient values are weighted by a Gaussian function, such that the Gaussian is centered on the inner boundary center. The second condition allows for slightly non-concentric boundaries but makes the detection of extremely non-concentric boundaries unlikely. Then the maximum of the weighted gradients is associated with the outer boundary.

At this point in the algorithm, we have coarse estimates for the location of both boundaries. In order to fine-tune these estimates, the detection process is repeated at a higher resolution. The higher resolution contour filters are more computationally expensive to apply, but we do not have to apply the entire filter bank. Instead, we only apply the few filters which have radii in the immediate neighborhood of the coarse estimates. This allows us to refine our estimates without adding significant computation. After deriving the final boundary estimates, the iris pattern is “unwrapped” into normalized pseudo-polar coordinates. We note that the objectives of cross-correlation-based segmentation algorithm are consistent with other iris segmentation algorithms that estimate the iris boundaries by finding the regions of high radial gradients. It is the use of correlation filter bank to efficiently obtain the Hough transform approximation, which differs from other implementations.

We tested the iris segmentation algorithm on an iris image database collected at CMU [12]. Some examples of iris images from CMU database are shown in Fig. 17.4. CMU iris image database contains high-resolution ($950 \times 1,419$) iris images acquired under visible wavelength illumination. This database contains 2,390 images from 101 different eyes with 20–25 images per class. Although the



Fig. 17.4 Sample images from CMU database




Fig. 17.5 Example cross-correlation-based iris segmentation results (*left*) and resulting unwrapped iris images (*right*)

CMU database images have high resolution, they tend to be more difficult for pattern matching because of greater intra-class variation (especially with regard to focus and occlusion). As can be seen from Fig. 17.4, the upper eyelid can cause partial occlusion in the iris image.

In Fig. 17.5, we show example segmentation results from the cross-correlation-based method as well as the corresponding unwrapped iris images. The white regions in the top-middle portion of the unwrapped images are due to the occlusion from the eyelids. The results of automatic segmentation were compared against manual segmentation, and it was observed that nearly 99% of the images were properly segmented by the cross-correlation-based algorithm. The few segmentation errors observed were mostly a result of heavy eyelid occlusion obscuring the iris boundaries.

This discussion about iris segmentation would not be complete without mentioning the real-world challenges in segmenting the iris images. Iris segmentation

is degraded by impairments such as non-frontal gaze, specular reflections, and occlusion due to eyelids and eyelashes.

17.4 Iris Matching

Cross-correlation is a powerful tool for quantifying the similarity between two images. So it can be used for matching iris images as illustrated in Fig. 17.6. Segmented and unwrapped training iris patterns from one class (i.e., one eye) are used to determine a correlation filter such as the MACE filter (Eq. 17.4) or OTSDF filter (Eq. 17.7). When a query iris image is presented, it is also segmented and unwrapped and then cross-correlated using the designed CF. The resulting correlation output should contain a sharp peak if the query image is from the same class as the training images used for designing the correlation filter, and no such peak if the query image is from an impostor, as depicted in Fig. 17.6.

The sharpness of the correlation peak can be quantified by the peak-to-correlation energy (PCE) defined as a ratio of square of the peak-to-correlation energy. Since the PCE is a ratio, multiplying the input query image by any constant will not affect the PCE, making it invariant to constant illumination changes. If the PCE is above a prespecified threshold, the input image is classified as authentic and otherwise as coming from an impostor. By varying the PCE threshold, one can trade off false accept rate (FAR) for false reject rate (FRR).

We demonstrated [13] that correlation filters can offer excellent iris recognition performance. We investigated the performance of OTSDF correlation filter on the CMU iris database. Three images from each iris class were used as reference images to define the class, and the rest were used for testing. The testing generated similarity scores for a number of authentic and impostor comparisons. We measured the equal error rate (EER), the point at which FRR equals FAR. The EER using correlation filters was 0.61%, whereas the EER using Gabor wavelet-based binary codes was 1.04%.

Since correlation filters are applied to unwrapped iris images, the shift-invariance of CFs corresponds to shifts in the polar domain, i.e., the CF method can handle




Fig. 17.6 Correlation-based matching of unwrapped iris images



Fig. 17.7 Sample ocular images from the Face and Ocular Challenge Set (FOCS) dataset [14]

in-plane rotations in the original image (before mapping to polar coordinates). However, the challenge is that iris regions have to be segmented from their surrounding regions before obtaining the unwrapped versions. One of the advantages of CFs is that they can be applied to the original eye images without any need for segmentation.

In Fig. 17.7, we show two example ocular images from the Face and Ocular Challenge Set (FOCS) [14] that contain the iris region as well as surrounding regions such as eyebrow, some part of nose bridge, skin near the eye, and part of the forehead region. Using these additional regions can improve the recognition rate. In Fig. 17.8, we show the correlation outputs for an authentic ocular image pair and an impostor ocular image pair. It is clear that correlation is stronger for the authentic pair. We also show the correlation output when the filter is same, but the probe image is an iris-occluded version of the authentic ocular image, and we can see that the correlation peak is still visible even though the iris is occluded.

17.5 Bayesian Graphical Models for Iris Recognition

One of the reasons for the degradation of match score between two iris images (acquired at different times) from the same eye is that the two images can exhibit nonlinear deformations, e.g., different regions of the unwrapped iris images may move differently, as illustrated in Fig. 17.9. Such nonlinear deformation can also be caused by slight differences in the segmentation boundaries produced for the two images. When segmentation boundaries differ, corresponding regions of the two unwrapped images may move by different amounts because the mapping from Cartesian coordinates to polar coordinates very much depends on where the inner and outer boundaries are.

Another challenge in matching two iris images from the same eye is that one may exhibit more occlusions caused by eyelid and the other may exhibit less occlusion, affecting the match score. If such occluded regions can be excluded or weighted less in determining a match score, then that should lead to a more robust match

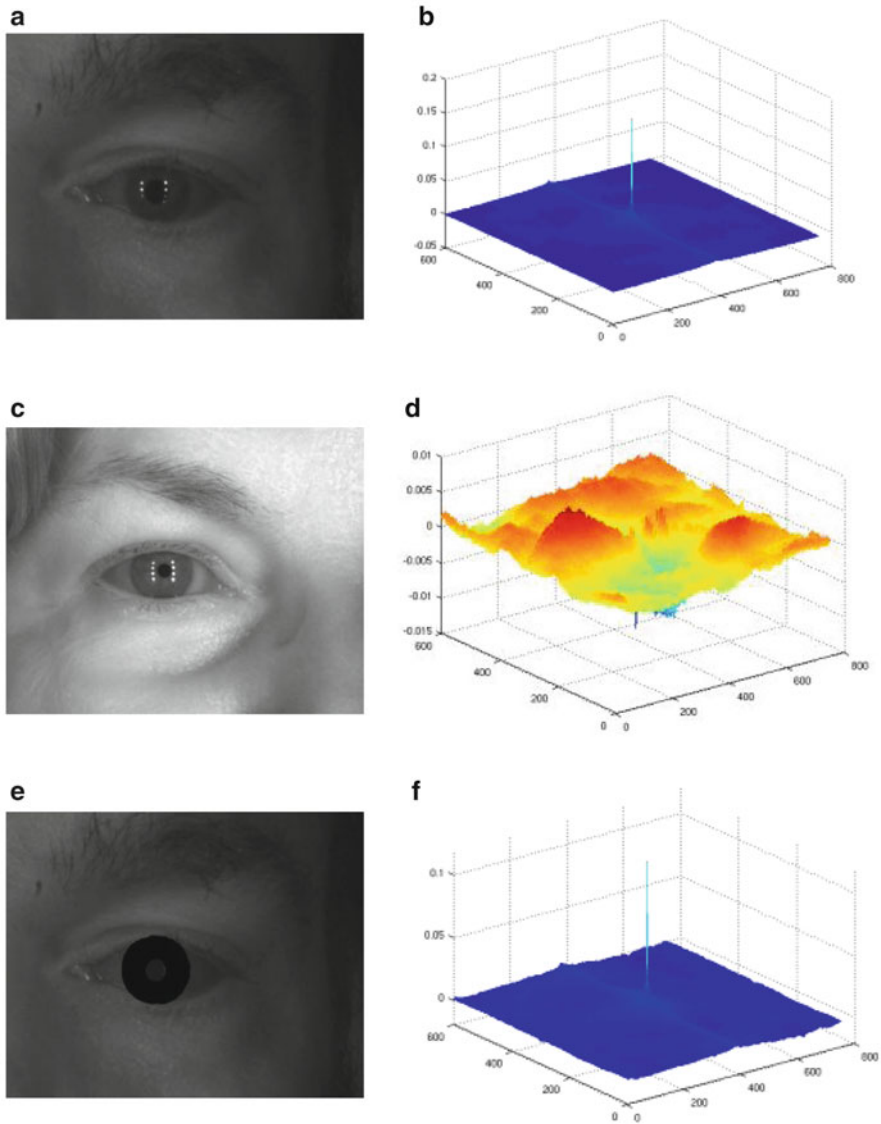


Fig. 17.8 (a) Authentic ocular image, (b) resulting correlation output, (c) an impostor ocular image, (d) resulting correlation output, (e) iris-occluded authentic ocular image, and (f) resulting correlation output

score. Toward this goal, Thornton et al. [15] proposed the use of Bayesian graphical models for improved iris matching.

The main idea of Bayesian graphical models can be summarized as follows. The two iris images being compared are divided into nonoverlapping patches as shown in




Fig. 17.9 Close-ups of segmented patterns from same eye (landmark points illustrate relative deformation)




Fig. 17.10 Hidden states of the model: Sample iris plane partition (*top*), deformation vector field (*center*), and binary occlusion field (*bottom*)

Fig. 17.10 (top) where the unwrapped iris image is divided into 36 patches. Different patches from the probe image may be shifted by different amounts compared to the corresponding patches of the template, as shown in Fig. 17.10 (middle) by white arrows. The length of the arrow indicates the magnitude of the shift, and the direction of the arrow indicates the direction of the shift. Also, some of the probe image patches may be occluded, as shown by the gray squares in Fig. 17.10 (bottom). To estimate these patch shifts and to estimate whether a patch is occluded or not, the corresponding patches from the two regions are cross-correlated. If both patches are unoccluded and from the same eye, the resulting correlation peak should be large,




Fig. 17.11 Graphical model structure, for a 3 by 12 iris plane partition. Shaded nodes represent observed variables and non-shaded nodes represent hidden variables

and the location of the correlation peak should indicate the relative shift between the corresponding patches from the two images. Thus, the cross-correlations between the patches from the two images provide clues about patch shifts between the two images and occlusions.

The graphical models corresponding to the patch structure in Fig. 17.10 are shown in Fig. 17.11. Shaded nodes O_1, O_2, \dots, O_{36} denote the observations (e.g., PCE values and peak locations from patches) and represent evidence about both the similarity between the template and query iris patterns and the presence of eyelids across the iris plane, whereas hidden variables d_1, d_2, \dots, d_{36} indicate the shift of probe image patches from corresponding template image patches and binary-valued hidden variables $\lambda_1, \lambda_2, \dots, \lambda_{36}$ indicate whether a patch is occluded (i.e., $\lambda = 1$) or not occluded (i.e., $\lambda = 0$).

Once the hidden variables (indicated by set H) and the observed variables (indicated by set O) are identified, the objective is to learn a joint probability distribution over these variables so that we may perform inference on the hidden states. In order to do this, we make some assumptions about the dependencies between variables. If we assume a fully connected model (i.e., every variable is directly dependent upon every other variable), learning and inference would be completely intractable. Therefore, we simplify the model by assuming direct dependence only between variables which have an intuitive or empirical statistical connection. This is a common practice in the field of probabilistic graphical models [16], which provides a general framework for working with complicated joint distributions. In keeping with conventional graphical model notation, a variable is represented visually as a node, and a direct dependence between two variables is represented as an edge connecting two nodes.




Fig. 17.12 Contrasting deformation field examples. *Top*: A “reasonable” field with higher likelihood in our model. *Bottom*: An “unreasonable” field with lower likelihood

We start by considering the relationship between the hidden deformation variables. Figure 17.12 shows two examples of deformation vector fields which might align a probe image to a template image. The alignment at the top is more “reasonable” than the alignment at the bottom, by which we mean that the top vector field is more likely to approximate the effect of real iris image movement. Part of what makes the second vector field more unlikely, besides the increased magnitude of the vectors, is the fact that many vectors from adjacent regions have conflicting directions. We would expect adjacent iris regions to exhibit similar motion, and we want a model capable of learning this tendency from iris data. Therefore, we allow for direct dependence between each vector and its spatial neighbors in the vector field. We formulate this dependence scheme as a Markov random field (MRF), which is an undirected graph structure on the variable set [17]. We arrange our variables in an MRF framework because it makes learning and inference more tractable. Specifically, we form a 2-D lattice MRF on the deformation variables, in which each node is connected to its neighbors on the lattice (illustrated in Fig. 17.11).

The graphical model parameters are learned using training images via expectation-maximization (EM) algorithm [18]. Use of the graphical model for an iris pattern comparison is a two-step process: (1) we infer distributions on the hidden variables of the model, and (2) we use this information to compute a match score between template and query. For space reasons, we do not discuss the details, but they can be found elsewhere [12].

The Bayesian graphical model approach was tested using CMU iris database containing high-resolution ($950 \times 1,419$) iris images acquired under visible wavelength illumination. This database contains 2,390 images from 101 different eyes with 20–25 images per class. Table 17.1 shows the false alarm rates (FAR) for

Table 17.1 False accept rate (*FAR*) for different false reject rates (*FRRS*)

Algorithm	FRR = 1%	FRR = 0.1%	FRR = 0.01%
Baseline	0.23%	1.16%	1.85%
BGM	0	0	0.46%

three different false reject rate (FRR) levels for the Bayesian graphical model (BGM) approach and the baseline algorithm. The baseline algorithm is the standard approach of iris segmentation followed by a binary code derived from the phases of Gabor wavelet inner products. As can be seen from Table 17.1, the Bayesian graphical model-based matching algorithm offers improved recognition performance.

17.6 Summary

Correlation filters (CFs) have long been researched for automatic target recognition applications where the targets can appear in different orientations, scales, locations, and with occlusions and obscurations. Advanced CF designs have been developed to deal with such image impairments. Also, CFs have built-in advantages such as shift-invariance and graceful degradation. One of the challenges in real-world iris recognition is that the iris images exhibit impairments such as non-frontal gaze, occlusions due to eyelids and eyelashes, specular reflections, and nonuniform illuminations. Correlation filters can be beneficial in dealing with such impairments. In this chapter, we discussed how correlation filters (CFs) can play an important role in iris recognition.

We discussed how CFs can provide an alternative method for iris segmentation and how they can be used for matching both unwrapped iris images (in polar coordinates) and iris images in the original Cartesian coordinates. The use of CFs can be extended to iris image patches so that the patch cross-correlations can provide information to a Bayesian graphical model (BGM) that outputs match scores that are adjusted for impairments such as nonlinear deformations and occlusions. We showed that BGM offers improved iris recognition performance compared to a baseline binary code-based matching.

