IEEE Signal Processing MAGAZINE

[VOLUME 32 NUMBER 5 SEPTEMBER 2015]

BIOMETRICS SECURITY AND PRIVACY PROTECTION

RECENT ADVANCES

SIGNAL PROCESSING-DRIVEN IMAGING TECHNOLOGIES

COMMEMORATING STEREO SOUND RECORDING

CREATING ANALYTIC DSP ONLINE HOMEWORK

IEEE Signal Processing Society

IEEE

Ultra Small 2x2mm

2W ATTENUATORS DC-20GHz from $1.99 ea. (qty. 1000)

Save PC board space with our new tiny 2W fixed value absorptive attenuators, available in molded plastic or high-rel hermetic nitrogen-filled ceramic packages. They are perfect building blocks, reducing effects of mismatches, harmonics, and intermodulation, improving isolation, and meeting other circuit level requirements. These units will deliver the precise attenuation you need, and are stocked in 1-dB steps from 0 to 10 dB, and 12, 15, 20 and 30 dB.

The ceramic hermetic ***RCAT*** family is built to deliver reliable, repeatable performance from DC-20GHz under the harshest conditions. With prices starting at only $4.95 ea. (qty. 20), these units are qualified to meet MIL requirements including vibration, PIND, thermal shock, gross and fine leak and more, at up to 125°C!

The molded plastic ***YAT*** family uses an industry proven, high thermal conductivity case and has excellent electrical performance over the frequency range of DC to 18 GHz, for prices starting at $2.99 ea. (qty. 20).

For more details, just go to minicircuits.com – place your order today, and you can have these products in your hands as soon as tomorrow!

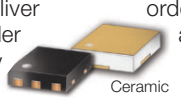

RoHS compliant

FREE Simulation Models! Modelithics®

http://www.modelithics.com/mvp/Mini-Circuits/

Mini-Circuits®

www.minicircuits.com P.O. Box 350166, Brooklyn, NY 11235-0003 (718) 934-4500 sales@minicircuits.com

515 rev E

CONTENTS

[VOLUME 32 NUMBER 5]

SPECIAL SECTION—BIOMETRICS SECURITY AND PRIVACY PROTECTION

Digital Object Identifier 10.1109/MSP.2015.2444511

COLUMNS

DEPARTMENT

IEEE SIGNAL PROCESSING magazine

COVER

©ISTOCKPHOTO.COM/RAWPIXEL
AND EYE INSET: ©ISTOCKPHOTO.COM/NEVARPP

IEEE SIGNAL PROCESSING MAGAZINE (ISSN 1053-5888) (ISPREG) is published bimonthly by the Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, 17th Floor, New York, NY 10016-5997 USA (+1 212 419 7900). Responsibility for the contents rests upon the authors and not the IEEE, the Society, or its members. Annual member subscriptions included in Society fee. Nonmember subscriptions available upon request. Individual copies: IEEE Members US$20.00 (first copy only), nonmembers US$213.00 per copy. Periodicals postage paid at New York, NY, and at additional mailing offices. Postmaster: Send address changes to IEEE Signal Processing Magazine, IEEE, 445 Hoes Lane, Piscataway, NJ 08854 USA. Canadian GST #125634188

Printed in the U.S.A.

Digital Object Identifier 10.1109/MSP.2015.2444512

SCOPE: *IEEE Signal Processing Magazine* publishes tutorial-style articles on signal processing research and applications, as well as columns and forums on issues of interest. Its coverage ranges from fundamental principles to practical implementation, reflecting the multidimensional facets of interests and concerns of the community. Its mission is to bring up-to-date, emerging and active technical developments, issues, and events to the research, educational, and professional communities. It is also the main Society communication platform addressing important issues concerning all members.

[from the EDITOR]

Min Wu
Editor-in-Chief
minwu@umd.edu

Is Signal Processing a New Literacy?

It is summer as I write this editorial. I am leaving soon for Chengdu, a southwestern city of China known as the Land of Abundance, to attend the third edition of the IEEE China Summit and International Conference on Signal and Information Processing (ChinaSIP). This conference was the IEEE Signal Processing Society's (SPS's) first major outreach initiative to the emerging economies where there has been a large and growing base of signal processing professionals. ChinaSIP was envisioned to help colleagues in China engage with the global community and to offer global colleagues opportunities to network and develop international collaborations.

With community building and member outreach as a primary mission, ChinaSIP has explored ways to bring exposure and benefit to attendees, for example, through a series of panel discussions on career development, education, and industry perspectives. I was fortunate to be entrusted by the SPS boards and leadership to chair the ChinaSIP Steering Committee to help shape the early effort and to work with the organizing teams and numerous volunteers whose tremendous contributions have led to successful events. It is great to learn that SPS is working on leveraging the experiences from ChinaSIP and extending the community building to Asia and other areas in the coming years.

Speaking of outreach of a different kind, I also attended an interesting workshop that brought together university and industrial researchers of a diverse range of engineering backgrounds. A keynote speech given by Jim Tung, a MathWorks fellow, discussed building a new literacy. Jim shared with the audience his visits to South Korea, where he learned that the government has been making a strong effort to make computer programming/coding part of the required middle school curriculum.

At that moment of the talk, we would have thought of computer coding as a new literacy. Not surprisingly, this is in line with the computer science community's effort to bring computer programming to elementary and middle schools and to the general public (as noted in [1]). Introductory computer coding basically teaches logical thinking and expresses a solution into logical steps using the syntax of a particular computing language or platform. However, there appeared to be something missing in the coding advocacy. Through an outreach activity at an elite high school, I interacted with the lead computer science teacher, who seemed to care mostly about teaching students a particular language (be it Java or Python) and building a web page. It is common at this level that programming is treated fashionably but separate from other traditional subjects such as math, physics, and biology.

While I value the early exposure of programming (and having started mine during my elementary school years), for most students and the general public, a computer is primarily a tool to accomplish tasks and solve problems. To solve problems, we need ideas and methods, which often need to be built on top of solid foundations as well as domain knowledge. Programming then comes as an expression or implementation of those ideas and methods in a specific language or platform. This resonated very well with Jim's keynote talk, as he walked the audience through the proliferations of algorithms in almost every gadget and system we have today (such as in cars) and noted that we use algorithms to help us understand many things that surround us.

The punch line of "creating algorithms is the new literacy" in Jim's talk brought us an insightful synergistic connection between computing and other science, technology, engineering, and mathematics (STEM) foundations. Developing ideas and methods requires us to model the problem with proper angles and abstraction, apply mathematical or scientific principles, and develop an effective way to solve the problem. The resulting algorithms will be expressed or implemented using a suitable programming language or higher-level tools and platforms that encapsulate some detailed tasks into convenient modules and packages. The hard work of learning math and science at grade school levels as well as the many foundation courses in engineering and science colleges would no longer be seen as out of fashion. In fact, they can and should be as fashionable and essential as coding, if not more.

Signal processing sits right at the intersection of a number of STEM disciplines. Traditionally, it arose from physics and mathematics, and it is increasingly connecting with a broad variety of disciplines, such as biomedicine, mechanical engineering, civil engineering, material science, social science, and, of course, computing. Modeling and algorithms have been essential in signal processing as are practical implementations involving both hardware and software. So I wonder, with its bridging capabilities, would signal processing be a new literacy?

REFERENCE

[1] M. Wu, "Sharing signal processing with the world," *IEEE Signal Processing Mag.*, vol. 32, no. 2, p. 4, Mar. 2015.