References

1. Wildes, R.: Chapter 3: Iris recognition. In: Wayman, J., Jain, A., Maltoni, D., Maio, D. (eds.) *Biometric Systems*, pp. 63–95. Springer, Berlin (2005)
2. Daugman, J.: Probing the uniqueness and randomness of IrisCodes: results from 200 billion iris pair comparisons. *Proc. IEEE* **94**(11), 1927–1935 (2006)
3. Daugman, J.: High confidence visual recognition of persons by a test of statistical independence. *IEEE Trans. Pattern Anal. Mach. Intell.* **15**(11), 1148–1161 (1993)

4. Vijaya Kumar, B.V.K., Xie, C., Thornton, J.: Iris verification using correlation filters. In: Proceedings of 4th International Conference on Audio- and Video-Based Biometric Person Authentication (AVBPA). LCNS 2688, pp. 697–705. Springer, Berlin/Heidelberg (2003)
5. Vijaya Kumar, B.V.K., Mahalanobis, A., Juday, R.: Correlation Pattern Recognition. Cambridge University Press, Cambridge (2005)
6. Mahalanobis, A., Ortiz, L., Vijaya Kumar, B.V.K.: Performance of the MACH/DCCF algorithms on the 10-class public release MSTAR data set. Proc. SPIE **3721**, 285–291 (1999)
7. Vijaya Kumar, B.V.K., Savvides, M., Xie, C.: Correlation pattern recognition for face recognition. Proc. IEEE **94**(11), 1963–1976 (2006)
8. Refregier, P.: Filter design for optical pattern recognition: multicriteria optimization approach. Opt. Lett. **15**(15), 854–856 (1990)
9. Mahalanobis, A., Vijaya Kumar, B.V.K., Casasent, D.: Minimum average correlation energy filters. Appl. Opt. **26**, 3630–3633 (1987)
10. Vijaya Kumar, B.V.K.: Minimum variance synthetic discriminant functions. JOSA-A **3**, 1579–1584 (1986)
11. Duda, R.D., Hart, P.E.: Use of the Hough transform to detect lines and curves in pictures. Commun. ACM **15**, 11–15 (1972)
12. Thornton, J.: Iris pattern matching: a probabilistic model based on discriminative cues. Ph.D. dissertation, Carnegie Mellon University (2007)
13. Thornton, J., Savvides, M., Vijaya Kumar, B.V.K.: Robust Iris recognition using advanced correlation techniques. In: Proceedings of International Conference on Image Analysis and Recognition (ICIAR), Image Analysis and Recognition, Lecture Notes in Computer Science #3656, pp. 1098–1105. Springer Berlin/Heidelberg (2005)
14. <http://www.nist.gov/itl/iad/ig/focs.cfm>, Face and ocular challenge series (FOCS).
15. Thornton, J., Savvides, M., Vijaya Kumar, B.V.K.: A unified Bayesian approach to deformed pattern matching of iris images. IEEE Trans. Pattern Anal. Mach. Intell. **29**, 596–606 (2007)
16. Frey, B.J.: Graphical Models for Machine Learning and Digital Communication. MIT Press, Cambridge, MA (1998)
17. Geman, S., Geman, D.: Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. IEEE Trans. Pattern Ana. Mach. Intell. **6**(6), 721–741 (1984)
18. Jordan, M.I.: Learning in Graphical Models. MIT Press, Cambridge, MA (1999)

Chapter 18

Iris Spoofing: Reverse Engineering the Daugman Feature Encoding Scheme

Shreyas Venugopalan and Marios Savvides

Abstract Biometric systems based on iridal patterns have shown very high accuracies in verifying an individual's identity due to the uniqueness of the iris pattern across individuals. For identity verification purposes, only the iris bit code template of an individual need be stored. In this chapter, we explore methods to generate synthetic iris textures corresponding to a given person for the purpose of bypassing an iris-based security system using these iris templates. We present analysis to prove that when this “spoof” texture is presented to an iris recognition system; it will elicit a similar response from the system as that due to the genuine iris texture to which the spoof corresponds. We embed this spoof texture within the iris of an imposter to achieve this end. Systems using filter-based feature extraction systems – such as Daugman style systems – may be bypassed using this technique. We assume knowledge of solely the feature extraction mechanism of the iris matching scheme and, as mentioned, the iris bit code template of the person whose iris is to be spoofed. We present a complete investigation into how one can get by an iris recognition system using this approach, by generating various “natural”-looking irises and hope to use this knowledge to incorporate several countermeasures into the feature extraction scheme of an iris recognition module.

18.1 Introduction

Due to uniqueness of iris patterns, recognition based on this biometric modality has gained a lot of popularity over recent years. The bit code template generated for this purpose, using iridal features, is generally considered to be unidentifiable data, that is, given such a template, one cannot determine the identity of the person

S. Venugopalan • M. Savvides (✉)
Cylab Biometrics Center, Department of Electrical and Computer Engineering,
Carnegie Mellon University, Pittsburgh, PA 15213, USA
e-mail: svenugop@andrew.cmu.edu; marioss@andrew.cmu.edu

to whom it belongs. Our work, outlined in this chapter, uses this template to design synthetic/spoof irises similar to the original. These spoof irises (“payloads”) can then be embedded within another person’s (host) iris pattern for the purpose of bypassing an iris recognition system. Recent works by Shah and Ross [1], Makthal and Ross [2], and Zuo et al. [3] also describe ideas of embedding iris texture within a background template, using Markov models and geometric/anatomical models. However, the aim in these works is the generation of synthetic irises to augment existing iris databases. To the extent of our knowledge, there has been no work similar to that presented in this chapter – spoofing a person’s iris pattern using only his/her iris bit code template for the purpose of bypassing a recognition system. We feel that works of this nature are necessary to determine weak points in current state-of-the-art systems deployed at sensitive facilities and to address them effectively.

Several papers discuss the generation of synthetic iris textures for realistic imagery, such as Lefohn et al. [4], or simply for the purpose of generating a large dataset of iris patterns to test the performance of novel iris recognition algorithms such as by Cui et al. [5], Shah and Ross [1], and Wecker et al. [6]. Lefohn et al. [4] stacks semitransparent synthetic textures to generate realistic-looking iris textures. The composite result of stacking several such artificial layers, based on cues used by ophthalmologists, provides a sense of depth to the iris. Cui et al. [5] were the pioneers in synthetic iris image generation. They modeled a PCA subspace using a small set of patches from a set of iris images. Similar iris images from the same class were then generated using the PCA components. Makthal and Ross [2] use Markov random field (MRF) models to describe the stochastic properties governing the intensity values of pixels in a specific neighborhood of the iris. This model is then used to create an iris-like structure from a randomly initialized starting texture. Wecker et al. [6] use reverse subdivision, a multiresolution technique to generate synthetic irises. The information contained within an iris texture at different resolution levels is combined to generate a new unique texture. Zuo et al. [3, 7] model the iris based on an extensive study of anatomical structures within the iris, their functions, and a study of iridology. They use information from these different studies to address the question of visual realism of a synthetic iris image.

As noted previously, all of these works aim to generate varying degrees of visually realistic iris textures in order to augment existing iris databases. This is to carry out iris recognition experiments on a larger database. There has not been any work that investigates spoofing iris patterns with the aim of “bypassing” existing iris recognition systems. Adler [8, 9] has done similar work in the domain of face recognition. He showed that using solely the match score values, one can recreate a sample image from a face recognition template. The regenerated image shows the significant facial features and hence compares with a high score to the original image. In the work described here, the motivation is to be able to generate a number of spoof irises to be used by an imposter, to elicit a response from an iris recognition system similar to the response from a genuine user’s iris pattern. The spoof iris codes generated were matched by traditional iris feature extraction and matching algorithms such those described by Masek [10] and Daugman [11, 12]. We show that every person has an underlying discriminating iris pattern that remains the same for

all irises of that person. We then simply embed this discriminating pattern belonging to a person A into a person B's pattern. Now, imposter B can present the spoof iris to an iris recognition system and can be erroneously identified as person A. Visual realism of the spoof iris can be decided based on the ratio discriminating pattern to host iris texture pattern.

In the next section, we briefly review the stages of a standard iris recognition algorithm pertinent to our method of spoofing. Following this, in Sect. 18.3, we describe our method of obtaining the discriminating iris pattern from the iris code. The effect of embedding this pattern into another person's iris texture is also discussed. We have generated spoof irises for the entire set of left eyes in the NIST ICE 2005 database [13] and have performed the relevant analyses. Results of comparing these spoofed irises to the genuine irises in the database are presented in Sect. 18.4. For a more detailed description of the techniques and the analyses presented in this chapter, the reader is directed to the work by Venugopalan and Savvides [14].

18.2 Brief Review of Iris Feature Extraction and Encoding

Once an eye image has been acquired, the iris region has to be segmented and represented in the polar domain. Following this, feature extraction and encoding are performed. Since our work deals with the iris feature extraction and encoding stages only, we will briefly outline these two stages. A more detailed review of iris recognition can be found in works by Daugman [11, 12] and in the detailed survey of iris recognition literature by Bowyer et al. [15]. Segmentation and masking in this work were performed manually so that they cause no error during our analyses described in the later sections. Iris segmentation and representation in the polar domain are illustrated in Fig. 18.1. The feature extraction used in our work is as described by Daugman [11] which projects the segmented and unwrapped iris texture onto a set of complex-valued 2D Gabor functions. This involves filtering the iris texture using a Gabor filter. A Gabor function $G(x, y)$ can be thought of as a complex harmonic function windowed by a Gaussian function (Eq. 18.1). The filter response is then encoded to obtain the iris bit code template corresponding to the person.

$$G(x, y) = \frac{1}{\sqrt{2\pi}\sigma_x\sigma_y} \exp\left[-\frac{1}{2}\left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2}\right)\right] \exp\left(j\frac{2\pi}{\lambda}(x+y)\right) \quad (18.1)$$

The quantization process is discussed by Daugman in [11]. In this work, we considered only the imaginary part of the complex Gabor function and optimized its parameters. Hence, our encoding requires only 1 bit per pixel which yields a more efficient classifier with similar performance. The final bit map for an iris in

Fig. 18.1 Segmentation of an eye image and representation in polar domain. (a) The *green circle* shows the limbic boundary and the *red circle*, the pupillary boundary. (b) The iris texture “unwrapped” into the polar coordinate system. The corresponding iris mask is also shown. In the mask, the *black region* represents iris texture, and the *white region* indicates non-iris region





Fig. 18.2 (a) A segmented iris image represented in the polar domain, (b) the filter response on filtering, (a) using the imaginary portion of a Gabor function, and (c) the corresponding iris code (see Ref. [11])

our work has a size of 60×360 . A segmented and unwrapped iris, as well as the corresponding encoded template, is shown in Fig. 18.2.

Given two such iris bit code templates I_1 and I_2 and the corresponding iris masks M_1 and M_2 , the match score – a hamming distance value – is computed as shown in Eq. (18.2).

$$H = \frac{\|I_1 \oplus I_2 \cap M_A \cap M_B\|}{\|M_A \cap M_B\|} \quad (18.2)$$

18.3 Algorithm Overview

In this section, we outline our approach to recover the iris pattern of a person, given the iris bit code template. We assume prior knowledge of only the individual's iris code and the Gabor function parameters used by the system. First, we identify those features in an individual's iris texture that discriminates it from another's iris. Then the idea is to embed this *person-specific discriminating (spoof) pattern* into another person's iris texture, in such a way that the iris matching system picks up this embedded pattern.

18.3.1 Generating the Spoof Pattern

As mentioned earlier, feature extraction is a filtering operation which is implemented by convolving the filter with the image. This process in essence measures the presence of the reversed version of a Gabor function, $G(r, \theta)$, at every pixel location (since convolution is correlation of one function with the reversed version of the other). Considering $R(r, \theta)$ to be the filter response, then we define the *person-specific discriminating iris pattern* to be the weighted sum of the shifted, reversed versions of the Gabor function, that is, $G(-r, -\theta)$. $R(r, \theta)$ is used as the weighting function. This pattern is person specific because it is unique to the class to which $R(r, \theta)$ corresponds to. This pattern $E_1(r, \theta)$ can be represented by the following integral:

$$E_1(r, \theta) = \int_{i=1}^{i=m} \int_{j=1}^{j=n} R(i, j) G(r-i, \theta-j) di dj \quad (18.3)$$

where the support of $R(r, \theta)$ is $m \times n$. If the pattern $E_1(r, \theta)$ is filtered using $G(r, \theta)$, the response is given by

$$\begin{aligned} R_E(r, \theta) &= E_1(r, \theta) * G(r, \theta) \\ &= \int_{i=1}^m \int_{j=1}^n \left\{ \int_{k=1}^m \int_{l=1}^n R(k, l) G(-k+i, -l+j) dk dl \right\} \\ &\quad G(-r+i, -\theta+j) di dj \\ &= \int_{k=1}^m \int_{l=1}^n R(k, l) g(r, \theta, k, l) dk dl \end{aligned} \quad (18.4)$$




Fig. 18.3 Iris Image textures and corresponding iris codes

where

$$g(r, \theta, k, l) = \int_{i=1}^m \int_{j=1}^n G(-k+i, -l+j) G(-r+i, -\theta+j) di dj \quad (18.5)$$

Optimum Gabor parameters for the database have a very narrow spread causing $g(r, \theta, k, l)$ to have a very small value, and hence, $R_E(r, \theta) \approx R(r, \theta)$.

From this analysis, we see that if we use $E_1(r, \theta)$ as input to the iris recognition system, we obtain a response similar to when the original pattern is presented to the system. This result forms the basis of our iris spoofing methodology. In the next section, we have experimentally verified that the iris matching scores are close to that when original irises are used. In a similar manner, we represent the iris bit code template (which is a quantized version of $R(r, \theta)$) as described in the previous section) as $R_Q(r, \theta)$. Then, another possible spoof pattern can be generated as

$$E_2(r, \theta) = \int_{i=1}^{i=m} \int_{j=1}^{j=n} R_Q(i, j) G(r-i, \theta-j) di dj \quad (18.6)$$

An example of $E_2(r, \theta)$ and the corresponding original unwrapped iris are shown in Fig. 18.3b, e, respectively. Also, in addition to using the reversed version of the Gabor function, $G(-r, -\theta)$ in order to generate the spoof irises, we can use patches, of the same size as the Gabor function, cropped out from within other irises such that the patch resembles $G(-r, -\theta)$. These patches $P(r, \theta)$ can be discovered by a cross-correlation operation between the iris texture and $G(-r, -\theta)$.

We label this spoof pattern as $E_3(r, \theta)$. An example of this pattern is shown in Fig. 18.3c with the iris code in Fig. 18.3f.




Fig. 18.4 For a spoof iris, this figure shows the effect of changing the values of p_1 and (Eq. 18.9). (a) Is the discriminating pattern to be added to the background texture (b). From (c) to (f) the value of p_1 decreases and that of p_2 increases

$$E_3(r, \theta) = \int_{i=1}^m \int_{j=1}^n R_Q(i, j) P(-r+i, -\theta+j) di dj \quad (18.7)$$

18.3.2 Embedding Spoof Pattern in a Background Texture

The next step in generating our spoof pattern is to embed the generated spoof pattern within a background texture. The background texture will be, for example, the iris of the imposter who wishes to bypass an iris-based access control system. By controlling the ratio of the embedded spoof pattern to the background iris texture, the spoof pattern can be hidden until it is visually indistinct from the background. We define the final texture as a weighted combination as in Eq. (18.8). In this equation, the background texture is represented by $S(r, \theta)$, and the weights $p_1, p_2 \in [0, 1]$ control the ratio of background texture to spoof pattern $E_3(r, \theta)$.

$$E_F(r, \theta) = p_1 S(r, \theta) + p_2 E_3(r, \theta) \quad (18.8)$$

Figure 18.4 shows a few images generated using varying amounts of p_1 and p_2 . The discriminating pattern becomes more evident when the p_1 term in Eq. (18.8) decreases and vice versa. If we convolve $E_F(r, \theta)$ with a Gabor function,

$$\begin{aligned} R_F(r, \theta) &= E_F(r, \theta) * G(r, \theta) \\ &= (p_1 S(r, \theta) + p_2 E_3(r, \theta)) * G(r, \theta) \\ &= p_1 S(r, \theta) * G(r, \theta) + p_2 E_3(r, \theta) * G(r, \theta) \\ &= p_1 R_S(r, \theta) + p_2 R_{E_3}(r, \theta) \end{aligned} \quad (18.9)$$

If $p_2 < p_1$, then the pattern can be kept hidden in the background texture $S(r, \theta)$. From our experiments, we have determined that $p_2 = 0.3$ and $p_1 = 0.7$ are optimum values in this respect. For these values, even during feature extraction (Eq. 18.9), the $R_{E_3}(r, \theta)$ term has more of an influence on the encoding stage than the first term. The results of verification experiments for spoof irises using this process are shown in the next section and Fig. 18.4c–f show the effect of changing the values of p_1 and p_2 on the appearance of the final spoof image.

18.4 Experiments and Results

In this section, we experimentally verify key ideas presented in the previous sections and present final results when using spoof iris textures in place of real textures. As mentioned before, the database we use is the NIST ICE 2005 database [13]. Spoof textures were generated for every image in this set using Eqs. (18.3), (18.6), (18.7), and (18.8). $S(r, \theta)$ in Eq. (18.8) is randomly chosen to be any iris image of a different class. Our verification experiments compare this spoofed dataset with the actual dataset and plotting the receiver operating characteristic (ROC) curve. The ROC curve plots the rate of genuine accepts (verification rate) against the rate of imposter accepts (false accept rate).

First, we compare the verification performance of spoof irises using Eqs. (18.3), (18.6), and (18.7) with each other and to the performance when the original iris textures are compared with each other. The relevant ROC curves are shown in Fig. 18.5.




Fig. 18.5 Receiver Operating Characteristic (ROC) curves generated during verification experiments using original irises and using spoof irises generated using Eqs. (18.3), (18.6), (18.7)




Fig. 18.6 ROC curves generated during verification experiments using different patch sizes in Eq. (18.7)

We see that even for very low values of false acceptance rates ($\sim 0.1\%$) the verification rates using the three spoof generation techniques remain above 90%. Next, in Fig. 18.6, we show the difference in performance when different patch sizes are used in Eq. (18.7). We settled on using a patch size of 13×13 for our purpose because we felt that the resulting spoof patterns satisfied the trade-off between visual realism and a high value of verification rate.

Next, we show a comparison of the ROC curves when using different values of p_1 and p_2 in Eq. (18.8). We decided to use $p_2 = 0.3$ and $p_1 = 0.7$ as the verification rate was high, and with these values, we felt the spoof patterns were not too obvious when hidden within the background texture.

An example of how spoof irises decrease the hamming distance between an imposter and genuine class is shown in Fig. 18.7. In addition, we also compared the verification performance of our spoofed database of irises against the genuine irises using the method proposed by Masek and Kovesi [10]. They use single-dimensional log-Gabor filters for feature extraction and encoding. The ROC curves generated using this method on the original irises and on our spoof irises are shown in Fig. 18.8. Even though the encoding algorithm is different, the verification rates shown by the spoof database remain high.

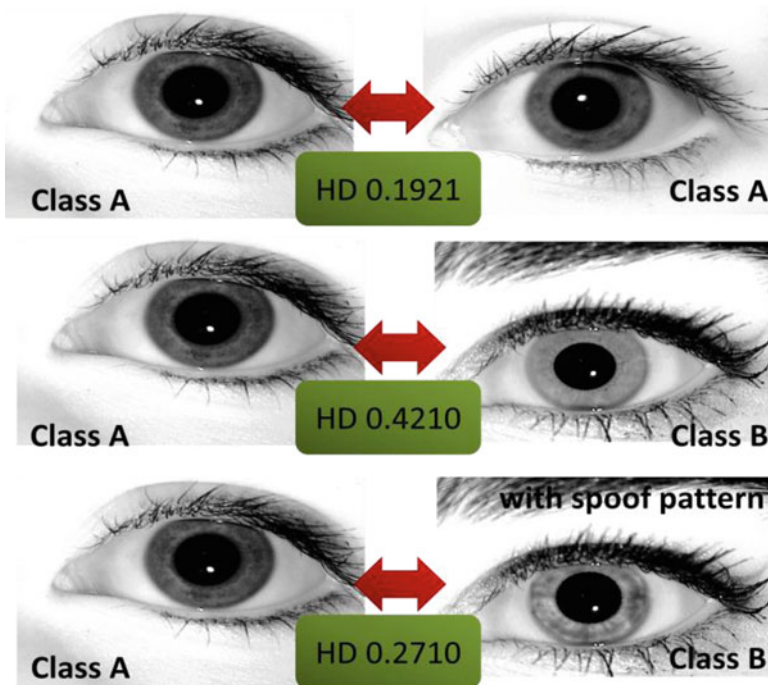


Fig. 18.7 Comparison of hamming distance (HD) between two images from *class A* (first row) and between an image from *class A* and an image from *B* (second row). The *third row* shows a comparison of hamming distance between a *class A* image and *class B* image that has an embedded spoof pattern of *A*. We see that the spoof pattern is able to decrease the calculated hamming distance

18.5 Summary

In this chapter, we explored a method of creating spoof irises to aid an imposter in bypassing an iris recognition system. We showed how such a spoof pattern of a “genuine” person may be generated using knowledge of only his/her iris bit code template and knowledge of the Gabor parameters used by the iris recognition system. The results in the previous section show natural-looking spoof iris images (spoof pattern embedded within another’s iris image) which give similar verification performances as the genuine pattern when presented to the recognition system. Exploring methods such as those presented in this chapter are essential to determine loopholes in state-of-the-art systems that have been deployed at high-security facilities and in order to develop countermeasures for the same.

Acknowledgments This research was supported by CyLab at Carnegie Mellon under grants DAAD19-02-1-0389 and W911NF-09-1-0273 from the Army Research Office.




Fig. 18.8 ROC curves generated using the method proposed by Masek and Kovesi [10]

References

1. Shah, S., Ross, A.: Generating synthetic irises by feature agglomeration. In: Proceedings of International Conference on Image Processing, pp. 317–320. Atlanta, GA, October 2006
2. Makthal, S., Ross, A.: Synthesis of Iris images using Markov Random fields. In: Proceedings of 13th European Signal Processing Conference, Antalya, Turkey, September 2005
3. Zuo, J., Schmid, N., Chen, X.: On generation and analysis of synthetic iris images. *IEEE Trans. Inform. Forensics Secur.* **2**(1), 77–90 (2007)
4. Lefohn, A., Caruso, R., Reinhard, E., Budge, B.: An ocularist’s approach to human iris synthesis. *IEEE Comput. Graph. Appl.* **23**(6), 70–75 (2003)
5. Cui, J., Wang, Y., Huang, J., Tan, T., Sun, Z.: An Iris image synthesis method based on PCA and super-resolution. In: Proceedings of 17th International Conference on Pattern Recognition, vol. 4, pp. 471–474. Cambridge, UK (2004)
6. Wecker, L., Samavati, F., Gavrilova, M.: Iris synthesis: a reverse subdivision application. In: Proceedings of the 3rd International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia, pp. 121–125. (2005). ISBN 1-59593-201-1
7. Zuo, J., Schmid, N., Chen, X.: On performance comparison of real and synthetic iris images. *IEEE Int. Conf. Image Process.* **1**, 305–308 (2006)
8. Adler, A.: Sample images can be independently restored from face recognition templates. *IEEE Can. Conf. Elect. Comput. Eng.* **2**, 1163–1166 (2003)
9. Adler, A.: Images can be regenerated from quantized biometric match score data. *IEEE Can. Conf. Elect. Comput. Eng.* **1**, 469–472 (2004)
10. Masek, L., Kovesi, P.: MATLAB source code for a biometric identification system based on Iris patterns. The School of Computer Science and Software Engineering, The University of Western Australia, 2003

11. Daugman, J.: How iris recognition works. *Proc. Int. Conf. Image Process.* **1**, I-33–I-36 (2002)
12. Daugman, J.: Probing the uniqueness and randomness of IrisCodes: results from 200 billion iris pair comparisons. *Proc. IEEE* **94**(11), 1927–1935 (2006)
13. <http://iris.nist.gov/ice/>
14. Venugopalan, S., Savvides, M.: How to generate spoofed irises from an iris code template. *IEEE Trans. Inform. Forensics Secur.* **6**(2), 385–395 (2011)
15. Bowyer, K., Hollingsworth, K., Flynn, P.: Image understanding for iris biometrics: a survey. *Comput. Vis. Image Underst.* **110**, 281–307 (2008)

Chapter 19

Optics of Iris Imaging Systems

David A. Ackerman

Abstract Iris imaging systems must capture iris images of sufficient quality to populate an enrollment database or to provide probe images that reliably match to existing enrollment images. From whatever distance they are taken, the iris images must therefore resolve information from the iris sufficient for the task of recognition. This chapter reviews concepts of optics and photography needed to specify requirements on the image acquisition components of systems which create iris images for the purpose of recognition. We consider fundamental and practical limitations of components of such systems and consider as examples, iris imaging systems that operate at 0.3 and 3 m on constrained and relatively unconstrained subjects.

19.1 Introduction

The goal of this chapter is to understand the optics of an iris imaging system. An iris imaging system captures a single or dual iris image at a given distance using illumination over a prescribed wavelength band. Design of an imaging system or camera to accomplish this task requires an understanding of the fundamental and practical capabilities and limitations of the camera components such as the lens and image sensor. The resulting iris image must satisfy requirements, for example, on spatial resolution and signal-to-noise ratio, so that it can adequately serve as an enrollment image to be encoded and stored in a database or as an identification

D.A. Ackerman (✉)

Distinguished Member of Technical Staff, Sarnoff Corporation, Princeton, NJ 08543-5300, USA
e-mail: dackerman@sarnoff.com

image to be encoded and compared with database iris codes for verification or identification. We will refer to the overall system that combines an iris imager and a means to identify an iris image as an iris recognition system.

We start by listing assumptions about a typical iris that we intend to image. An iris, the colored portion of an eye surrounding the pupil, is about 11 ± 1.5 mm in diameter [1] and sits behind the partially reflective cornea. We note that the diffuse reflectivity of an iris, referred to as its albedo, is dependent on illumination wavelength and is typically low, around 10%, in the near-infrared (NIR) band. In the NIR band of 700–900 nm throughout which most iris imagery is acquired, the spatial variation in iris reflectivity which gives each iris its unique appearance results from subtle shadings of a few percent in contrast over a range of length scales from microns to millimeters.

The level of fidelity of an iris image needed to distinguish it from other irises is a topic of great interest to manufacturers and users of iris recognition systems. In this chapter, we will not recommend specific image requirements or threshold values for image metrics. Requirements on iris images are listed in ISO/IEC standards and are, at the time of this writing, undergoing careful scrutiny and revision [2]. We will, however, describe a framework for understanding the connection between components of an iris imaging system and measurable attributes of the images expected from such a system. In Sect. 19.4, we will briefly discuss the relationship of an imaging system and the algorithm used to analyze its images, that is, to segment or isolate irises from the larger image, to encode them, and to match the resulting numerical templates to those of other irises.

Finally, we will list some attributes of typical iris imaging systems that form the recipe for most of the successful instruments that are available today. Currently, available iris cameras that operate at approximately 30 cm or less from the subject typically resolve at least 200 pixels across an iris diameter. Those few systems operating at longer distances do not all utilize this level of spatial resolution. Most iris imaging systems use near-infrared illumination between 700 and 900 nm in order to minimize absorption by pigmentation on the front surface of the iris but enable the use of silicon imagers with reasonable sensitivity to image the structural component of the iris that becomes visible in NIR behind the front surface [3]. All iris imaging systems conform to accepted eye-safety standards [4] which limit the irradiance and radiance of near-infrared illumination upon subjects' eyes, whether continuous or pulsed.

While most iris imaging systems share common elements, each system is optimized for particular use conditions according to constraints on size, weight, power, operating distance, ambient conditions, and the expected behavior of subjects. In the next section, we will discuss some basic photographic concepts that will allow us to connect intended use conditions to requirements on iris images in order to design the imaging parts of an iris recognition system.

19.2 Review of Photographic Concepts

In this section, we discuss the interaction of light with the optical components of an iris imaging system or camera. We will treat light using a variety of equally valid models.

19.2.1 Geometrical Optics

Light can be modeled as massless particles, that is, photons, waves, or rays. Geometrical optics builds upon a ray model in order to understand the propagation of light from an illuminated object through an optical system to an image sensor such as in a camera. We next consider some simple concepts of geometrical optics that will facilitate our examination of iris imaging systems. While commercially available ray-tracing programs can predict the performance of complicated lens arrays, we will use the simplest thin lens approximations to obtain predictions of lens behavior. These predictions based on simple approximations that model complex multi-element lenses as single thin lenses are still useful.

19.2.1.1 Focal Length of a Lens

A lens is specified by a number of parameters; we will start by using only a focal length, defined as the distance from the plane of a thin lens to the point of convergence of parallel rays through the lens. Figure 19.1 shows a thin lens focusing an object onto an image plane. Two important rays are traced by dashed lines. A marginal ray leaves the object approaching the lens in a direction parallel to the optical axis. When the marginal ray crosses the plane of the thin lens, it is refracted so that it crosses the optical axis at a distance from the lens equal to the focal length of the lens. A chief ray leaves the object in the direction of the center of the thin lens, emerging on the opposite side of the lens with its direction unchanged. Both rays carry information about the object as they converge at the focused image.




Fig. 19.1 An object (blue arrow on left, d_o from the lens) is imaged onto an image plane (as inverted blue arrow on right, d_i from the lens.) The so-called marginal and chief rays converge in the image plane and, in this case, predict the position of the tip of the image arrow

When an object at distance d_o creates a focused image at distance d_i through a lens of focal length f as shown in Fig. 19.1, the geometry of similar triangles yields the lens maker's formula:

$$\frac{1}{f} = \frac{1}{d_o} + \frac{1}{d_i} \quad (19.1)$$

The magnification of an imaging system is the ratio of the size of the image h_i to the size of the object h_o , which can be rewritten in terms of the image and object distances, $M = d_i/d_o$. Starting with d_o and f and using the lens maker's formula (19.1), we find

$$M \equiv \frac{h_i}{h_o} = \frac{d_i}{d_o} = \frac{f}{d_o - f} \quad (19.2)$$

In a typical camera setup, the numerical value of M is most often less than one while in a microscope, M is typically greater than one. For the setup shown in Fig. 19.1, when the object distance is much greater than the focal length of the lens, $M \sim f/d_o$. In other words, M is roughly the reciprocal of the number of focal lengths between the object and the camera.

19.2.1.2 Object-Referred Pixel Size and Field of View

The digital image sensor of a camera is made from an array of photosensitive picture elements or pixels. Thus, an image in a digital camera is cast upon a typically square array of elements which convert the image light intensity to photo-induced charge which is then converted to voltage and measured for each pixel in order to record the image. A useful concept in comparing the size of features in an image to the corresponding features of the object is the object-referred pixel size L_{orp} , defined as the dimension on the object which is imaged onto a single pixel on the imager. If the size of a single square pixel is L_p , then

$$L_{orp} = \frac{L_p}{M} \sim \frac{L_p \cdot d_o}{f} \quad (19.3)$$

in which the second expression for L_{orp} applies to cases in which $d_o \gg f$. For example, if an image sensor with 6 μm pixels receives a focused image of a 2 m distant object using a 100-mm focal length lens, $L_{orp} \sim 120 \mu\text{m}$. In the spirit of geometrical optics, this means that square blocks of the object 120 μm on a side are imaged onto single pixels as shown in Fig. 19.2 in which an iris (object) is imaged onto an image sensor.