[SP]

Digital Object Identifier 10.1109/MSP.2015.2449285
Date of publication: 13 August 2015

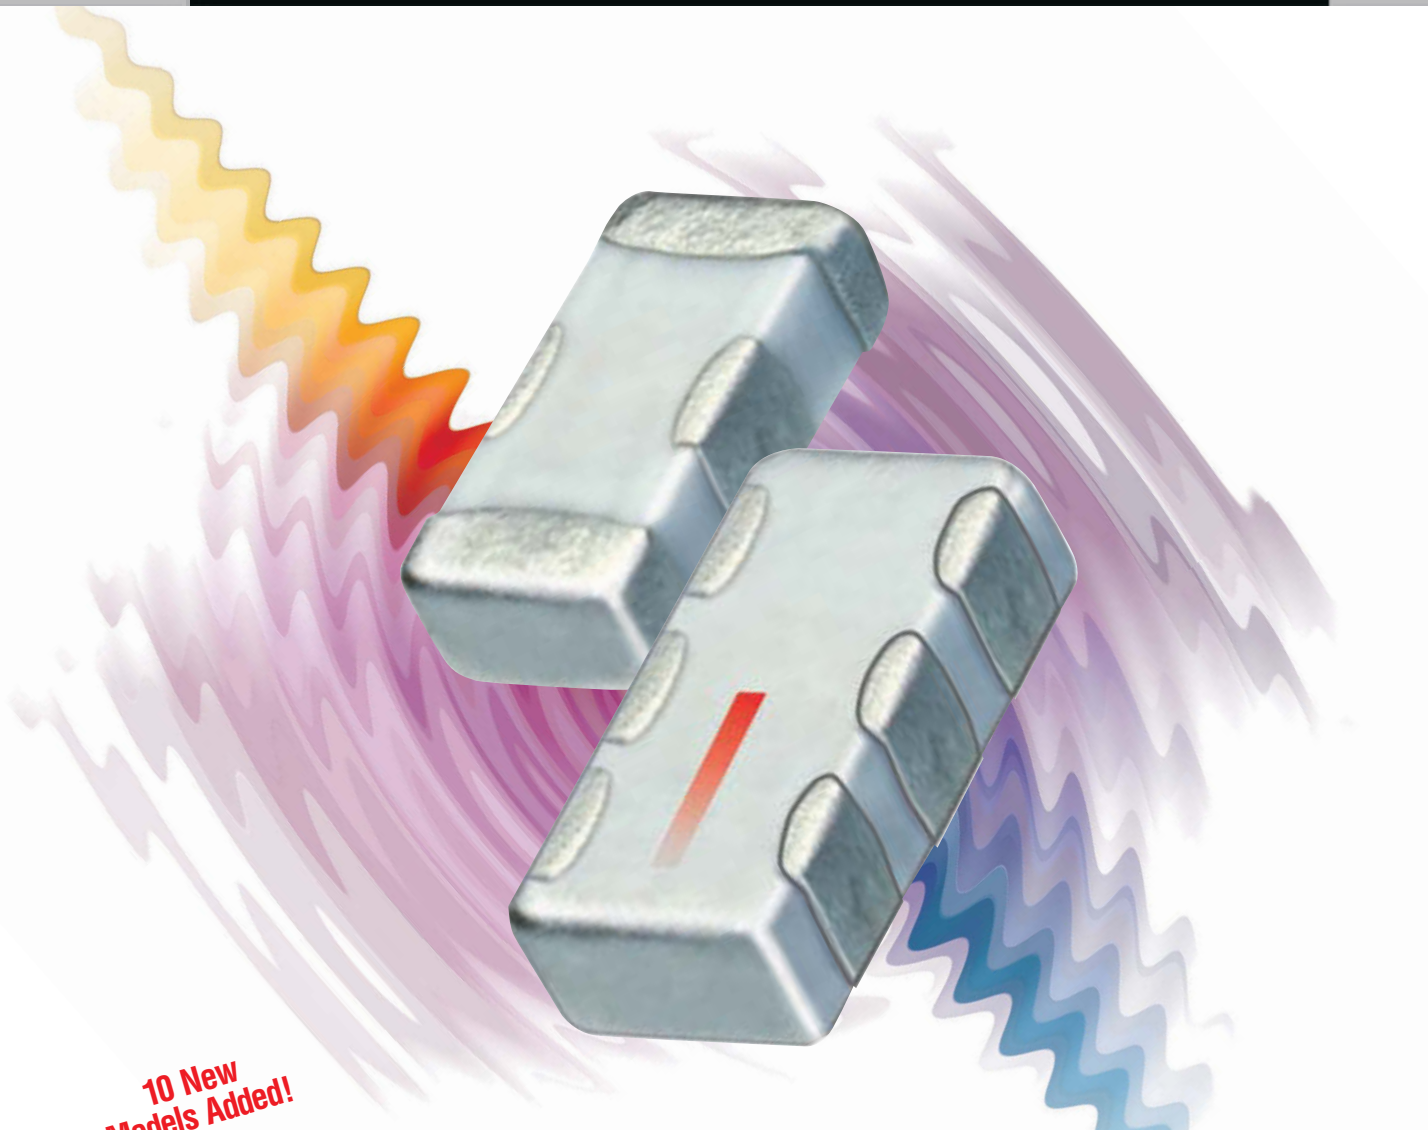

10 New Models Added!

CERAMIC FILTERS

LOW PASS BANDPASS HIGH PASS DIPLEXERS

Covering DC to 18.3 GHz *from* 99¢ *ea. qty.3000*

Over 200 models as small as 0.06 x 0.03"! These tiny, hermetically sealed filters utilize our advanced Low Temperature Co-fired Ceramic (LTCC) technology to offer superior thermal stability, high reliability, and very low cost. Supporting a wide range of applications with high stop band rejection and low pass band insertion loss in tiny packages, they're a perfect fit for your system requirements. Visit minicircuits.com for comprehensive data sheets, PCB layouts, free high-accuracy simulation models, and everything you need to choose the model for your needs. Order direct from our web store, and have them in your hands as soon as tomorrow!

Now available in small-quantity reels *at no extra charge:*
Standard counts of 20, 50, 100, 200, 500, 1000 or 3000.
Save time, money, and inventory space!

Wild Card Filter Kits, *KWC-LHP, only* $98

- *Choose any 8 LFCN or HFCN models*
- *Receive 5 of each model*
- *A total of 40 filters for a great value*
- *Order your KWC-LHP Filter Kit TODAY!*

RoHS compliant U.S. Patents 7,760,485 and 6,943,646

Free, High-Accuracy Simulation Models for ADS
Modelithics®
www.modelithics.com/mvp/Mini-Circuits.asp

Mini-Circuits®

www.minicircuits.com P.O. Box 350166, Brooklyn, NY 11235-0003 (718) 934-4500 sales@minicircuits.com

504 Rev F

president's **MESSAGE**

Alex Acero
2014–2015 SPS President
a.acero@ieee.org

Should We Experiment with New Peer-Review Models?

My children cannot understand that I didn't have a cell phone growing up, that the first TV we had at home was a black-and-white 19-inch set without a remote, or that I had to drop off film to be developed before I could see my pictures. Technology has changed our lives profoundly.