The field of view of a digital camera is determined by the size of the image sensor assuming that the optical system is designed to create an image that covers the sensor. If the sensor in Fig. 19.2 is made of N_h pixels in the horizontal direction

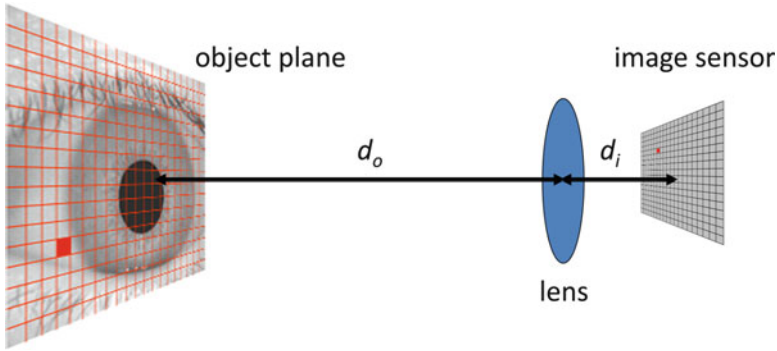


Fig. 19.2 An object (*left*, d_o from the lens) is imaged onto an image plane (*right*, d_i from the lens.) The portion of object shown by the *red square* is imaged onto a single pixel shown as another *red square* on the image plane

and N_v pixels in the vertical direction, then the dimensions of the region on the object that can be imaged by the sensor are determined by the magnification of the optics and the pixel count:

$$\text{FOV}_j = N_j \cdot L_{orp} \quad \text{where } j = h, v \quad (19.4)$$

Thus, in our previous example in which $L_{orp} \sim 120 \mu\text{m}$, an image sensor with $1,280 \times 960$ pixels will provide a field of view of $14.6 \times 10.9 \text{ cm}$. A sensor with more pixels will create a larger field of view as will an optical system with lower magnification, for example, shorter focal length lens and/or greater object distance.

Sometimes, it is useful to consider the field of view in terms of angle rather than area. In this case, we simply convert the horizontal and vertical dimensions of the field of view to angles subtended by the field of view with respect to the optical axis. For angular field of view, Eq. (19.4) becomes


$$\text{angular FOV}_j = 2 \arctan \left(\frac{N_j \cdot L_{orp}}{2} \right); \quad \text{where } j = h, v; \quad \text{units of radians} \quad (19.5)$$

19.2.1.3 F-Number and Lens Aperture

In addition to focal length, camera lenses are often characterized by their so-called F-number, where F is defined by Eq. (19.6) in terms of focal length f and lens aperture D and illustrated in Fig. 19.3.

$$F = \frac{f}{D} \quad (19.6)$$

Fig. 19.3 Schematic of a thin lens of focal length f and aperture diameter D which would then have $F = f/D$



Thus, a lens with 100-mm focal length and a 25-mm diameter clear aperture for gathering light would be conventionally labeled $F/4$ and spoken as “ $F4$.” Many lenses have variable apertures and thus variable F-number. Changing the diameter of the lens aperture varies the amount of light collected so that in dim lighting, a low F-number would gather light at a greater rate and necessitate a shorter exposure than at high F-number. In the days of film cameras, a lens with a larger aperture that would admit more light than a lens of the same focal length with a smaller aperture was referred to as a “fast” lens since the rate at which the film would be exposed was relatively fast. A lens with comparatively smaller aperture was referred to as a “slow” lens.

19.2.1.4 Depth of Field

Aperture size affects the photographic depth of field of an imaging system. Depth of field refers to a range in object distance that an object can occupy for a given lens and given image distance that results in an in-focus image. It is the subjective judgment of focus quality that makes depth of field a subjective quantity. To quantify depth of field requires specifying the dimension of a blur spot in the image plane beyond which a point image is judged out of focus. The size of such a blur spot is referred to as the “circle of confusion” in the context of depth of field. Figure 19.4 schematically shows three objects spread out in object distance along the optical axis. Only the middle object is placed at an object distance that is focused for the lens and image distance, that is, the object distance of the middle object together with the image distance and lens focal length satisfy the lens maker Eq. (19.2).

Figure 19.4 illustrates the cases of large and small aperture lenses and graphically illustrates why a lens set to larger aperture exhibits relatively shallow depth of field compared to the same lens set to a smaller aperture. In the insets of Fig. 19.4, the low F-number case (top) creates images using rays, some of which impinge on the image sensor at steeper angles than in the case of higher F-number (bottom). An imaging system with high F-number (small aperture) has a larger range in object positions for a given circle of confusion because the rays from the various positions on the object plane approach the imager at relatively shallow angles and diverge more gradually in the vicinity of the image plane. The effect of F-number on image focus is shown dramatically in Fig. 19.5 for two cases that feature objects spread out in object distance.




Fig. 19.4 Three objects (*blue, green, and red arrows*) at three different object distances create images at corresponding image distances. The image distance of only the middle object (*green arrow*) corresponds to the position of the focal plane at which the image sensor is placed. Thus, only the *green arrow* is in focus. The *upper figures* show the case of a large aperture (low F-number) in which rays from the base of the arrows impinge on the image sensor at relatively large angles. Thus, as shown in the *upper right inset*, the *blue arrow* is blurred because rays from its base cross in front of the imager and have diverged at the position of the sensor. Similarly, rays from the base of the red arrow have yet to converge as they hit the image sensor and so the *red arrow* is also out of focus. The lower figures show the case of small aperture (high F-number) in which rays from the base of the arrows impinge upon the imager at relatively low angles. In this case, although the blue and red arrows are also out of focus, the degree of blur is smaller. A larger aperture results in shallow depth of field, whereas smaller aperture creates relatively larger depth of field



Fig. 19.5 Example of a scene that extends over a wide range in object distance photographed with shallow (*left*) and large (*right*) depth of field. The focal plane is set at the labeled bottle. While the bottle is in focus in both scenes, the foreground and background objects on the *left* are out of focus as imaged with aperture set to $F/2.8$ while the same features are sharper on the *right* as imaged with aperture set to $F/16$

The ray geometry of Fig. 19.4 provides an expression for depth of field (DoF) that is presented in (19.5) and which can be approximated by Eq. (19.7):

$$DoF \approx 2C \cdot F \cdot \left(\frac{d}{f}\right)^2 \quad (19.7)$$

in which d is object distance, f is lens focal length, F is lens F-number, and C is the subjective quantity circle of confusion. As an example, if we pick circle of confusion C to equal 1 pixel dimension = 6 μm , (certainly, C cannot be less than 1 pixel) and if $d = 2$ m, $f = 100$ mm, and $F = 4$, we find $DoF \approx 2$ cm. Specifically, DoF represents the distance measured along the optical axis of the camera from the nearest to the furthest in-focus object position. In this particular case, the DoF is approximately 1% of the object distance. Changing the F-number of the lens from $F/4$ to $F/8$ by decreasing the lens aperture by a factor of two increases the DoF by a factor of 2 to about 4 cm. If we accept as focused an image with a blur spot in the image plane of two pixels instead of one, DoF would double. (This might be the case with a color image in which each individual color is interpolated over ~ 2 pixels.) Changing object distance or lens focal length by a factor of 2 would change DoF by a factor of 4.

Geometrical optics that we have presented so far predicts a field of view determined by lens and image sensor characteristics and a depth of field set by the lens focal length and aperture for a given object distance. Together, a volume with lateral dimensions (measured perpendicular to the optical axis) equal to that of the field of view and with axial dimension (measured along the optical axis) equal to that of the depth of field defines a capture volume. Inside the capture volume, an in-focus object is imaged onto the sensor. We will use the concept of capture volume in designing an iris imager in Sect. 19.3.

19.2.2 Spatial Resolution

Spatial resolution measures the response of an image capture device to the spectrum of spatial frequencies that comprise an object. Spatial resolution of an imaging system is a function of spatial frequency (pitch and orientation) as well as lens- and image sensor-related parameters including position in the field of view (lens-related), illumination spectrum (lens- and sensor-related), object distance (lens-related), and brightness (sensor-related). Because spatial resolution is fundamentally limited by the wave phenomenon of diffraction, we will start the discussion of spatial resolution using a wave model of light rather than the ray model of geometrical optics. We will then use ray tracing to consider aberrations due to imperfect lenses which further degrade spatial resolution. Finally, we will consider quality of focus as it affects spatial resolution.



Fig. 19.6 Image of water waves interacting with a breakwater taken from Ref. [5]

Ultimately, the spatial resolution of an entire system determines the fidelity with which an iris imaging system renders the large and small features of an iris. However, looking ahead to Sect. 19.4, it is important to keep in mind that a given algorithm exploits a particular subset of the iris information, determined by its own spatial frequency response. Therefore, accurate prediction of biometric performance of an iris recognition system requires knowledge of both the spatial resolution of the capture device and spatial frequency bandwidth of the matching algorithm.

19.2.2.1 Diffraction Fundamentally Limits Spatial Resolution

A measure of spatial resolution of an imaging system is its point spread function (PSF), the image of a point object. Observing a small LED pilot light across a darkened room without corrective lenses provides a view of the PSF of the observer's eyes at the LED wavelength. Crudely, the PSF is a blur spot of size that characterizes the smallest feature resolvable by an imaging system. An ideal image convolved with the system PSF produces the actual captured image (allowing for variation of PSF over the field of view). Diffraction fundamentally limits the PSF, dependent upon the size (and shape) of the aperture of the collecting lens, the object distance, and the illumination wavelength.

Loosely speaking, diffraction results from the interaction of waves with an object illuminated by these waves. In this sense, water waves, sound waves, and light waves behave similarly with an interaction characterized by the ratio of object size to wavelength. A clear example of diffraction of water waves is shown in Fig. 19.6. The wavelength of near-infrared light is about seven orders of magnitude shorter than that of ocean waves of Fig. 19.6, but diffraction of light nevertheless bears a close resemblance to that of water or sound waves. To the extent that we




Fig. 19.7 Monochromatic plane waves of wavelength λ impinge normally on a grating of spatial wavelength Λ from the right. The grating scatters the waves along three directions described in terms of their angle θ from the normal

can approximate the diffraction of light as a scalar phenomenon, thereby ignoring the vector nature of electromagnetic waves, diffraction of light can be computed relatively easily. In the so-called Fraunhofer approximation [6], in which the object size must be comparable to or greater than the wavelength of light but must also be much smaller than the distance between the object and viewer and the field of view, we find that the distribution of light diffracted by an object closely approximates a scaled Fourier transform of the object, where the scaling relates to the optical wavelength. A mathematical description of scalar diffraction theory that underlies the Fraunhofer approximation is given in Refs. [7] and [5].

Reference [5] provides an approximate mathematical description of diffraction, from which we can gain an intuitive understanding of the diffraction phenomenon in the Fraunhofer diffraction regime by thinking of objects that diffract light in terms of the effects upon light waves of each of the individual spatial frequency components of an object. A simple object to start with is a sinusoidal grating of the type found in an optical spectrometer, manufactured holographically and illustrated schematically in Fig. 19.7.

Light normally incident on the grating (from the front in the case of a reflective grating or the back in the case of a transmissive grating) is scattered by each line in the grating. The scattered light constructively interferes as it radiates from the grating with its energy concentrated in three directions, one normal to the grating and two at equal and opposite directions relative to the normal direction. The directions, in terms of angle θ from the normal direction, are given by the formula

$$n\lambda = \Lambda \sin \theta \quad (19.8)$$

where λ = wavelength of light, Λ = wavelength of grating, and $n = -1, 0,$ and 1 . Each angle corresponds to a direction of constructive interference that can be assigned to a spatial frequency of the grating: $n = \pm 1$ for the two sine waves of wavelengths $\pm\Lambda$ that combine to create a standing wave of the proper wavelength and phase, and $n = 0$ for an average value that simply offsets the standing wave. Here, we avoid the issue of the finite extent of the grating and think of it as very




Fig. 19.8 Diffraction patterns created by 850-nm light on a screen placed 2 m in front of circular apertures of different sizes, 200- μm diameter (*left*) and 600- μm diameter (*right*). The red circles suggest the clear opening of a lens (in this example, 70-mm focal length set to $F/4$) that could receive the diffracted light and reimage it on a sensor

large, containing many lines. In this approximation, we can say that the sinusoidal grating diffracts light in three discrete directions. Note that the ± 1 order diffraction angle increases with λ/Λ , that is, with longer wavelength light or shorter wavelength (higher spatial frequency) grating. Because we can regard any object as a sum of its spatial frequency components, we can analyze Fraunhofer diffraction as the net effect of diffraction from the set of individual gratings that comprise an object, each with a particular grating period, phase, amplitude, and orientation. Objects with small features possess high spatial frequencies which diffract light at large angles. Objects that are comprised of comparatively lower spatial frequencies will diffract light at smaller angles.

An illustration of Fraunhofer diffraction from two similar objects appears in Fig. 19.8. Here, light strikes circular apertures from the back and diffracts forward where it hits screens in front of the apertures. In the case of a small aperture, comprised of relatively high spatial frequencies, the diffraction pattern is coarse and spreads over a relatively large area because the diffraction due to its high spatial frequencies occurs at large angles. In comparison, the diffracted pattern appearing on the screen due to the larger aperture is concentrated over a relatively smaller area. The larger aperture contains relatively lower spatial frequency content which diffracts light of a given wavelength at smaller angles. The key concept is that small or sharp features, made of high spatial frequencies, diffract light at larger angles relative to the normal direction. In fact, we can regard the diffraction pattern in the Fraunhofer regime as a scaled Fourier transform of the object with the lowest spatial frequency components traveling straight to the screen and the higher spatial frequency components diffracted at angles away from the center of the diffraction pattern.

A direct consequence of diffraction is illustrated by the red circles in Fig. 19.8 which represent the input clear apertures of camera lenses that might be used to focus on and create images of the two illuminated circular apertures that are





Fig. 19.9 Top row: Images of 600- μm aperture with simulated diffraction due to 850-nm illumination, 2.5 m object distance, 70-mm focal length lens set to $F/4$ (17.5-mm clear lens aperture, *left*) and $F/11$ (6.4-mm clear lens aperture, *right*) showing increasing diffraction effects with smaller aperture. Lower row: Images of human iris with simulated diffraction under same imaging conditions with 70-mm lens set at $F/4$ (*left*) and $F/11$ (*right*)

diffracting the illumination. Because a lens cannot admit light beyond its clear opening, it cannot use rejected light to rebuild the image of the object. Since the higher spatial frequencies of an object result in larger angle diffraction, the red circles represent a cutoff that effectively low-pass filters the spatial frequency spectrum of the object. For different lens apertures, as shown in Fig. 19.9, the same object imaged through a smaller aperture suffers more filtering than imaged through a larger aperture. The effects of the filtering can be visualized at the image plane of the lens where the spatial frequency components are reconstituted into a filtered representation of the original object. The more of the original components that are included in the image, the more faithfully the image represents the original. In Fig. 19.9, we simulate the 600- μm aperture shown in Fig. 19.8, filtered by $F/4$ and $F/11$ apertures. The larger aperture ($F/4$) produces a sharper image. In a second example using a human iris as an object, the same effect of (simulated) diffraction is illustrated in lower row of Fig. 19.9 with aperture also set to $F/4$ and $F/11$. Realistic loss of resolution of the fine iris details due to diffraction is visible in the $F/11$ iris image.

Fig. 19.10 Diffraction blur spots on an image sensor showing pixel grid. Large blur spot (*right*) results in diffraction-limited resolution while small blur spot (*left*) results in pixel-limited resolution. Center blur spot is match to pixel pitch



Diffraction limits the size of the smallest object that can be imaged by an optical system. The diffraction limit can be derived in the Fraunhofer approximation (19.4) and (19.8) and simply stated:

$$\Delta x_d \cong \frac{\lambda d}{D} \tag{19.9}$$

in which Δx_d = the diameter (full width at half max, FWHM) of the smallest spot that can be resolved in the object plane, λ = illumination wavelength, d = object distance, and D = aperture diameter. Spots that are smaller than the diffraction limit are blurred to the diffraction-limited spot size. We can use Eqs. (19.2) and (19.6) to convert the diffraction-limited spot size in the object plane to that appearing in the image plane (indicated with a prime) from which we get

$$\Delta x'_d \approx \frac{\lambda \cdot d_o}{D} \cdot M \approx \frac{\lambda \cdot d_o}{D} \cdot \frac{f}{d_o} = \lambda \cdot F \tag{19.10}$$

in which F = F-number of the lens.

Equation (19.10) relates the image plane blur spot to the F-number of the lens for a given illumination wavelength so that, for example, with a wavelength of 850 nm typical of near-infrared light and an $F/4$ lens, we would expect an image plane blur spot of approximately $\Delta x'_d \approx 3.4\mu\text{m}$. It is instructive to compare the image plane blur spot size to that of the pixel pitch of the imager as pictured schematically in Fig. 19.10. The ratio of image plane spot size to pixel pitch is simply $\Delta x'_d / L_p \approx (\lambda / L_p) \cdot F$, which is “tuned” by F . A large value of F (small aperture, strong diffraction-related filtering) creates a blur spot on the imager that is larger than a pixel dimension, while a small value of F (large aperture, weak diffraction filtering) creates a blur spot that is smaller than a pixel dimension. Figure 19.10 illustrates both of these cases and the intermediate case in which $F \approx L_p / \lambda$ which matches the blur spot and pixel dimension.

In the case of a very small blur spot, the image resolution is pixel limited. In other words, the smallest image feature is determined by the spatial quantization of the image into chunks corresponding to individual pixels. In this case, the lens resolution (and cost!) exceeds the capabilities of the image sensor. The opposite case in which the diffraction-induced blur spot size covers many pixels creates a diffraction-limited image. In this case, the sampling density (and cost!) of the pixel array on the image sensor exceeds the capability of the lens. A practical design optimizes performance bearing in mind the cost of components. By such reasoning, the intermediate case in which F-number, pixel dimension, and illumination wavelength are arranged to match the image plane blur spot to the pixel dimension is optimal. As an example, for 6- μm pixels and 850-nm illumination, $F/7$ provides an optimal match. For the same illumination but with smaller pixels $\sim 2\ \mu\text{m}$, as would be found in cell phone imagers, $F/2.4$ is required to optimize performance. From the previous discussion of depth of field, we see that optimizing blur spot size for smaller pixels (for a given illumination wavelength) forces use of a lower F-number which reduces DoF. Balancing design considerations of image resolution, depth of field with practical considerations of system performance and cost involves trade-offs that are the basis of camera design. As we shall see, system performance criteria are determined by the end use of the camera, be it for wedding photography, aerial surveillance, or iris recognition. In any design, however, the fundamental (unavoidable) effects of diffraction must be considered.

19.2.2.2 Lens Quality Practically Limits Spatial Resolution

Fundamental limitations such as diffraction are compounded by practical limitations related to lens quality. A mass-produced molded plastic lens might be lower in cost compared to a diffraction-limited multi-element aspheric glass lens, but it might also produce a poorer quality image. Careful ray-tracing programs follow the trajectory of light through any arrangement of refracting materials and can determine parameters such as the quality of focus. Sacrifices in design and materials can result in lens aberrations that compromise performance as illustrated [8] in the ray tracing of Fig. 19.11. Pictured in Fig. 19.11 is a thin lens that creates a sharp focus. Close inspection of the focus reveals that parallel rays, especially those entering the lens at its periphery do not cross at a single point. This phenomenon is graphed as a scatter pattern thereby quantitatively connecting the lens design to a measure of spatial resolution. Certainly, we cannot expect to create a blur spot that is smaller than the scatter of rays at a point of best focus. By means of ray tracing, we discover a level spatial resolution that can achieve the diffraction limit in the case of high-quality lenses but for other lenses can be considerably poorer in spatial resolution than predicted by the fundamental diffraction limit.

Referring back to Fig. 19.10, we see that opening the aperture in an attempt to achieve pixel-limited spatial resolution might result in aberration-limited resolution as suggested by the red circle created by a collection of rays such as shown in Fig. 19.11 that is larger than the theoretically achievable diffraction-limited spot. In




Fig. 19.11 Parallel rays impinging on a lens (*left*) are refracted to a focus within the *black rectangle*. An expanded view of the area within the rectangle shows that rays do not all focus to the same point due to lens aberration which is most apparent in rays passing through the outer portions of the lens. A collection of points corresponding to rays in the image plane is plotted (*right*) showing the scatter which produces a blurred focus in the image plane

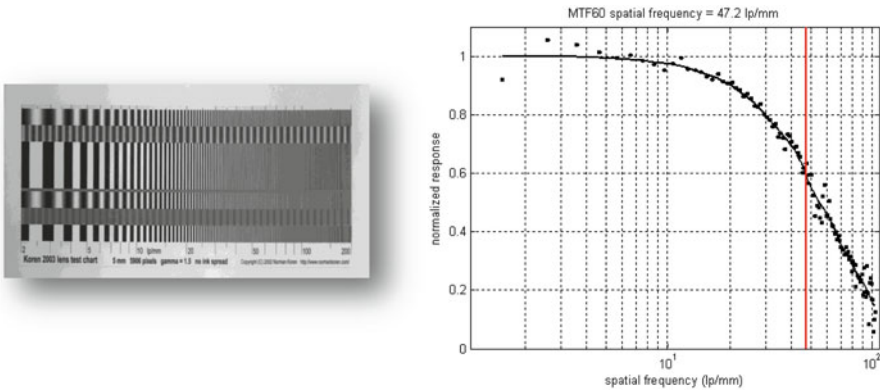



Fig. 19.12 A test target due to Ref. [9] (*left*) features sine and square waves of decreasing spatial frequency which can be imaged resulting in contrast (normalized peak to valley variation) which is plotted as a function of image plane spatial frequency (*right*). *Red line* indicates image plane spatial frequency at which contrast drops to 60% in keeping with ISO/IEC image quality specification of Ref. [2]

practice, there is an optimum aperture size for a given lens to achieve best spatial resolution. Spatial resolution can be measured by a number of techniques and can be characterized in terms of modulation transfer function (MTF) which is the Fourier transform of the point spread function. One technique for measuring MTF uses a grating with increasing pitch and uniform contrast, as illustrated in Fig. 19.12 and detailed in Reference [9].

Roughly speaking, it takes a distance equal to that of four blur spots to discern a single line pair (or sinusoidal cycle) at the limit of resolution, for example, when contrast in the perceived modulation pattern has dropped to half of its original value. By such a measure, the data of Fig. 19.13 indicate that at $F/5.6$, the lens under test should be producing a blur spot at 850 nm of approximately $(1/MTF_{50})/4 \sim 4.5 \mu\text{m}$. The image plane diffraction-limited spot for this particularly high-quality lens is

Fig. 19.13 MTF₆₀ as a function of F-number for a 70-mm lens showing a peak at around $F/5.6$ at which spatial resolution is an optimum trade-off of aberration and diffraction



approximately $\Delta x_d \approx F \cdot \lambda \sim 4.8 \mu\text{m}$, in good agreement within the spirit of the approximations used. If we measure MTF as a function of F-number, we trace out a lens performance curve that conforms to the suggestion of Fig. 19.10, namely, that high F-number performance suffers due to diffraction and low F-number performance suffers due to lens aberration. Performance of the lens under test peaks at $F/5.6$. Thus, we have combined the wave concept of diffraction and the ray concepts that describe lenses to get a picture of how lenses and image sensors work together to provide spatial resolution.

For concreteness, we approximate the measured MTF of a lens at a particular wavelength and F-number such as that represented in Fig. 19.12 with a Gaussian profile and 50% contrast point which we will call MTF₅₀. We assume that the effects of diffraction and aberration contribute independently to the measured MTF curve. (With a perfect lens, we would measure the effects of diffraction only.) The width of the PSF due to diffraction and aberration can be approximated by Eq. (19.11) which assumes that the PSF, of diameter $x'_{1/2}$, is due to independent Gaussian contributions of diffraction and aberration. In Eq. (19.11), the broadening due to aberration $\Delta x'_{ab}$ and that due to diffraction λF adds in quadrature.

$$x'_{(\frac{1}{2})} = \sqrt{(\Delta x'_{ab})^2 + (\lambda \cdot F)^2} \tag{19.11}$$

Finally, we can account for defocus by further broadening the PSF diameter. The details of defocus broadening relate to the size and shape of the lens aperture. To avoid incurring defocus broadening, we must ensure that an object remains within the depth of field, that is, remains in the capture volume.

Fig. 19.14 Schematic lens, aperture, and image sensor





Fig. 19.15 Analogous representation of image sensor pixels and light particles as buckets catching rain. A given bucket catches, on average, S photons. The bucket-to-bucket variance in photon number due to shot noise is also S

19.2.3 Signal and Noise

Thus far, we have discussed light in terms of waves and rays. Now, we switch to a photon description of light as a particle similar to a drop of rain or a billiard ball. This picture is valuable in discussing light as it hits an image sensor and is transformed into electrical charge which is counted in order to profile intensity as a function of position in an image. Figure 19.14 shows a schematic camera with lens, aperture, image sensor, and transducer that creates a digital representation of an image. A row of pixels from the image sensor, represented as buckets,¹ is shown in Fig. 19.15. Photons from a flat field image (spatially uniform intensity) impinge upon the lens and are imaged onto the sensor. After an exposure time when the shutter is closed and the sensor is no longer exposed to light, each bucket or pixel holds a certain number of photoelectrons, each resulting from a photon. In the case of a flat field image, the number of photons in the buckets would be expected to be uniform since the object that fills the field of view is itself uniformly irradiated with light. However, experiments show that the photoelectron number varies from pixel to pixel due to the random rate of photons arriving from the object. The variance of the photoelectron number across the imager is equal to the average number of photons when the variation of photon rate is determined only by the random rate of photon arrival. The process that determines the variance of photon number is known as shot noise and is a consequence quantum mechanics.

¹In this discussion, we assume that photons that end up in one bucket remain in that bucket. In reality, we need to consider the possibility of photons finding their way into neighboring buckets, an effect that would further degrade spatial resolution.




Fig. 19.16 Number of photoelectrons captured as a function of pixel (bucket) number in an area of an image exposed to a low light level (*left*) and higher light level (*right*). At higher light level, the absolute noise level is relatively higher, but its fractional effect on the image is less than in the lower light portion of the image. We can say that the SNR of the brighter part of the image is higher than the dimmer part of the image

The result is that after an exposure time, the j th pixel (bucket) contains S_j photoelectrons with an average over all pixels of $\langle S_j \rangle$. In the case of shot noise, the variance in S_j is equal to $\langle S_j \rangle$ and, therefore, the standard deviation in the photon count in the buckets is $\sqrt{\langle S_j \rangle}$. If we consider an image with dim and bright regions, we find that the signal-to-noise ratio (SNR), defined as the ratio of the average photoelectron number to the standard deviation in photoelectron number, becomes

$$\text{SNR} = \frac{\langle S_j \rangle}{\sqrt{\langle S_j \rangle}} = \sqrt{\langle S_j \rangle} \quad (19.12)$$

In a dim region in which the average photoelectron count is 100, $\text{SNR} = 10$ while in a bright region in which average photoelectron count is 10,000, $\text{SNR} = 100$ as shown schematically in Fig. 19.16 in which the noise in the dim region is fractionally higher than in the brighter region due to shot noise.

Other sources of noise can contribute to variation in photoelectron count in a given pixel including read noise, a variation occurring in the transformation of charge to voltage, and dark current which creates photoelectrons in the absence of signal from an object. Processes beyond fundamental shot noise contribute independently to the total noise which is computed as the square root of the sum of the squares of each independent contribution. As an example, SNR in the presence of read noise N_r becomes

$$\text{SNR} = \frac{\langle S_j \rangle}{\sqrt{\langle S_j \rangle + N_r^2}} = \sqrt{\langle S_j \rangle} \cdot \frac{1}{\sqrt{1 + \frac{N_r^2}{\langle S_j \rangle}}} \quad (19.13)$$

from which it is clear that SNR reduces to that shown by Eq. (19.12) when read noise is insignificant. However, in the case of very dim images, read noise and other contributions become lower bounds on noise and decrease SNR.




Fig. 19.17 Schematic view of imaging system showing object (*right*), a distance d_o from the lens and aperture. Pixels from the imager located a distance d_i from the lens, are projected onto the object and appear with size L_{orp}

We now estimate SNR for an imaging system starting with an object that is illuminated at wavelength λ by a level of irradiance (power/area) of I . We assume that the light scatters from the object in a Lambertian pattern, that is, with an angular distribution proportional to the cosine of the angle between the normal and the object and the angle of scattering and with a diffuse reflection coefficient or albedo of a . A Lambertian model is reasonable for a wide class of non-shiny materials including the human iris. Next, we assume an optical system with an aperture and lens which creates an image on a sensor with individual pixels. For simplicity, we assume that we are in the shot noise regime in which contributions to noise other than shot noise are insignificant due to adequate illumination. Finally, we assume that the sensor has a quantum efficiency QE (ratio of photoelectrons created to incident photons) and that the exposure time is τ_{exp} . The imaging system is shown in Fig. 19.17.

The number of photoelectrons per pixel per exposure time is given by Eq. (19.14), using Eqs. (19.2) and (19.6), in terms of pixel dimension L_p , magnification M , albedo α , collection cone half-angle θ , exposure time τ_{exp} , quantum efficiency QE, photon energy $E_\gamma = hc/\lambda$, F-number F , and wavelength λ , as well as fundamental constants, h and c .

$$S = I \cdot \left(\frac{L_p}{M}\right)^2 \cdot a \cdot \sin^2\theta \cdot \frac{\tau_{exp} \cdot QE}{E_\gamma} = I \cdot \left(\frac{L_p}{2F}\right)^2 \cdot a \cdot \frac{\tau_{exp} \cdot QE \cdot \lambda}{hc} \quad (19.14)$$

Equation (19.15) gives an expression for signal-to-noise ratio SNR for the imaging system as

$$SNR = \left(\frac{L_p}{2F}\right) \cdot \sqrt{\frac{I \cdot a \cdot \tau_{exp} \cdot QE \cdot \lambda}{hc}} \quad (19.15)$$

Notably, lens focal length and object distance do not appear in Eq. 19.11. SNR is proportional to pixel size and inversely proportional to F-number; larger pixels and a faster lens improve SNR. In addition, SNR is proportion to the square root of a factor that relates to photoelectron collection. Increased irradiance on target, higher

target albedo, longer exposures, and higher sensor quantum efficiency all improve SNR (but only proportional to the square root of the increase).

It should be noted that on some digital still cameras, a setting labeled as ISO (from the International Standards Organization) produces higher sensitivity to light. The equivalent video camera setting is more aptly labeled signal gain. In both digital still and video cameras, increasing the ISO or gain settings boosts the transduction ratio of charge to voltage and therefore does not increase SNR. In fact, charge-to-voltage amplification noise typically increases at higher gain settings so that high ISO or gain might reduce SNR.

19.2.4 Trade-Offs and Constraints

Camera design requirements often lead to trade-offs which produce design constraints. If illumination is dim and long exposures are impractical due to subject motion, a large aperture can be used at the expense of shallow depth of field and increased susceptibility to lens aberration. Alternatively, high signal gain can be used at the expense of SNR. If a large depth of field is required, then a small lens aperture might reduce spatial resolution due to diffraction, in which case a reduced cost lens might do just as good a job as an expensive one. Some design trade-offs are fairly obvious, such as exposure time and signal-to-noise ratio. However, under some situations, design trades produce interesting and nonobvious design constraints. Such is the case with depth of field under the constraint that matches the diffraction-limited spot size to a pixel dimension as is practical for high spatial resolution iris imaging.

We start by considering Eq. (19.7) in which circle of confusion C matches a single pixel dimension L_p . In other words, while the blur in the image plane due to defocus is comparable to the size of one pixel, we consider the object to be in focus. Furthermore, using Eq. (19.10), we match the diffraction-limited blur to the size of one pixel, which, given the illumination wavelength λ and the pixel pitch, sets the F-number of the lens. Finally, we recognize that the quantity (dlf) in Eq. (19.7) is approximately the inverse magnification given in Eq. (19.2) for $d \gg f$, which can be written as $(dlf) \sim M^{-1} \sim L_{orp}/L_p$. Putting these pieces together, we get Eq. (19.16), which predicts depth of field under the conditions that (1) we maintain the diffraction-limited blur spot comparable to a pixel size and (2) we maintain object-referred pixel size at a given value dictated by the imaging application, both quite reasonable for an iris imaging system.

$$DoF \approx 2C \cdot F \cdot \left(\frac{d}{f}\right)^2 \sim 2L_p \cdot \frac{L_p}{\lambda} \cdot \left(\frac{L_{orp}}{L_p}\right)^2 = 2 \frac{L_{orp}^2}{\lambda} \quad (19.16)$$

Equation (19.16) yields a constraint, namely, that under the given conditions, depth of field is approximately constant, independent of object distance.