Technology has also changed the world of scholarly papers. I wrote essays in college using a typewriter and plenty of Wite Out. In 1990, my first International Conference on Acoustics, Speech, and Signal Processing (ICASSP) paper was written with the help of a workstation, but I had to print it on paper, cut it, glue it onto a double-sided mat, and mail it between two pieces of cardboard. As Publications chair for ICASSP 1998, I set up a system for authors to make electronic submissions for the first time. As recently as 2001, I had to go to IEEE's headquarters to review ICASSP papers, whereas now we do all paper reviews from home via a website. Today, anyone can make the look and feel of conference papers as good as that of journal articles using easily available software, which was not the case in the 1980s. My office shelves were full of thick conference proceedings and lots of journals, but I stopped getting the paper versions over a decade ago.

Technology has dramatically altered the way authors write manuscripts, the way reviewers review them, and how readers consume journal articles and conference papers. Although we now have web tools to make the process easier, conferences and journals in our Society follow the same peer-review model they've had since the middle of the 20th century. Journal editors skim the submission, choose a few reviewers and give them a deadline, then look at the reviews and make a decision. The decision is to 1) accept unconditionally, 2) accept with minor revisions, 3) ask authors to revise and resubmit for another round of reviews, or 4) reject outright. Conferences skip option 3) since they have a more rigid schedule, as well as journals like *IEEE Signal Processing Letters* that want a faster turnaround.

One of the most interesting aspects of the web 2.0 revolution is social media and its impact on web pages. Many blogs allow people to add comments and "likes," which can influence authors to update their article in a matter of days or even hours. But, by and large, peer reviews of scholarly papers haven't adopted any of those changes.

Some conferences in machine learning are posting submissions on a website and allowing readers registered with their Google Scholar account to write comments visible to everyone. Authors can then improve the paper rapidly since they get more comments and get them more quickly than in traditional reviews. Program committees then use those comments as well as blind review feedback for their acceptance decisions. Some researchers want conferences to publish not only the accepted papers but also the rejected submissions, together with paper downloads/views, likes, and comments.

Conference organizers typically follow an iterative process to match the list of reviewers to the submissions, which can take a long time when the number of papers is large. I've heard many editorial boards saying how hard it is to find good reviewers and how reviewer fatigue makes it difficult with the ever-increasing number of conferences and journals.

Some machine-learning conferences let reviewers see the title and abstract for all submissions in an area and then bid on a few. Then a machine-learning system recommends assignments using features such as reviewers' bids, similarity between the text of each submission and the text of reviewers' published papers, and common citations. Some conferences ask authors to provide scores for each reviewer, which can provide an indication of reviewer quality after proper normalization. Other ideas include a credit system where an author of a published paper owes credits to a journal or conference and has to pay those credits by agreeing to review papers. Program committees still control these decisions, but these technologies can play an important role in the peer-review process.

Several conference technical chairs and journal editors tell me how happy they are with these changes, but I also hear from researchers in our field that feel there is potential for gaming and manipulation in open review and social media. Shall we experiment with a new form of peer-review system that could supplement our current review process? And if so, who would volunteer to assist us in such experimentation?

[SP]

Digital Object Identifier 10.1109/MSP.2015.2449286
Date of publication: 13 August 2015

[special REPORTS]

John Edwards

Signal Processing Opens New Views on Imaging

The importance of signal processing in imaging is growing rapidly as technologies continue to develop and mature and as various fields begin to recognize the value of innovative new imaging and image analysis systems. By enabling people to clearly observe and detect things that are not ordinarily visible or not readily apparent to the unaided eye, signal processing-driven imaging technologies are helping to save lives and property from hazards lurking both on the ground and below the earth's surface.

TURNING THE CORNER

Researchers often claim that they have "turned the corner" whenever they make a major advancement or reach a new level of understanding in their work. Yet, scientists at the University of Bonn and the University of British Columbia can uniquely claim that they have developed an imaging technology that is specifically designed to turn its field of view around corners to see objects that would otherwise remain hidden from view. Their new camera system is designed to see around bends and turns without any help from a mirror. Using diffusely reflected light, the new camera can reconstruct the shape of objects located far outside of the field of view of human eyes and ordinary cameras.

"Being able to look around corners can potentially benefit a range of applications, from traffic safety to search and rescue operations," says Matthias B. Hullin, a professor at the University of Bonn's Institute for Computer Science. To image objects that are blocked to human viewers by walls or other physical obstacles, the camera shines a laser dot on a nearby vertical surface and then records both the direction of the light and the time it takes to reach the camera's imaging sensor (Figure 1). As the image is processed and reprocessed, the outline of the object the camera is trying to see becomes apparent to the viewer [see (Figure 2(a) and (b)].

SIGNAL PROCESSING-DRIVEN IMAGING TECHNOLOGIES ARE HELPING TO SAVE LIVES AND PROPERTY FROM HAZARDS LURKING BOTH ON THE GROUND AND BELOW THE EARTH'S SURFACE.

"We are recording what is essentially a light echo, and the time-resolved data allows us to reconstruct the object," Hullin explains. "Part of the light has also come into contact with the unknown object, carrying with it essential information about its shape and appearance."

The system takes advantage of multipath scattering interference. The approach is directly opposite to standard practice in telecommunications and remote sensing, where multipath in most applications is considered an annoyance rather than a source of additional data. "The key insight from our work is that, using accurate models for multipath scattering, a lot of information...can be

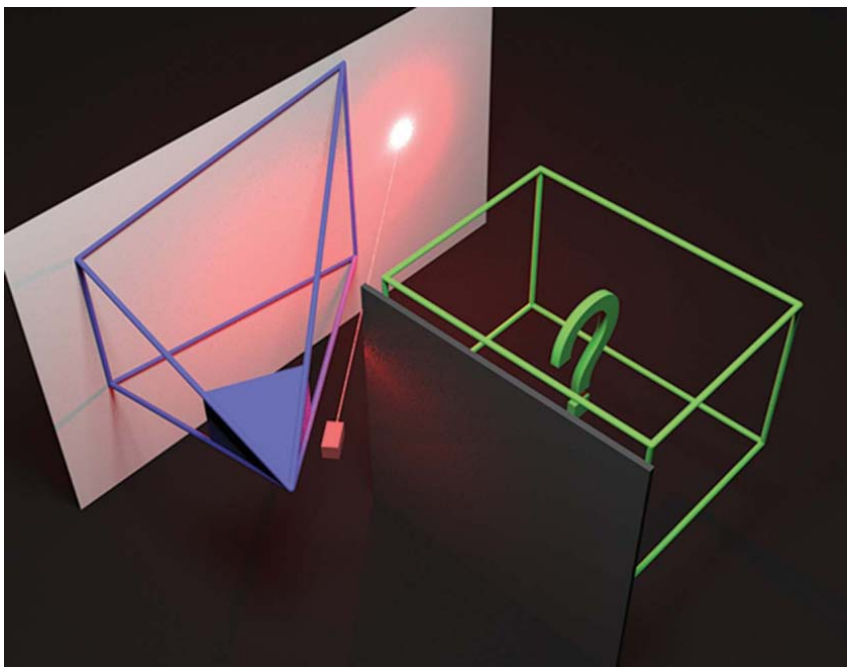

MATTHIAS B. HULLIN, UNIVERSITY OF BONN.

[FIG1] A rendering of the imaging scenario. The light source (red) and camera (field of view marked in blue) both look at a white wall. The objects to be captured in a reconstructed image (green box) are hidden behind an occluder and are only accessible through indirect reflections off of the diffuse wall.