If $L_{\text{orp}} = 100$ mm and $\lambda = 850$ nm, $\text{DoF} \sim 3$ cm. (This is, of course, because lens focal length must increase with object distance to maintain constant L_{orp} .) Nevertheless, Eq. (19.16) suggests that an ordinary iris recognition camera working at 30 cm, using a diffraction-limited 10-mm lens and a 30-m long-distance iris recognition system using a high-quality 1-m focal length telescope will have comparable depths of field. Equation (19.16) also is useful in determining whether an autofocus mechanism is needed for an iris imaging system or whether a fixed focus imager, for which the subject or operator adjusts the subject-to-camera distance, is adequate.

Another interesting case that arises under similar conditions predicts signal-to-noise ratio when lens aperture is chosen to match diffraction blur to a pixel dimension. In this case, Equation (19.15) yields

$$\text{SNR} = \left(\frac{L_p}{2F} \right) \cdot \sqrt{\frac{I \cdot a \cdot \tau_{\text{exp}} \cdot \text{QE} \cdot \lambda}{hc}} \sim \sqrt{\frac{I \cdot a \cdot \tau_{\text{exp}} \cdot \text{QE} \cdot \lambda^3}{4hc}} \quad (19.17)$$

independent of lens parameters and object distance (assuming that irradiance on target is a constant). Assuming that the irradiance is limited by eye-safety considerations, that the exposure time is limited by subject motion considerations, and that QE is typically 0.1–0.2, Eq. (19.17) can be used to predict an upper bound on SNR in a given application. Equations (19.16) and (19.17) express design constraints that arise from reasonable conditions encountered in design of an iris imaging system.

19.3 Camera Design

The previous Section reviewed photographic concepts and laid the groundwork for design of an iris imaging system. In this Section, we will use the concepts to design two iris cameras that work in different scenarios. The first is more typical of iris cameras available at the time this Chapter was written, working at around 0.3 m. Such systems include wall-mounted and handheld units, some of which capture a single iris and some of which capture dual iris images, often with multiple cameras. The second iris camera captures images from ten times further, 3 m, and is useful in applications with fewer subject constraints. As such, the second system must allow for position and motion of the subjects. In both systems, we will assume that illumination operates with an approximate wavelength of 850 nm at an eye-safe level. Eye safety is a primary concern for active illumination systems. A full discussion of eye safety is out of the scope of this Chapter; the reader is referred to Ref. [4]. One important aspect of eye safety is worth mentioning, however. In the near-infrared wavelength band, shorter duration exposures are safer than longer ones with the threshold level of irradiance given in Ref. [4] by Eq. (19.18):

$$I_{\text{TLV}} = \frac{1,800}{t^{0.75}} \quad \text{in units of mW/cm}^2 \quad (19.18)$$

The implication is that higher irradiance exposure is safe in short doses. If short pulse illumination coordinated with an equally short camera exposure times can supply enough photons to meet SNR requirements while remaining eye safe, it can address the problem of motion blur. On the other hand, continuous and less bright illumination can avoid the complexity of pulsed illuminators at the expense of lower motion tolerance.

19.3.1 0.3-m Iris Imaging System

Design of a $d_o = 0.3$ m iris imaging system optimizes spatial resolution and signal-to-noise ratio for the prescribed use, in this case assuming a cooperative and motionless subject aligned with the field of view. For concreteness, let $\text{FoV} = 0.13$ m in width which allows for capture of two irises in one image and let illumination wavelength $\lambda = 850$ nm, ideal for penetrating dark iris pigmentation but still detectable by a silicon image sensor [12].

We shop for image sensors with several factors in mind. First, we want good quantum efficiency at 850 nm. At 850 nm at the time of this writing, 10% QE is good for a front-illuminated silicon image sensor and 20% would be excellent. Next, we want enough pixels to give adequate resolution across an iris since we do not want to be pixel limited at a resolution that is inadequate for iris recognition. At the time of writing this Chapter, the ISO/IEC standard body has yet to finalize what resolution is adequate for iris recognition with proposals ranging from 7 to 20 samples per mm in the object plane. Finally, we want enough pixels in the imager to cover the field of view at the chosen sample density. As an example, we choose an object referred pixel size L_{orp} of 100 μm which gives 10 samples/mm and roughly 100 pixels across an iris (19.2). Knowing L_{orp} and knowing the width of the FoV dictates the number of pixels across the image sensor, in this case, $0.13 \text{ m}/100 \mu\text{m} = 1,300$. This particular example would suggest a 1.3-Mpixel imager with $1,300 \times 1,000$ pixels. (Of course, the imager could have more pixels so long as we could read images from it fast enough and could afford its cost.)

With an image sensor in mind, specifically one that has 1.3 MPx, adequate QE and operates with a global shutter to avoid image shearing, we shop for a lens, preferably one with good transmission and MTF at the illumination wavelength. A current ISO/IEC proposal for MTF of an iris imaging system suggests a minimum of 2 line pairs/mm on an object with at least 60% contrast. If we assume about 4 pixels per line/pair, 2 lp/mm translates to 8 pixels/mm which, for a conservative design, we round up to a pixel density of 10 px/mm on the object plane. L_{orp} is therefore 100 μm consistent with our sensor plans from above. Given the pixel size on the chosen imager, say 3.5 μm , and the object sampling density given by L_{orp} , we immediately know the required magnification $M = L_p/L_{\text{orp}} = 3.5/100$. Given the object distance d_o , M , and Eq. (19.2), we calculate the required lens focal length. In this case, we manipulate Eq. (19.2) to give

$$f = d_o \cdot \frac{1}{M^{-1} + 1} \quad (19.19)$$

which, for $d_o = 0.3$ m and $M = 0.035$, requires a 10-mm focal length lens.

The required MTF at the image plane is given by the ISO/IEC MTF standard for the object plane converted using magnification from 2 lp/mm to $(2 \text{ lp/mm} \times M^{-1}) = 58 \text{ lp/mm}$. If the chosen 10-mm focal length lens can far outperform this level of spatial resolution, we can regard the lens as diffraction limited. Two subtleties are worth mentioning. First, manufacturers' MTF specifications are frequently stated as a single number, for example, 100 lp/mm. In this case, the specification typically refers to the peak performance (see Fig. 19.13) in visible light. Since we are concerned with a near-infrared wavelength, we must use caution in accepting the manufacturer's specification. However, it is possible that the specification is overly conservative if it averages over a broad spectrum. The second consideration is lens transmission which is often tuned using lens coatings for visible light at the expense of near-infrared performance. We do not want a lens with elements coated in such a way to reduce contrast or transmit only 40% of 850-nm light! It is always best to measure the MTF of a candidate lens. We will proceed by assuming a diffraction-limited lens for our design.

Since it is cost-efficient to match the blur spot of the lens to the pixel size, assuming that the lens is diffraction limited or at least close, we use Eq. (19.10) to calculate the F-number for the desired "matched" condition, in this case $F = 3.5 \text{ } \mu\text{m}/850 \text{ nm} = 4.1$. With an $F/4$, 10-mm lens at an object distance of 0.3 m, we can compute the depth of field assuming a value of circle of confusion C . For now, we will take $C = 1 \text{ pixel} = 3.5 \text{ } \mu\text{m}$. Using Eq. (19.7) gives $\text{DoF} = 2.5 \text{ cm}$, quite close to the estimate of Eq. (19.16), representing approximately 10% of the object distance and requiring careful focusing. If we allow images to be taken outside the DoF, defocus will degrade spatial resolution so this might be a situation that would benefit from an autofocus mechanism, despite its complexity and cost.

With a depth of field and field of view, we can consider the "capture volume" of the camera as a volume in object space inside of which any object will appear in focus. Simplistically, we consider the capture volume to be the field of view $13 \times 10 \text{ cm}$ in directions perpendicular to the optical axis \times depth of field = 2.5 cm along the optical axis.

Finally, we calculate the SNR for the imaging system by assuming a given level of irradiance on the object, an object albedo, and an exposure time. The irradiance is determined by the available light from all sources of 850-nm illumination. These might include the sun but most commonly include one or more light-emitting diodes (LEDs) tuned to the desired wavelength. Iris albedo is roughly 10%. Exposure time is determined by the motion of the object. Sports photography requires very short exposures to capture rapidly moving subjects while avoiding motion blur. Astrophotography requires long exposures for its dim subjects that move slowly and predictably. For iris photography, we can arrange that subjects remain fixed in place by providing chin or forehead rests. If we allow freestanding, nominally stationary subjects, we must account for postural sway. For this example, we choose

an exposure time of 30 ms. A subject moving a few $\times L_{\text{orp}}$ in an exposure time will produce a motion-smear image. In this case, subject motion perpendicular to the optical axis must remain below $\sim 100 \mu\text{m}/0.03 \text{ s} \sim 0.5 \text{ cm/s}$ which is quite still.

Assuming a safe and readily attainable level of irradiance of 50 W/m^2 and an image sensor $\text{QE} = 0.1$, we use Eq. (19.15) to predict a value of $\text{SNR} = 110$. An acceptable value of SNR for iris recognition is not well established with a provisional ISO/IEC recommendation of ~ 100 (19.2). Note that SNR is often quoted in units of dB. However, this unit is ambiguous since some disciplines use $20 \times \log_{10}(R)$ while others use $10 \times \log_{10}(R)$ to convert the dimensionless number R to dB – it is unambiguous to use dimensionless signal-to-noise ratio R .

The example camera design will work well for a stationary subject, capturing both irises at once with adequate spatial resolution and signal-to-noise ratio. In the next section, we consider requirements of an iris imaging system used to capture irises at 3 m. Some of the requirements need to change while others stay fixed.

19.3.2 3-m Iris Imaging System

With an iris imaging system at 3 m from the subject, we can imagine a less constrained capture scenario. Therefore, as an added requirement, we specify that a subject may be approaching the system at 1 m/s with a component of velocity perpendicular to the optical axis of up to 0.1 m/s.

Again, we address illumination with an emphasis on eye safety. Placing the illuminators closer to the subject than 3 m allows lower power to be used which reduces the risk of approaching an unsafe level of irradiance or radiance. Whether at 3 m or closer, we will assume that illuminators are distributed in such a way to provide uniform lighting for all subject positions and heights required for the system without detailing the mechanism. Indeed, lighting is an important part of an iris imaging system but for now, we will simply stipulate that over the capture volume, we have an illuminator design that provides 50 W/m^2 for all subject heights and positions. Examples of manufacturers who have addressed the issue of illumination at a distance include those of Refs. [10] and [11].

In a similar fashion to designing the 0.3 m system in Sect. 19.3.1, we start by assuming $L_{\text{orp}} = 100 \mu\text{m}$. However, because of the less constrained scenario, we demand a field of view that is larger than the 0.3 m system, in this case, 25-cm wide and 20-cm high, necessitating an image sensor with a $2,500 \times 2,000$ pixel count. Such a 5-MPx image sensor will cost more than the 1.3-MPx sensor used in the 0.3-m system, but the larger FoV enables more tolerant alignment of the camera to the user. We again use an image sensor with $3.5\text{-}\mu\text{m}$ pixel pitch.

At 3 m, the focal length of the lens, given by Eq. (19.19), is approximately ten times longer or 100 mm since the only parameter in Eq. (19.19) that has changed is d_o . Similarly, we specify the same $F/4$ aperture since we are working with the same illumination wavelength and pixel pitch as in the 0.3 m system. We also note that in accord with the prediction of Eq. (19.16), Eq. (19.7) gives a very similar

DoF = 2.5 cm for the 3 m iris imaging system. Even if we allow for a larger circle of confusion equal to $3 \times L_p$ and close the aperture from $F/4$ to $F/5$, the DoF barely reaches 10 cm. A subject approaching the imaging system at 1 m/s passes through the DoF in 30–100 ms depending on the subjective definition of DoF that is applied. In any case, catching an in-focus image is a challenge for such a system. To address this challenge, we must either use a form of tracking autofocus or capture images fast enough to ensure that some subset of captured images are adequately focused for the purpose of iris recognition. Assuming that the imager can be run at a video rate of 30 frames per second, we obtain an image every 33 ms, which is adequate to capture a few in-focus images for a DoF of 3–10 cm and a 1 m/s approach speed.

However, we must also account of a transverse motion (perpendicular to the optical axis) of 0.1 m/s. At this speed, a subject traverses a distance of one L_{orp} in 1 ms. If we can apply a 1-ms optical pulse coordinated with a 1-ms camera shutter exposure, the transverse subject motion can be effectively frozen. The combination of short pulse and fast video frame rate is required for a system that operates at 3 m with a fixed capture volume and the reduced constraints of an approaching subject. An altogether different system which uses predictive tracking autofocus could address the challenges differently by translating the capture volume along with the subject while operating with continuous illumination and lower frame rate. The decision between a fixed capture volume with pulsed illumination and high video frame rate or tracking autofocus with lower frame rate and continuous illumination is clearly a difficult one with different manufacturers taking different approaches.

In the case of fixed capture volume, Eq. (19.15) predicts a low SNR with the same irradiance as the 0.3-m system and a short 1-ms exposure time. As a consequence, the fixed capture volume 3-m system must use a higher irradiance for a shorter time, consistent with the need to suppress motion smear and remain eye safe. Pressured to capture as many photoelectrons in the short exposure as possible, we seek a sensor with higher QE of 0.2, extend the exposure to 1.5 ms, and increase irradiance to 600 W/m^2 on the subject, ultimately yielding $\text{SNR} \sim 100$. Tuning such a camera to capture in-focus images of moving subjects demonstrates the engineering compromises of a 3-m iris imager design.

19.4 Epilogue: Requirements of an Iris Recognition System

So far, we have discussed principles of optics and designed iris imaging systems without reference to the segmenting (isolating the iris from a large field of view), encoding, and matching algorithms that perform the actual iris recognition functions. The design rules that we have used are photographic in nature and equally applicable to wedding photography as they are to aerial photography where requirements upon spatial resolution and SNR are deeply understood.

In this last section, we ask a series of questions without answering them: What are the actual spatial resolution and SNR requirements for iris recognition? How

do these requirements depend on the use case, for example, number of subjects in the database? How do they depend on the specific algorithm used to perform the matching? If the requirements are algorithm specific, is it advantageous to co-optimize the imaging system with the algorithm to achieve the best overall system performance? Or is a fixed standard to which all iris imaging systems must comply a better strategy? Certainly, application of standards affects iris imaging device manufacturers, iris matching algorithm developers, and the customers who buy iris recognition systems comprising both components so the question is not simply academic.

Questions of iris imaging system requirements divide the user community into two segments. In the first segment are iris recognition system users who rate either high performance or low cost as their top priority. A high-performance system, for example, would operate in a high-security environment with stringent standards upon failure to enroll and especially on false non-match rates. A low-cost system would, for example, enable cost-effective proliferation through a low-security infrastructure such as an enterprise time-in-attendance or a national entitlement system. A second segment of users places lifetime maintenance of an iris recognition system as the top priority. For the second set of users, multiple sourcing, interoperability, and forward/backward compatibility trump purchase cost and performance.

It is the second segment of users who drive the application of a single standard on iris imaging systems. For purchase of iris imaging systems across multiple branches of government, for example, the value of a single standard by which equipment is qualified ensures that competitive bidding, replacing, and upgrading can occur without reliance on a single supplier or a narrowly available system. By the same token, for cost-driven commercial installations, a single standard on iris imaging systems might preclude a low-cost alternative, with adequate performance that fits a budget. In this case, a single standard might keep a potential user from choosing iris recognition to fit their needs. A single standard on iris imagers might also reduce the incentive to develop a high-performance system based on an integrated sensor/algorithm approach.

Sections 19.2 and 19.3 of this Chapter provide the basic design tools needed to create an iris imager to meet the imaging requirements of any user. The future of iris imaging systems will be driven by the use cases, cost constraints, and algorithms that motivate their design. As iris imaging requirements evolve, the design principles outlined herein will provide the basis of understanding the optics of iris imaging systems.

References

1. Lefohn, A., et al.: An ocularist's approach to human iris synthesis. *IEEE Computer Graphics and Applications* **23**(6), 70–75 (2003)
2. Iris Image Quality Standard, SC 37 N 3331. ISO/IEC 29794-6 Annex A (2009)

3. Muron, A., Pospisil, J.: The human iris structure and its usages. *Acta Univ Plalcki Physica* **39**, 87–95 (2000)
4. ACGIH: TLVs and BEIs: Based on the Documentation of the Threshold Limit Values for Chemical Substances and Physical Agents and Biological Exposure Indices. American Conference of Governmental Industrial Hygienists. Cincinnati, OH, <http://www.acgih.org/Resources/press/TLV2005list.htm> (2005)
5. Trebino, R.: Diffraction and the Fourier Transform. [Online] www.physics.gatech.edu/frog/lectures (2008)
6. Born, M., Wolf, E.: *Principals of Optics*, vol. 7. Cambridge University Press, Cambridge (1999), see Section 8.3.3., Fraunhofer and Fresnel diffraction
7. Jackson, J.D.: *Classical Electrodynamics*, vol. 2. Wiley, New York (1975), See Section 9.8, Scalar Diffraction Theory
8. Zemax.: Software for Optical System Design. See <http://www.zemax.com/>
9. Koren, N.: Understanding image sharpness. Norman Koren Photography Page. (Online) <http://www.normankoren.com/Tutorials/MTF6.html> (2011)
10. Sarnoff, SRI International.: Iris On the Move (R). (Online) <http://www.sarnoff.com/demos/iris-on-the-move> (2011)
11. AOptix.: Insight VM. (Online) <http://www.aoptix.com/iris-recognition/product/insight-vm>.
12. Burge, M.J., Monaco, M.K.: Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XV. In: Shen, S.S., Lewis, P.E. (eds.) *Proceedings of SPIE*, Vol. 7334, pp. 73341D-1-73341D-8. SPIE, Bellingham, WA (2009)

Afterword

Since its introduction in 1992, the iris modality has earned acceptance as a highly accurate biometric. As evidenced in the collection of articles in this handbook, the iris itself continues to captivate the inquisitive minds of top researchers in the field of biometrics who seek to unearth all of the secrets this phenomenology has to offer. Probably more importantly, for those who have been leading the charge to promote its utility in a number of scenarios involving security, the iris remains able to boast extremely high rates of accuracy in large-scale national deployments.

Iris recognition systems generally accept for “match processing” images that have passed a predetermined quality threshold. The requirement for meeting this quality threshold has been instrumental in allowing the iris modality to build and preserve credibility as a highly accurate biometric. Image samples not meeting the required quality threshold are generally rejected by the iris recognition system – a result commonly referred to as a “failure to acquire.” In many deployments, the “failure to acquire” events are immediately detected and resolved by simply acquiring another image. This process can usually be implemented in such a way that it goes unnoticed by the individual being imaged – they just find themselves waiting a little longer for the success indicator, without necessarily understanding why.

In order to meet the demand to produce high-quality data (i.e., images) sufficient to pass the quality threshold, many of today’s iris acquisition systems require highly cooperative or conditioned individuals. In addition to the constraints placed on the individual whose iris is being acquired, the environment where the acquisition takes place must also be controlled or steps taken to mitigate the adverse effects introduced. That said, there have been some successes in the development of iris recognition systems (e.g., Iris on the move (IOM)) that have allowed for greater standoff distances of operation and freedom of movement by the individual being imaged. These advancements have mainly been limited to maximizing the performance of mainstream optical designs and sensor technologies through creative engineering of the acquisition platform to meet the desired signal quality of existing algorithms. This approach shifts a significant portion of the burden of conformance

from the individual onto the acquisition system. And while the performance and operation of these systems are quite impressive, the level of sophistication involved in the design has the adverse consequence of substantially increasing the cost to produce the acquisition system.

In 2008, I was invited to give a presentation entitled “The Science of Biometrics with Relaxed Constraints” at the 2nd Annual IEEE International Conference on Biometrics, Theory, and Systems. During this presentation, I discussed the possibility of exploiting features derived from the periorbital region of the eye to enhance the performance of iris recognition algorithms when confronted with data that is less than ideal (i.e., data of a quality level that falls just short of a predetermined threshold for processing by conventional algorithms). In an attempt to uphold the reputation of iris recognition for being a highly accurate biometric, and because the exploitation of all features in the vicinity of the eye was being encouraged, the term ocular recognition was used.

The proper implementation of an ocular recognition system would have at its core an iris recognition algorithm. In addition, the data quality requirements for enrollment in an iris system would not only be maintained but expanded to preserve information contained in the periorbital region of the eye. Methodologies employed to extract and subsequently match features derived from the periorbital region of the eye are only necessary when the data quality is insufficient for processing by conventional iris algorithms. Thus, ocular recognition – not to be confused with face recognition by parts – has the potential to be just as accurate as conventional iris recognition algorithms when processing high-quality data and yet more robust to minor degradations in signal quality that render them ineffective.

In the years following BTAS 2008, there have been a number of publications at mainstream IEEE conferences on biometrics addressing topics germane to ocular recognition. As is often the case, research progress by scientists on the leading edge of this area has been hampered by the dearth of data that exists. Much of the data that does exist to support investigations into the efficacy of the iris lacks sufficient detail in the periorbital region of the eye due (1) to cropping techniques designed to isolate the iris or (2) pixels in the periorbital region of the eye being overexposed due to the wavelength used to illuminate the iris and the reflectance properties of human skin. Alternatively, the ocular region was extracted from high-resolution photos of the face at visible wavelengths. These photos are not ideally suited for exploiting iris information.

Developments occurred to help to spur new innovations in the area of ocular recognition. First, the National Institute of Standards and Technology released an Ocular Challenge data set as part of its “Face and Ocular Challenge Series (FOCS)” in November of 2010. The stated goal of FOCS “is to engage the research community to develop robust face and ocular recognition algorithms along a broad front.” Second, NIST also recently published a report detailing an “Ocular and Iris Baseline Algorithm” (November, 2011). And finally, the Intelligence Advanced Research Projects Activity (IARPA) sponsored research efforts to investigate the utility of advanced ocular recognition algorithms to improve the robustness of biometric exploitation systems on suboptimal (iris) data, as well as next generation

sensor technology and innovative optical systems to relax the constraints on iris/ocular acquisition as part of its Biometrics Exploitation Science and Technology (BEST) program.

As I close, I would like to thank the authors for all of their efforts to assemble this outstanding handbook. One thing is for certain, there are a number of exciting questions that remain unanswered pertaining to understanding the discriminative power of the eye – and all of its complements. These questions, coupled with the quest to build ocular/iris acquisition systems that “adapt to the user, instead of requiring the user to adapt to the technology,” will continue to drive bold innovation in this area of biometrics for years to come.

Washington, DC, USA

Michael C. King

Index

A

- Acquisition module, 130
- Active contours without edges, iris segmentation, 297
 - foreground and background regions, 299–300
 - lookout regions, 300
 - technique/contour formulation, 300–301
- Active sectors, 111
- Authentic distribution, 210–211
- Average correlation energy (ACE), 341

B

- Bayesian fusion technique, 190
- Bayesian graphical methods
 - false alarm and false reject rates, 352–353
 - graphical model structure, 351
 - image patches, 349–350
 - Markov random field, 352
 - nonlinear deformation, 348, 350
 - occlusion, 348, 350
 - parameter learning, 352
- Bayesian graphical model, 312
- Bitwise hamming distance, 323, 328–329
- Blur estimation, 73–74
- Borda count, 229

C

- Cambridge (CAM-2-Q)
 - boxplots, 95–96
 - histograms of quality scores, 91, 92
 - Pearson and Spearman correlation coefficients, 91, 94
 - scatterplots, 93, 95

- Camera design
 - eye safety, 387
 - 0.3-m iris imaging system
 - diffraction limited lens, 389
 - iris photography, 389–390
 - ISO/IEC MTF standard, 388, 389
 - lens focal length, 388, 389
 - lens transmission, 389
 - pixel size, 388
 - quantum efficiency, 388
 - signal-to-noise ratio, 389–390
 - sports and astrophotography, 389
 - 3-m iris imaging system, 390–391
 - threshold level of irradiance, 387
- Canadian Border Services Agency, 41
- Cancelable iris biometrics, 185–186
- Cancelable iris templates
 - random matrix, 193
 - random permutation-based, 194–196
 - restricted isometry property, 192–193
- Cauchy-Euler equations, 135–136
- Central matching facility, 60
- Chimeric datasets, 219–220
- Chinese Academy of Sciences Institute of Automation (CASIA), 24
- Circular banding artifact, 4
- Clarity measure, 68
- COIR ocular collection, 310
- Compression
 - JPEG, 20, 61
 - standard iris storage formats
 - effect on matching accuracy, 62–63
 - formats, 61–62
 - guidelines, 64
- Confidence in scores (CS), 77–78, 80–81

- Consensual reflex, 38
- Cornea, 183
- Correlation filters (CF), 12–13
 - advantages, 353
 - average correlation energy, 341
 - Bayesian graphical methods
 - false alarm and false reject rates, 352–353
 - graphical model structure, 351
 - image patches, 349–350
 - Markov random field, 352
 - nonlinear deformation, 348, 350
 - occlusion, 348, 350
 - parameter learning, 352
 - correlation peak, 339
 - iris matching, 347–348
 - iris recognition method, 338
 - iris segmentation
 - CMU database, 345–346
 - cross-correlation-based circular boundary detection, 344
 - cross-correlation method, 346
 - Hough transform, 343–345
 - minimum average correlation energy filter, 341–342
 - minimum variance synthetic discriminant function filter, 342
 - near-infrared illumination levels, 338
 - optimal trade-off synthetic discriminant function filter design, 340–341
 - output noise variance, 341–342
 - similarity measure, probe and reference images, 339–340
 - texture-rich iris region, 338
- Correlation peak, 339
- Cosmetic contact lenses, 2
- Cryptographic techniques, 186
- CSCI, 190–191

- D**
- Daugman's integrodifferential operator, 289, 290
- Decision-level fusion, 221
- Density function, 107
- Detection error tradeoff (DET) curves
 - Casia 2.0, 122, 123
 - ICE1 iris database, 123, 124
 - MBGC, 125
- Dilation factor, 72
- Directional ray detection method
 - circle-fitting schemes, 301
 - description, 302
 - eye center detection, 303
 - eyelid boundary detection, 304–305
 - iterative approach, 304
 - pupil segmentation, 303
 - starburst algorithm, 302
 - variations/directional gradients, calculus of, 301
- Doddington zoo, 31, 40
- D-prime index, 77

- E**
- Eigen-iris approach, 25
- Elastic similarity metrics
 - binary feature maps
 - accumulative scores, 113–116
 - Daugman approach, 111–112
 - elastic similarity, 112–113
 - local feature maps, 116–120
- Encoding module, 130
- Equal error rate, 312
- Eye color, 94
- Eye safety, 387

- F**
- Face Ocular Challenge Series (FOCS) dataset, 12
 - characterization, 309
- COIR ocular collection, 310
- description, 310
- IARPA BEST program, 310
- IrisCodes, 310
 - Bayesian graphical model, 312
 - dynamic range, 311–314
 - equal error rate, 312
 - fusion method, 317, 318
 - Gabor wavelets, 316
 - iris-pupil and-sclera boundary, 311
 - keypoint descriptors, 314
 - local binary patterns, 315
 - performance, 311
 - segmentation examples, 311, 312
 - sped-up robust features, 314
- local discriminant analysis, 317
- periocular image
 - challenges, 286
 - eye center detection, 287–289, 306
 - illumination normalization scheme, 287, 288, 306
 - padding with zero pixels, 286, 287
 - principal component analysis, 317
 - 4SF@ algorithm, 316–317
- Failure to enroll (FTE) rates, 173, 180
- False match rate, 63

- False non-match rate, 63
 - False reject rate (FRR)
 - causes of, 212–215
 - short and long time lapse, 209–210
 - Feature extraction and encoding
 - Gabor function, 357–358
 - Hamming distance, 358
 - image segmentation, 357–358
 - match score, 358
 - Feature-level fusion, 221
 - Feedforward neural network (FFNN), 79
 - Field of view, 370–371
 - Finite element method (FEM), 140
 - Flutter shutter, 18
 - Fourier-based approximation, 259–261
 - Fragile iris code bits, 29
 - Fraunhofer approximation, 376
 - Fraunhofer diffraction, 377–378
 - Fusion of face and iris biometrics, 11
 - experiments
 - detection results, 231–232
 - fusion results, 233–235
 - image, LG4000 and IOM, 231
 - matching results, 232–233
 - iris on the move (IOM) sensor
 - design, 223
 - eye detection, 227
 - face detection, 226
 - face matching, 227
 - frames from, 224
 - fusion, 229–230
 - iris matching, 227–228
 - preprocessing, 225–226
- G**
- Gabor function
 - definition, 323–324
 - feature extraction and encoding, 357–358
 - and phase-steerable filter, 326–327
 - physical meaning, 326
 - GAC. *See* Geodesic active contours (GAC)
 - Gaussian distribution, 160, 162
 - Generic rectilinear format, 57–58
 - Geodesic active contours (GAC)
 - contour initialization, 253, 254, 294, 295
 - description, 292
 - embedding function, 251, 253–255, 292, 293, 296, 297
 - evolution of, 254–256, 297, 298
 - final contours, 255, 256
 - level sets, 251–253, 292
 - performance, 307
 - segmentation accuracy, 270
 - stopping function, 252, 253, 293–295
 - successful and failed segmentation outputs, 297, 299
 - Geometrical optics, iris imaging systems
 - depth of field, 372–374
 - field of view, 370–371
 - F-number, 371–372
 - focal length of lens, 369–370
 - lens aperture, 371–372
 - object-referred pixel size, 370
 - Gray-level co-occurrence matrices (GLCM), 26, 175–176
- H**
- Hamming distance
 - bitwise, 323, 328–329
 - feature extraction and encoding, 358
 - fractional, 162–164
 - genuine and imposter distributions, 78
 - IrisCodes, 323
 - iris spoofing algorithm, 364
 - Hard lenses, 2–3
 - Hill function, 133
 - Hough transform, 248–250, 289–292, 343–345
- I**
- Identity credential, 59
 - Illumination level, 71
 - Image compression, 20
 - Image hashing technique, 33
 - Image projection technique, 227, 228
 - Impostor distribution, 63
 - Integro-differential operator, 130, 246–248
 - Intelligence Advanced Research Projects Activity (IARPA), 310
 - Interlacing effect, 70–71
 - Interoperability, 56–57
 - Intra-eye color distribution, 177, 178
 - Iris biometrics research survey, 6–7
 - applications
 - fingerprint, iris, retina and DNA, 39
 - hardware implementations, 40
 - datasets and evaluations, 36–38
 - iris image acquisition
 - flutter shutter, 18
 - image compression, 20
 - issues, 17
 - non-ideal images and quality metrics, 18–20
 - wavefront-coded imagery, 17
 - iris region segmentation
 - MBGC portal videos, 23
 - phase congruency analysis, 21