Digital Object Identifier 10.1109/MSP.2015.2437291
Date of publication: 13 August 2015

1053-5888/15©2015IEEE

unlocked that is not contained in the direct reflection," Hullin says.

The method, Hullin explains, is based on space-time impulse responses of light. "Imagers with sub-nanosecond temporal resolution...are rare, expensive, and unpractical, so experimental results were out of reach until we found a way to circumvent the need for high-end gear."

To circumvent the cost problem, the researchers settled on a type of camera that long ago entered the mass market: a Creative Senz3D, which currently sells for about US$115. The "time-of-flight" camera resolves distance based on the known speed of light, measuring the time-of-flight of a light signal between the camera's sensor and the subject for each point of the image.

Yet, before the inexpensive camera could begin imaging around corners, it required significant modification. "It is surprisingly challenging to find research-grade hardware that offers free choice of modulation frequencies and raw data readout," Hullin says. "To construct our imaging system, we therefore had to hack into a time-of-flight development board, literally drilling into the sensor chip to combine it with a custom signal generator."

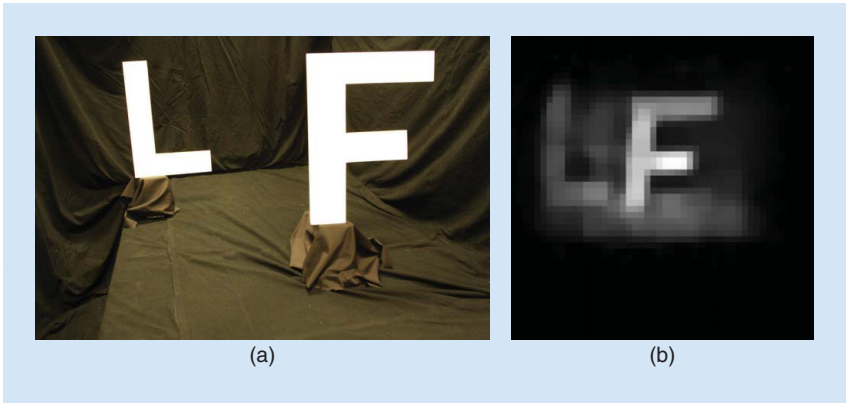

MATTHIAS B. HULLIN, UNIVERSITY OF BONN.

[FIG2] A regular photograph of cardboard cutout letters and the reconstructed indirect image captured by the modified Creative Senz3D camera.

Signal processing is central to the project. "Light transport and image sensors are linear time-invariant systems, so we can draw from a rich pool of signal processing methodologies to deal with the data at hand," Hullin remarks. "Nothing we do would be possible without the help of compressed sensing."

Hullin notes that the challenge of imaging around corners is not unlike trying to enhance a poorly focused conventional photograph. "In many ways, what we want to achieve is closely related to the problem of deblurring an image," he says. "Once certain

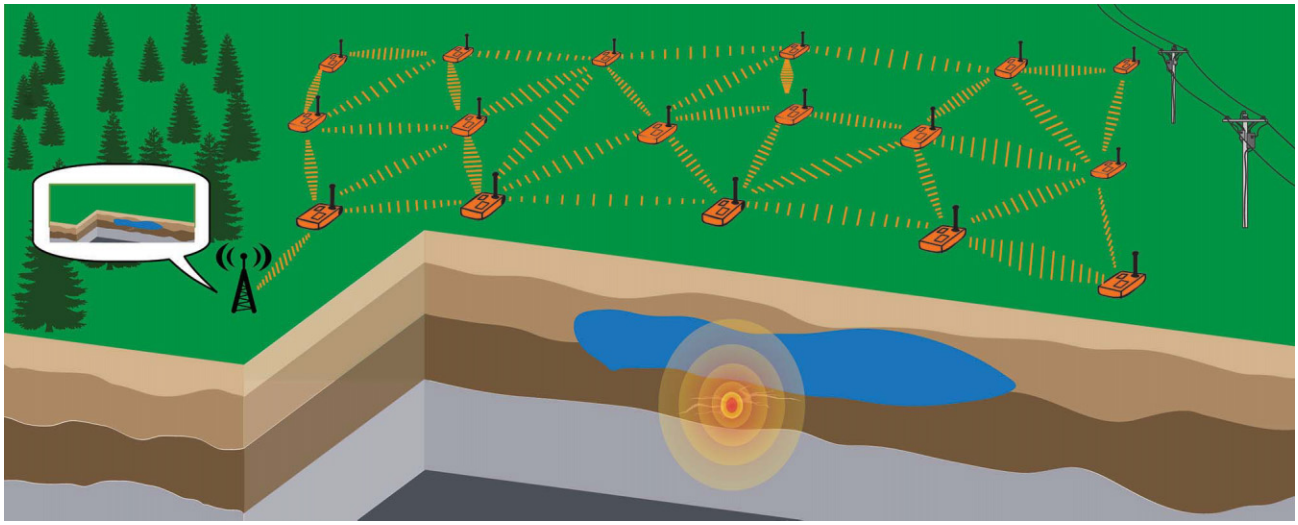

INTELLIGENT DOTS.

[FIG3] An RISI seismic imaging system in a passive layout configuration.

frequency bands have been destroyed by blur, plausible reconstructions can only be obtained by making additional assumptions about the signal we are interested in."

Image deblurring is often formulated as a least-squares optimization problem regularized with a total-variation gradient penalty to avoid unnecessary structure within the image. Hullin says his team models hidden scenes as a volumetric distribution of scattering densities. Then, using a combination of sparsity priors, they constrain the solution space to solutions that are zero almost everywhere except at object surfaces. "In other words, opaque objects are preferred over 'cloud-like' ones," Hullin says. "As is often the case when using compressed sensing methods, we found development of good prior models the most challenging task," he adds.

"The accuracy of our method has its limits," Hullin admits. Imaging remains restricted to rough outlines. Yet, the researchers believe that the rapid development of technical components and mathematical models will soon allow a higher resolution to be achieved. A similar around-the-corner imaging technology developed at the Massachusetts Institute of Technology (MIT), using direct reflection of lasers creates similar results: blurry, yet discernible images of objects from outside the camera's field of view.

Another challenge, Hullin says, is making the system real-time capable. "So far, we need to capture hundreds of individual input images and the reconstruction takes hours to complete," he says. In an early prototype, a student has been able to demonstrate that objects can be detected and tracked around a corner within a few milliseconds per frame.

HULLIN NOTES THAT THE CHALLENGE OF IMAGING AROUND CORNERS IS NOT UNLIKE TRYING TO ENHANCE A POORLY FOCUSED CONVENTIONAL PHOTOGRAPH.

The new imaging technology is already generating interest from various commercial and government organizations. "We are in touch with people from various industries including automotive, defense, and entertainment," Hullin says.

IMAGING HIDDEN DANGERS

Wen Zhan Song, a Georgia State University computer science professor, wants to get a picture of what is happening underground with the help of a real-time seismic imaging system that uses ambient noise to reveal shallow earth geologic structures in detail. The system promises to allow users to study and monitor the sustainability of an area's subsurface as well as pinpoint potential hazards hidden inside geological structures. Song and his collaborators, Yao Xie of the Georgia Institute of Technology and Fan-Chi Lin of the University of Utah, are planning to test their system by imaging the subsurface of geysers in Yellowstone National Park.