- Iris biometrics research survey (*cont.*)
 - starburst method, 22
 - liveness detection, 35–36
 - multi-biometrics
 - BioSecure project, 31
 - ocular biometrics, 32
 - posterior union model, 30
 - performance, 38–39
 - privacy-enhancing techniques, 33–34
 - security, 34–35
 - texture coding and matching
 - CASIA v1 dataset, 24
 - eigen-iris approach, 25
 - fragile iris code bits, 29
 - iris analysis in parts, 27–28
 - sparse representation techniques, 29
 - speed iris matching, 28–29
 - texture analysis, 26–27
 - texture filter formulations, 25–26
 - theoretical analyses, 40–41
- Iris boundary characteristics
 - circularity of, 240
 - intensity variations, 240–241
 - nasal inclination of pupil, 240
- Iris Challenge Evaluation (ICE) 2006, 8
 - algorithm performance results, 90
 - data, 87
 - protocol, 88–89
 - quality scores
 - correlation study, 91
 - and demographics, 94–97
 - and performance, 97–101
 - right and left irises, 91–94
- IrisCodes, 12, 310
 - binomial distributions, 322
 - clustering process, 324–325
 - coded winning indexes, 325
 - filter-generating function, 325–327
 - FOCS dataset
 - Bayesian graphical model, 312
 - dynamic range, 311–314
 - equal error rate, 312
 - fusion method, 317, 318
 - Gabor wavelets, 316
 - iris-pupil and-sclera boundary, 311
 - keypoint descriptors, 314
 - local binary patterns, 315
 - performance, 311
 - segmentation examples, 311, 312
 - sped-up robust features, 314
 - Gabor filters, 322, 323
 - generalized frameworks, 329–330
 - hamming distance, 323
 - phase distance vs. bitwise hamming
 - distance, 328–329
 - precise phase representation
 - accuracy and speed, 333
 - CASIA-1 iris database, 331
 - equal error rate, 333
 - genuine and impostor distributions, 331, 332
 - ROC curves, 331, 332
 - properties, 322
 - speed and accuracy, 322
- Iris Exchange (IREX) evaluation, 58, 63
- Iris imaging systems (optics), 13–14
 - camera design
 - eye safety, 387
 - 0.3-m iris imaging system, 388–390
 - 3-m iris imaging system, 390–391
 - threshold level of irradiance, 387
 - description, 367
 - fidelity level, 368
 - high-performance systems, 392
 - near infrared illumination, 368
 - photographic concepts
 - geometrical optics, 369–374
 - spatial resolution, 374–382
 - trade-offs and constraints, 386–387
 - requirements, 391–392
 - single standard application, 392
- Iris matching, 347–348
- Iris on the move (IOM) sensor
 - design, 223
 - frames from, 224
 - fusion of face and iris biometrics
 - BordaIris and BordaFace method, 229
 - eye detection, 227
 - face detection, 226
 - face matching, 227
 - iris matching, 227–228
 - MinFace and MinIris approach, 229
 - preprocessing, 225–226
- Iris recognition
 - cancelability results
 - normalized hamming distance, 200, 201
 - ROC, 199, 200
 - cancelable iris biometric, 185–186
 - components, 184–185
 - ICE 2005 dataset, 199
 - MBGC videos, 201
 - method, 338
 - ND dataset, 196, 197
 - parts of eye, 183, 184
 - performance, 198–199
 - publicly available datasets, 187–188
 - SCI variation, 197–198

- sparsity-motivated selection
 - algorithm, 188
 - cancelable iris templates (*see* Cancelable iris templates)
 - recognition from video, 191
 - sector-based recognition, 189–191
 - sparsity concentration index, 189
- videos
 - ROC plots, 202
 - three gallery images, 186–187
- Iris segmentation
 - accuracy, 270
 - anterior surface anatomy, of human iris, 240
 - challenges
 - congenital iris abnormality, 245, 246
 - eye glass/contact lens, 244, 245
 - eyelids/eyelashes, occlusions due to, 241–242
 - image sensors, 245
 - motion blur, 243, 244
 - nature of interacting population, 245
 - off-angled iris images, 243, 244
 - poor illumination, 242, 243
 - specular reflections, 243
 - standoff distance, 244
 - user cooperation, 243–244
 - classical algorithms
 - Hough transform, 248–250, 343–345
 - integro-differential operator, 246–248
 - CMU database, 345–346
 - correlation filter approach, 306
 - cross-correlation-based circular boundary detection, 344
 - cross-correlation method, 346
 - definition, 239, 285
 - error prediction
 - decision tree-based machine learning approach, 275
 - geometric iris measure, 275
 - likelihood ratio test, 275
 - local image analysis, 274
 - Naive-Bayes tree classifier, 277
 - pupil segmentation measure, 275
 - scale parameter, 275
 - under-segmentation score, 276
 - Fourier-based approximation, 259–261
 - geodesic active contours, 250–256
 - image acquisition, 261
 - under constrained conditions, 262–265
 - IOM system, 265–267
 - public use iris recognition system, 262–265
 - user cooperation, 262
 - iris boundary characteristics
 - circularity of, 240
 - intensity variations, 240–241
 - nasal inclination of pupil, 240
 - nonideal periocular image, 283
 - over-/under-segmented limbus boundary, 270, 271
 - performance, 269–270
 - periocular image
 - acquisition factors, 283
 - active contours without edges, 297, 299–301
 - challenges, 283, 284
 - Daugman’s integrodifferential operator, 289, 290
 - directional ray detection method, 301–305
 - FOCS, 286–289
 - GAC (*see* Geodesic active contours (GAC))
 - Hough transform, 289–292
 - pupillary and limbus boundaries, 285
 - sample, 283
 - techniques, 285–286
 - periocular region
 - benefits, 282
 - description, 284
 - IOM systems, 282
 - primary task, 239
 - refinement
 - eyelash removal in iris images, 271–273
 - Libor Masek’s MATLAB package, 273, 274
 - roles, 285
 - segmentation accuracy, 306
 - variational level sets, 257–259
- in visible wavelength
 - color components, 268
 - image acquisition, 267, 268
 - multilayered perceptron feed forward neural networks, 269
 - neural pattern recognition approach, 268
 - proportion of sclera, 269
 - sclera map, 268–269
- Iris spoofing, 13
 - algorithm
 - Hamming distance, 364
 - ROC curves, 362–363
 - spooft pattern embedding in background texture, 361–362
 - spooft pattern generation, 359–361

Iris spoofing (*cont.*)
 feature extraction and encoding
 Gabor function, 357–358
 Hamming distance, 358
 image segmentation, 357–358
 match score, 358
 multiresolution reverse subdivision
 technique, 356
 spoof iris codes, 356–357
 synthetic iris texture generation, 356

Iritech (IRTCH-2-Q)
 boxplots, 95–96
 histograms of quality scores, 91, 92
 Pearson and Spearman correlation
 coefficients, 91, 94
 scatterplots, 93, 95

ISO/IEC standards, 57

Isotropic deformation, 140, 143–145

J

JPEG compression, 20, 61

K

Kendall's correlation coefficient, 213

Keypoint descriptors, 314

KIND 7 iris record, 60, 62–65

L

Less constraining acquisition frameworks
 degradation factors, 159
 image assessments
 coarse approximation function, 157
 entropy, 158–159
 sensitivity
 decidability index, 162
 fractional hamming distances, 162,
 163
 match and non-match distributions,
 164
 spectral absorbance, 165, 166
 specificity, 159–161

LG EOU 2200, 87

Libor Masek's MATLAB package, 273, 274

Lighting factor, 71

Liveness detection, 35–36

Local binary patterns (LBP)
 liveness detection, 35
 ocular biometrics, 32
 periocular recognition, 315

Longtin's model, 133

Lossy compression, 61

M

Mahalanobis cosine metric, 227

Markov random field (MRF), 352

Matching module, 130

Match variability, 155

MEC. *See* Multispectral Enhanced irisCode (MEC)

Melanin, 172, 177, 215

Merck Manual of Geriatrics, 215

Meshwork skeleton, 134–135

Minimum average correlation energy (MACE)
 filter, 341–342

Minimum variance synthetic discriminant
 function (MVSDF) filter, 342

Modulation transfer function (MTF), 381

Motion blur, 243

Multi-biometrics
 characteristics of, 220–221
 description, 219–220
 levels of fusion, 221–222

Multiple Biometrics Grand Challenge
 (MBGC), 23, 37, 222–223

Multispectral Enhanced irisCode (MEC), 179,
 180

Multispectral iris fusion, 10
 cameras
 multispectral acquisition system, 173,
 175
 spatially registered four-channel iris
 image, 173, 174
 MEC, 179, 180
 modeling NIR iris images, 177–179
 texture, 175–177

Multistage supervised learning process, 177

Multivariate adaptive mapping, 76

N

Naive-Bayes tree classifier, 277

Nearest neighbor-based recognition algorithm
 (NN), 196, 198

Network-centric application, 59

NIR illumination, 152

Noisy Iris Challenge Evaluation (NICE), 9, 21

Non-ideal images, 18–20

Non-match variability, 155

Normalization module, 130, 131

O

Object-referred pixel size, 370

Occlusions
 Bayesian graphical methods, 348, 350
 due to eyelashes, 241–242

due to eyelids, 241
 metric, 72
 Ocular biometrics, 32
 Off-angled iris images, 243
 Off-angle quality factor, 73
 Open CV, 23, 226
 Optimal trade-off synthetic discriminant
 function (OTSDF) filter design,
 340–341
 Orthotropic deformation, 138–143

P

Pan-and-tilt device, 156
 Pattern recognition approaches, 134
 Peak-to-correlation energy (PCE), 347
 Pearson and Spearman correlation coefficients,
 91, 94
 Periocular biometrics, 167
 Periocular recognition. *See* Face Ocular
 Challenge Series (FOCS) dataset
 Persistent pupillary membrane, 5, 6
 Personalized iris matching strategy, 28
 Phase-based iris recognition algorithms, 104
 Phase congruency analysis, 21
 Phase-only correlation (POC) function, 26
 Phasors, 172
 Photographic concepts, iris imaging systems
 geometrical optics, 369–374
 spatial resolution, 374–382
 trade-offs and constraints, 386–387
 Pixel count, 72
 Point spread function (PSF), 18, 375, 382
 Poisson's ratio, 136–137
 Polar storage format, 57–58
 Posterior union model (PUM), 30
 Precise phase representation algorithm
 accuracy and speed, 333
 CASIA-1 iris database, 331
 equal error rate, 333
 genuine and impostor distributions, 331,
 332
 ROC curves, 331, 332
 Privacy-enhancing techniques, 33
 Proportion of sclera, 269
 Public use iris recognition system, 262–265
 Pupil dilation, 132–135
 Pupillary latency, 133

Q

Q-stack classifiers, 68
 QST method, 81–82
 Quality factors, 8

adaptive authentication
 confidence in scores, 77–78
 experimental results, 79–82
 multivariate regression problem, 76
 quality of sample, 77
 quality sample and template features,
 78–79
 blur, 73–74
 dilation, 72
 ICE 2005 dataset, ROC curves for, 74–76
 illumination, 71
 interlacing, 70–71
 iris segmentation scores, 69–70
 lighting, 71
 occlusion, 72
 off-angle, 73
 pixel count, 72
 Q-stack classifiers, 68
 Quality metrics, 18–20
 Quality of sample (QS) evaluation, 77, 80
 Quality scores, ICE 2006
 correlation study, 91
 and demographics, 94–97
 and performance, 97–101
 right and left irises, 91–94

R

Random permutation-based cancelable iris
 biometrics, 194–196
 Rank-level fusion, 221
 Rank one recognition rates, fusion approaches,
 233–234
 Receiver operating characteristic (ROC) curves
 fusion methods, 234
 ICE2005 dataset, 74–76
 iris spoofing algorithm, 362–363
 ND dataset, 200
 spoof iris texture, 13
 Restricted isometry property (RIP), 192–193
 ROC curves. *See* Receiver operating
 characteristic (ROC) curves
 Rubber-sheet model, 148

S

Sagem-Iridian (SI-2-Q)
 boxplots, 95–96
 histograms of quality scores, 91, 92
 Pearson and Spearman correlation
 coefficients, 91, 94
 scatterplots, 93, 95
 Scale invariant feature transform (SIFT), 222
 Sclera, 183

- Score-level fusion, 221, 222
 - Segmentation. *See also* Iris segmentation
 - accuracy, 306
 - module, 130
 - process, 104–105
 - scores, 69–70
 - Signal-level fusion, 221
 - Sign test, worse mean match score, 211
 - Single-lens multispectral camera, 173–175
 - Smart cards, 59
 - Sobel edge filter, 271, 272
 - Soft contacts, 2
 - Sparsity concentration index, 189
 - Sparsity-motivated selection and iris recognition
 - algorithm, 188
 - cancelable iris templates
 - random matrix, 193
 - random permutation-based, 194–196
 - restricted isometry property, 192–193
 - recognition from video, 191
 - sector-based recognition, 189–191
 - sparsity concentration index, 189
 - Spatial resolution, iris imaging systems
 - description, 374
 - diffraction
 - blur spots, 379, 380
 - consequences, 377–378
 - different lens apertures, 378
 - F-number, 379, 380
 - Fraunhofer diffraction, 377–378
 - limit, 379
 - pixel dimension, 379, 380
 - sinusoidal grating, 376–377
 - of water waves, 375
 - lens quality
 - aberration-limited resolution, 380, 381
 - defocus broadening, 382
 - Gaussian profile, 382
 - modulation transfer function, 381–382
 - ray-tracing programs, 380
 - point spread function, 375
 - Spectrally sampled structural subspace features
 - algorithm, 316–317
 - Spectral reflectance and radiance, iris, 153
 - Sped-up robust features (SURF), 314
 - Speed iris matching, 28–29
 - Standard iris storage formats, 7
 - compression
 - effect on matching accuracy, 62–63
 - formats, 61–62
 - guidelines, 64
 - generic rectilinear format, 57–58
 - KIND 7, 60, 62–65
 - KIND 3 and KIND 16, 60
 - polar storage format, 57–58
 - role and importance of, 56–57
 - storage size constraints, 59–60
 - Starburst algorithm, 302
 - Starburst method, 22
 - Stereo rig multispectral camera, 173–175
- T**
- Taylor expansion, 8–9
 - definition, 106
 - iris features
 - binary features, 108–110
 - density function, 107
 - derivatives and averaging parameters, 108
 - elastic similarity metric (*see* Elastic similarity metrics)
 - estimation technique, 106
 - local extrema, 110–111
 - MBGC-31 dataset, 120
 - phase-based iris recognition algorithms, 104
 - segmentation process, 104–105
 - template size, 121
 - verification performance, 122–125
- Template aging, 10–11
 - description, 205–206
 - false reject rates
 - causes of, 212–215
 - short and long time lapse, 209–210
 - image data set and algorithms, 207–209
 - iris matching algorithms, 209
 - lifetime enrollment, 206
 - time lapse effect, 207
- Template-matching approach, 226, 227
- Tissue dynamics, 9
 - advantages, iris unwrapping, 130–131
 - deformation, 145–149
 - Hill function, 133
 - image-based model, 134
 - integro-differential operator, 130
 - mathematical formulation
 - boundary condition, 137
 - Cauchy-Euler equations, 135–136
 - isotropic deformation, 140
 - orthotropic deformation, 138–140
 - Poisson's ratio, 136–137
 - meshwork skeleton, 134–135
 - numerical results and simulation
 - FEM, 140
 - isotropic deformation, 143–145
 - orthotropic deformation, 141–143

- parametric model, 133
- pupil dilation, 132, 133

U

- UBIRIS.v2, 154–155
- Unwrapping, iris, 130–131
- User cooperation, 243–244

V

- VeriEye algorithm, 209, 210
- Viola-Jones cascade face detector, 226
- Visible wavelength (VW), 9–10
 - feasibility, 153
 - iris segmentation
 - color components, 268
 - image acquisition, 267, 268
 - multilayered perceptron feed forward neural networks, 269
 - neural pattern recognition approach, 268

- proportion of sclera, 269
- sclera map, 268–269
- less constrained iris recognition, 153–154
- less constraining acquisition frameworks
 - degradation factors, 159
 - image assessments, 157–159
 - sensitivity, 161–166
 - specificity, 159–161
 - UBIRIS.v2 database, 156
- NIR illumination, 152
- periocular biometrics, 167
- spectral reflectance and radiance, 153
- UBIRIS.v2, 154–155

W

- Wedgelet, 27–28
- Worse mean score, 210–211

Z

- Z-score normalization, 317, 318