The technology is designed to create images of underground structures in settings where there is no active source, such as earth tremors. "We're using background noise," Song says. At Yellowstone, for instance, vibrations created by visitors and their cars feed the system. "We essentially use that type of information to tap into a very weak signal to infer the image of [what is] underground," Song explains.

The real-time in situ seismic imaging (RISI) technology images and monitors subsurface geophysical structures and dynamics in real time. "An RISI system is a wireless seismic network that senses and processes seismic noises from natural earthquakes and oil/gas exploration and production activities," Song says. "Instead of collecting data to a central place for postprocessing, the distributed seismic data processing and tomographic computing are performed in the in situ network and the evolving 3-D image is computed and delivered in real-time for visualization" (Figure 3).

Each solar-powered RISI station is equipped with an Advanced RISC Machine-based microprocessor, a mesh network

UNIVERSITY OF DAYTON.

[FIG4] Image enhancement allows the extraction of additional details within the target area, enabling the automated surveillance technology to accurately detect and classify objects that pose a possible threat to a nearby pipeline.

transceiver, a global positioning system (GPS) receiver, and a geophone, a device that converts ground movement (displacement) into voltage. Each station is capable of acquiring and processing data, and communicating with other nearby RISI stations, via the mesh network.

Song says that RISI promises to be useful in a wide range of georelated applications, including oil and gas exploration and production, safety and environmental surveillance, and volcano, landslide, and earthquake monitoring. "It can provide real-time awareness—in less than a second—of oil/gas exploration and production processes and potential subsurface hazards so that certain mitigation or optimization action can be taken promptly," he notes. The technology could also be used to alert homeowners to a subsurface change below their residence, an indication that their house may be sinking.

According to Song, ambient seismic events generally have an extremely low signal-to-noise ratio (SNR), mandating the use of a noise suppression algorithm. Yet even after denoising, the SNR of specific events may remain too low for reliable interpretation. "In ambient noise seismic imaging, it is basically impossible to identify events from individual data streams," Song says. A cross-correlation among different seismic channels and data streams is often needed to identify the events window and find the relative time difference. "The cross-correlation based signal processing needs to be done in a distributed fashion and under severe communication constraints," Song notes.

Song is optimistic the RISI has significant scientific and commercial potential. "We are helping industry lower costs greatly, because previously they only knew what was going on in the subsurface only days or even months later and we can reduce this time to just seconds," he says.

"We are essentially making automatic and fast video cameras for 3D/4D subsurface images while existing industry techniques and practices are [comparable to] manual painting with paintbrushes," Song says.

IMAGE ANALYSIS OF PIPELINE THREATS

Millions of kilometers of pipelines lay buried underground worldwide, transporting water, fuel and other types of essential liquids and gasses. All of these pipelines are vulnerable to breach by various types of construction and drilling activities. Unintentional pipeline damage often has serious and immediate consequences, including explosions, environmental contamination and the interruption of vital services. To lower the risk of pipeline damage, pipeline property owners must often limit or outright prohibit other activities on their land.

Recent advances in sensor technologies have led to the use of aircraft-based video acquisition systems to monitor activities on vulnerable properties. Such systems, however, typically generate massive amounts of data that human analysts must pore over visually to identify pipeline right-of-way threats. The process is both expensive and time-consuming.

(continued on page 18)

reader's CHOICE

Top Downloads in IEEE *Xplore*

The "Reader's Choice" column in *IEEE Signal Processing Magazine* contains a list of articles published by the IEEE Signal Processing Society (SPS) that ranked among the top 100 most downloaded IEEE *Xplore* articles. This issue is based on download data through March 2015. The table below contains the citation information for each article and the rank obtained in IEEE *Xplore*. The highest rank obtained by an article in this time frame is indicated in bold. Your suggestions and comments are welcome and should be sent to Associate Editor Michael Gormish (gormish@ieee.org).

TITLE, AUTHOR, PUBLICATION YEAR IEEE SPS PUBLICATIONS	ABSTRACT	RANK IN IEEE TOP 100						*N* TIMES IN TOP 100 (SINCE JAN 2011)
		MAR 2015	FEB 2015	JAN 2015	DEC 2014	NOV 2014	OCT 2014	
AN OVERVIEW OF MASSIVE MIMO: BENEFITS AND CHALLENGES Lu, L.; Li, G.Y.; Swindlehurst, A.L.; Ashikhmin, A.; Zhang, R. *IEEE Journal on Selected Topics in Signal Processing* vol. 8, no. 5, 2014, pp. 742–758	Equipping cellular base stations with a very large number of antennas potentially allows for orders of magnitude improvement in spectral and energy efficiency. This paper presents an extensive overview and analysis of massive MIMO systems.	31	22	**20**	37			6
IMAGE QUALITY ASSESSMENT: FROM ERROR VISIBILITY TO STRUCTURAL SIMILARITY Wang, Z.; Bovik, A.C.; Sheikh, H.R.; Simoncelli, E.P. *IEEE Transactions on Image Processing* vol.13, no. 4, 2004, pp. 600–612	This paper introduces a framework for quality assessment based on the degradation of structural information. Within this framework a structure similarity index is developed and evaluated. MATLAB code available.	53	80	47	59	**29**	33	30
AN INTRODUCTION TO COMPRESSIVE SAMPLING Candes, E.J.; Wakin, M.B. *IEEE Signal Processing Magazine* vol. 25, no. 2, 2008, pp. 21–30	This article surveys the theory of compressive sampling, also known as compressed sensing or CS, a novel sensing/sampling paradigm that goes against the common wisdom in data acquisition.	64	68	**52**				42
IMAGE SUPER-RESOLUTION VIA SPARSE REPRESENTATION Yang, J.; Wright, J.; Huang, T.S.; Ma, Y. *IEEE Transactions on Image Processing* vol. 19, no. 11, 2010, pp. 2861–2873	This paper presents an approach to single-image super-resolution based upon sparse signal representation of low and high resolution patches.	78		**57**	91	83	89	18
IMAGE QUALITY ASSESSMENT FOR FAKE BIOMETRIC DETECTION: APPLICATION TO IRIS, FINGERPRINT, AND FACE RECOGNITION Galbally, J.; Marcel, S.; Fierrez, J. *IEEE Transactions on Image Processing* vol. 23, no. 2, 2014, pp. 710–724	This paper uses 25 general image quality features extracted from the authentication image to distinguish between legitimate and imposter samples for fingerprint, iris, and 2D face biometrics.	96	67	**56**	97		69	11
NEW CHALLENGES FOR IMAGE PROCESSING RESEARCH Pappas, T.N. *IEEE Transactions on Image Processing* vol. 20, no. 12, 2011, p. 3321	The editor-in-chief of *IEEE Transactions on Image Processing* addresses the direction of the journal and image processing.		**51**	**51**		87		7

Digital Object Identifier 10.1109/MSP.2015.2438991
Date of publication: 13 August 2015

TITLE, AUTHOR, PUBLICATION YEAR IEEE SPS PUBLICATIONS	ABSTRACT	RANK IN IEEE TOP 100 MAR 2015	FEB 2015	JAN 2015	DEC 2014	NOV 2014	OCT 2014	*N* TIMES IN TOP 100 (SINCE JAN 2011)
PRIVACY PRESERVING DATA SHARING WITH ANONYMOUS ID ASSIGNMENT Dunning, L.A.; Kresman, R. *IEEE Transactions on Information Forensics and Security* vol. 8, no. 2, 2013, pp. 402–413	This paper offers an algorithm and analyzes multiple algorithms to assign ID numbers ranging from 1 to N to N parties without using a trusted central authority and is still resistant to collusion among other members.		**53**	65			**59**	5
REFLECTIONS ON SIGNAL PROCESSING [FROM THE EDITOR] Zoubir, A. *IEEE Signal Processing Magazine* vol. 30, no. 6, 2013, p. 4	Abdelhak Zoubir, *IEEE Signal Processing Magazine*'s editor-in-chief, answers the questions: How do you explain signal processing, and where is it going?		**70**					1
MODELING AND OPTIMIZATION FOR BIG DATA ANALYTICS: (STATISTICAL) LEARNING TOOLS FOR OUR ERA OF DATA DELUGE Slavakis, K.; Giannakis, G.B.; Mateos, G. *IEEE Signal Processing Magazine* vol. 31, no. 5, 2014, pp. 18–31	This article contributes to the cross-disciplinary efforts in data science by putting forth models capturing a range of SP-relevant data analytic tasks, such as principal component analysis (PCA), dictionary learning (DL), compressive sampling (CS), and subspace clustering. It offers scalable architectures and optimization algorithms for decentralized and online learning problems.			**62**				1
SUPER-RESOLUTION IMAGE RECONSTRUCTION: A TECHNICAL OVERVIEW Park, S.C.; Park, M.K.; Kang, M.G. *IEEE Signal Processing Magazine* vol. 20, no. 3, 2003, pp. 21–36	This article introduces the concept of super-resolution (SR) algorithms and presents a technical review of various existing SR methodologies and models the low-resolution image acquisition process.			**89**			93	20
MODULATION FORMATS AND WAVEFORMS FOR 5G NETWORKS: WHO WILL BE THE HEIR OF OFDM? Banelli, P.; Buzzi, S.; Colavolpe, G.; Modenini, A.; Rusek, F.; Ugolini, A. *IEEE Signal Processing Magazine* vol. 31, no. 6, 2014, pp. 80–93	This article provides a review of some modulation formats suited for 5G enriched by a comparative analysis of their performance in a cellular environment, and by a discussion on their interactions with specific 5G ingredients.			93	95	**80**	91	4
PERMISSION USE ANALYSIS FOR VETTING UNDESIRABLE BEHAVIORS IN ANDROID APPS Zhang, Y.; Yang, M.; Yang, Z.; Gu, G.; Ning, P.; Zang, B. *IEEE Transactions on Information Forensics and Security* vol. 9, no. 11, 2014, pp. 1828–1842	This paper presents a dynamic analysis platform for analysis of Android apps and their use of permission to access sensitive system resources. Performance is assessed by applying the platform to 1,249 apps.			**97**				1

[SP]

Digital Object Identifier 10.1109/MSP.2015.2464392

sp HISTORY

Anthony C. Davies

Stereo Sound Recording and Reproduction—Remembering the History

Abbey Road Studios have been an iconic landmark in London ever since The Beatles decided to name their latest album according to the street where their recording studio was located—and illustrate its cover with its now-famous pedestrian crossing. However, on 1 April 2015, there was another kind of success celebrated there. A new IEEE Historic Milestone plaque has been unveiled to commemorate the numerous inventions of engineer Alan Dower Blumlein on stereo sound recording and reproduction. The unveiling took place in the presence of over 100 people including Blumlein's family descendants, recording engineers from Abbey Road Studios, and representatives of the IEEE (Figure 1)—including IEEE President Howard Michel, IEEE Past President Roberto De Marca, IEEE Region 8 Past Director Martin Bastiaans, and IEEE History Committee representative Antonio Perez-Yuste.

IEEE President Howard Michel and Abbey Road Studios Managing Director Isabel Garvey unveiled the bronze plaque in Studio Two at Abbey Road. The plaque citation celebrates a seminal patent filed by Blumlein on 14 December 1931 that discloses several inventions related to stereo sound reproduction [1]. "It included a 'shuffling' circuit to preserve directional sound, an orthogonal 'Blumlein pair' of velocity microphones, the recording of

Digital Object Identifier 10.1109/MSP.2015.2440191
Date of publication: 13 August 2015

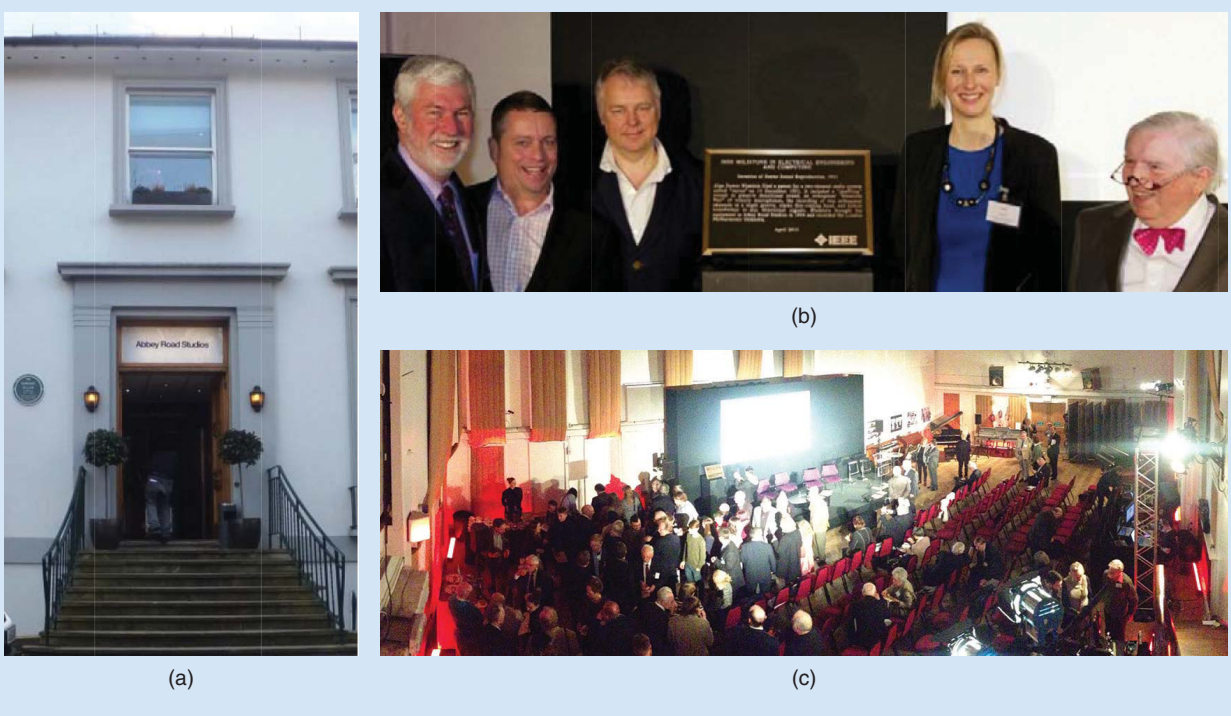

(a) (b) (c)

[FIG1] The IEEE Historic Milestone plaque unveiling ceremony at Abbey Road Studios. (a) The entrance to Abbey Road Studios. (b) From left: IEEE President Howard Michel, Alan Blumlein (grandson), Peter Cobbin (Abbey Road Studios director of engineering), Isabel Garvey (Abbey Road Studios managing director), and Simon Blumlein (son). (c) Studio Two during the ceremony.

two orthogonal channels in a single groove, stereo disc-cutting head, and hybrid transformer to mix directional signals." The citation also explains why Abbey Road Studios have been selected to host this commemorative plaque. This is the location where Blumlein brought his stereo equipment in 1934 to record the London Philharmonic Orchestra conducted by Sir Thomas Beecham. The plaque was moved outside to take its place at the entrance to the building, opposite a Westminster City Council plaque recognizing the opening of the Studios (then called EMI Studios) in 1931 by Sir Edward Elgar.

The ceremony included a display of historic items from the EMI Archive Trust and from Abbey Road Studios (Figure 2). A number of classic sound recordings and historic videos were played, including the well-known *Walking and Talking* video [2], in which tests of walking to and fro on a stage while speaking demonstrate how Blumlein's inventions allow the sound heard to track the speaker's movements. Lectures were given by Simon Blumlein about his father; by Robert Alexander, the author of a biography of Blumlein [3]; and by David Fisher, who explained in layman's terms the operation of some of Blumlein's inventions. The event concluded with a technical seminar involving several speakers and a panel discussion.

PIONEERING INNOVATIONS IN STEREO RECORDING AND REPRODUCTION

Blumlein's interest in improved sound reproduction originated from experience in local cinemas when watching the first movies with spoken dialogs also known as *talkies*. Early sound systems normally had a only a single set of speakers and could yield counterintuitive effects, e.g., with a character appearing on one side of the screen while his voice originated from the other side. For improved realism, Blumlein wanted a sound reproduction system where the sound would come from the position of the actor and would follow the actor as he moved around the screen.

Previous methods relied on using two widely separated microphones to mimic the two ears of the human head. The time difference of arrival of the sounds creates the directional sensation for the listener. While this solution achieves reasonable results when using headphones, it is not as effective as might be expected with loudspeakers because the listener receives both signals in each ear. In his 1931 patent, Blumlein proposed using two colocated microphones oriented at 90° to one another. Ideally, the microphones should occupy the same physical space. Since it cannot be achieved in practice, they are placed as close to each other as physically possible, for instance, with one centered directly on top of the other. The two microphone signals are then processed to convert phase differences into amplitude differences. Commonly, these microphones, referred to as a *Blumlein pair,* each have a figure-of-eight polar diagram. However, in his early experiments at EMI, Blumlein did not have access to such microphones and used two omni-directional microphones instead, separated by a baffle-board to simulate a human head [1]. The resulting apparatus delivers a high degree of stereo separation in the source signal as well as the room ambiance.

The name *stereo*, as universally used nowadays, is inaccurate since it conveys the concept of a three-dimensional representation, which is not the case. It is not an attempt to recreate the original sound field for the listener, and there is no front-to-back information and no up-and-down information. Nevertheless, the procedure creates an improvement in realism and listening experience, as well as providing the left-to-right directional information that was Blumlein's objective. Blumlein himself used the term *binaural sound* and his aim and achievement was to capture directional information as intensity differences only between the two channels.

Once the acquisition problem had been solved, the next challenge was about recording such stereo signals on some physical medium, namely a vinyl disk. Blumlein suggested cutting the sound track at 45° of the horizontal disc to record the two stereo channels in a single groove, one channel per "wall" of the groove [4]. Side-to-side movement of a conventional pick-up stylus produces (L+R), where L and R are the left- and right-hand channel signals, and up-and-down movement produces (L–R). The (L+R) signal is what would be produced from mono-equipment and records (discs), while (L–R) is the directional component. Typical pick-ups have significant crosstalk between their two channels, but

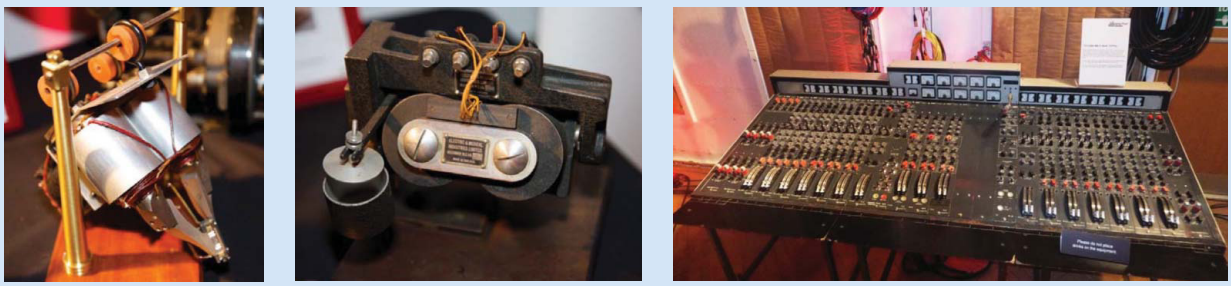

(a) (b) (c)

[FIG2] The EMI Archive Collection. (a) A Blumlein-pair ribbon microphone consisting of two colocated microphones oriented at 90° to one another. (b) A stereo cutting tool to record two audio channels in a single grove of a vinyl disk. (c) A mixing desk of the 1970s (TG 12345 Mk II) indicating the complexity that audio engineers faced to combine several sources and produce a professional-quality audio recording track.

the directional information is substantially preserved for the listener. Advantageously, this method was backward compatible with the mono equipment already in use. Mono recording indeed only uses the side-to-side movement of the stylus to reproduce the sound and up-and-down movements carry no useful information. Finally, this groove cutting technique also guaranteed that defects, such as low-frequency turntable rumble and high-frequency distortion, were equally distributed over both channels.

These basic ideas were almost forgotten until the introduction of FM broadcasting and the vinyl LP record in the 1950s led to some public interest in "hi-fi" music reproduction and affordable equipment started to became available for home use by enthusiasts. Even so, it was not until the mid to late 1960s that stereo transmissions and stereo recordings became routinely available and widespread, and Blumlein's inventions re-emerged as very important concepts. By then, costs of consumer electronic equipment had been substantially reduced and continued to fall with the improvements of solid-state electronics, integrated circuits, and low-cost manufacturing. Advances in the design of loudspeaker systems had enabled much smaller units to be made, compatible in cost and size with typical home environments. While other ways of creating more realistic sound reproductions have been devised (surround sound, ambisonics, quadraphonic, etc.), they have not yet found the widespread use and acceptance that stereo sound has. Interested readers should check [5, Sec. 4 and 6] for more details on these topics as well as a thorough historical survey of sound recording and reproduction.

A PROLIFIC INVENTOR

Blumlein was a prolific inventor of electronic circuit configurations and techniques, many of which are still familiar to today's designers (though the connection with Blumlein is usually forgotten). Examples include the cathode (and now emitter or source) follower, the "long-tailed pair," and many others. Using negative feedback to make circuits with predictable behavior despite wide component tolerances was one of his contributions, as was making use of the Miller effect, rather than just trying to minimize its effects. He worked for several companies before becoming involved in the sound-recording business, which began when he joined the Columbia Graphophone Company in 1929 (which merged with the Gramophone Company to form EMI a couple of years later). He also made many contributions to the EMI high-definition television system, which was the basis for all television systems for many years from the start of public TV broadcasting from Alexandra Palace in 1936. In preparation for and the early stages of World War II, he was actively involved in novel radar developments as well as navigational and direction-finding techniques. The successful "oboe" navigational aid invented by Reeves was developed with the aid of a team at EMI led by Blumlein. It was while testing his work on the very successful H2S blind-bombing aid that he died in a Halifax bomber aircraft crash in June 1942. The IEEE plaque is unrelated to all these other achievements but may help to increase awareness of them by electronic engineers and historians of technology, and to the 128 patents that this one talented person achieved [3], [6].

ACKNOWLEDGMENTS

The support and enthusiasm from Universal Music Ltd. and Abbey Road Studios, and their provision of Studio Two, was essential to the success of this event, as was the assistance from members of the Blumlein family and Joanna Hughes of the EMI Archives Trust. The friendly assistance of the employees at Abbey Road Studios and those who contributed to the demonstrations and presentations made this a memorable occasion.

AUTHOR

Anthony C. Davies (tonydavies@ieee.org) is an Emeritus Professor at King's College London, United Kingdom, a Life Fellow of the IEEE, and a former director of IEEE Region 8.

REFERENCES

[1] A. D. Blumlein, "Improvements in and relating to Sound-transmission, Sound-recording and Sound-reproducing Systems," British Patent 394325, June 14, 1933.

[2] A. D. Blumlein. Walking and talking. [Online]. Available: http://www.youtube.com/watch?v=rqaMiDqE6QQ.

[3] R. C. Alexander, *The Inventor of Stereo: The Life and Works of Alan Dower Blumlein*. Oxford, U.K.: Focal Press, 1999.

[4] EMI Archive Trust. (9 June 2015). Alan Blumlein's Stereo Model. [Online]. Available: https://vimeo.com/81729598

[5] P. Copeland. (2008). Manual of analogue sound restoration techniques. [Online]. The British Library. Available: http://www.bl.uk/reshelp/findhelprestype/sound/anaudio/analoguesoundrestoration.pdf

[6] R. W. Burns, *The Life and Times of A.D. Blumlein*. London: Inst. Eng. Technol., 2000.

[SP]

ERRATA

In the article "Challenges in Content-Based Image Indexing of Cultural Heritage Collections" by D. Picard, P. Gosselin, and M. Gaspard in the July 2015 issue of *IEEE Signal Processing Magazine* [1], the subtitle, "Support vector machine active learning with applications to text classification" was incorrect. The correct subtitle of [1] is "Automatic labeling and interactive search."

Reference

[1] D. Picard, P. Gosselin, and M. Gaspard, "Challenges in content-based image indexing of cultural heritage collections," *IEEE Signal Processing Mag.*, vol. 32, no. 4, pp. 95–102, July 2015.

Digital Object Identifier 10.1109/MSP.2015.2447371

from the GUEST EDITORS

Nicholas Evans, Sébastien Marcel,
Arun Ross, and Andrew Beng Jin Teoh

Biometrics Security and Privacy Protection

Biometrics is the science of recognizing individuals based on their behavioral and biological characteristics such as face, fingerprints, iris, voice, gait, and signature. A typical biometric system may be viewed as a pattern classification system that utilizes advanced signal processing schemes to compare and match biometric data.

The past decade has witnessed a rapid increase in biometrics research in addition to the deployment of large-scale biometrics solutions in both civilian and law enforcement applications. Example applications that incorporate biometric recognition include: logical and physical access systems; surveillance operations to fight against fraud and organized crime; immigration control and border security systems; national identity programs; identity management systems; and the determination of friend or foe in military installations.

Since an individual's biometric data is personal and sensitive, issues related to biometric security and privacy have been raised. These include: spoofing, where an adversary presents a falsified biometric trait to the system with the intention of masquerading as another person; evasion, where a person attempts to obfuscate or modify a biometric trait to avoid being detected by the system; database alteration, where the templates stored in a database are modified to undermine system integrity; and template compromise, where the stored biometric data is perused or stolen and exploited for illegitimate means.

The advent of cloud computing technology and personal mobile devices has broadened the application domain of biometrics; however, at the same time, it has brought to the forefront the need for dedicated security technologies to protect biometric data from being misappropriated and used for purposes beyond those intended. Similarly, the use of surveillance systems in public areas presents new challenges with respect to privacy.

The research community has responded to these concerns with new security and privacy enhancement and protection technologies. There are numerous indicators of the increasing interest, e.g., a number of special sessions in conferences, evaluation campaigns, tutorials, large-scale collaborative projects, and ongoing efforts toward standardization. A number of signal processing methods have been developed to analyze the vulnerability of biometric systems and design solutions to mitigate the impact of these vulnerabilities. At the same time, privacy-preserving constructs have been developed by signal processing researchers to ensure that stored and/or transmitted biometric data is adequately protected from misuse.

This special section was conceived to champion recent developments in the rapidly evolving field and also to encourage research in new signal processing solutions to security and privacy protection. After a rigorous preselection and peer-review process, eight articles were selected.

The first contribution from Hadid, Evans, Marcel, and Fierrez focuses on the security side of biometrics, providing a gentle introduction to spoofing and countermeasures and a methodology for their assessment. The article also provides a case study in face recognition.

The next contribution discusses how adversarial machine-learning techniques can be harnessed to protect biometric systems from sophisticated attacks. Biggio, Fumera, Russu, Didaci, and Roli argue that security is best delivered with adaptive, security-by-design solutions.

Itkis, Chandar, Fuller, Campbell, and Cunningham report the challenges in designing effective cryptosystems for iris-recognition systems. Their work also illustrates the shortcoming of the more traditional performance metrics used in biometrics and promotes the use of a new entropy metric.

The article by Patel, Ratha, and Chellappa reviews different approaches to cancelable biometric schemes for template protection. The aim of such techniques is to preserve privacy by preventing the theft of biometric templates through the application of noninvertible transforms.

Barni, Droandi, and Lazzeretti describe a different approach to template protection based on cryptographic technology. They illustrate how secure, two-party computation and signal processing in the encrypted domain can be combined to enhance security and protect privacy.

Still on the theme of template protection, Lim, Teoh, and Kim describe their work on biometric feature-type transformation. Such transformations are typically used as a precursor to many forms of biometric cryptosystems that demand specific input formats such as point-set or binary features.

The final article on template protection discusses the practical implications of biometric security and offers a fresh perspective to the problem. Nandakumar and Jain argue that improvements to security and privacy seldom come without degradations to recognition performance and that, consequently, there remains a significant gap between theory and practice.

The special section rounds out with an article by Bustard on the privacy and legal concerns surrounding the collection, storage, and use of personal biometric data. In particular, the article discusses recent European legislation on this issue and its potential impact on the adoption of biometrics technology.

Digital Object Identifier 10.1109/MSP.2015.2443271
Date of publication: 13 August 2015

ACKNOWLEDGMENTS

We wish to express our gratitude to Prof. Fulvio Gini, *IEEE Signal Processing Magazine's* then-area editor of special issues, for the support and advice provided throughout the preparation of this special section. We also thank Rebecca Wollman, IEEE Signal Processing Society publications administrator, for her valuable assistance. Finally, we acknowledge the many anonymous reviewers whose outstanding contributions have ensured and helped to enhance the quality of the articles that follow.

ABOUT THE GUEST EDITORS

Nicholas Evans (evans@eurecom.fr) is with the Department of Multimedia Communications, EURECOM, France.

Sébastien Marcel (marcel@idiap.ch) is with Idiap Research University, Switzerland.

Arun Ross (rossarun@cse.msu.edu) is with Michigan State University, United States.

Andrew Beng Jin Teoh (bjteoh@yonsei.ac.kr) is with Yonsei University, Seoul, South Korea.

[SP]

Using multiple object detection and recognition algorithms, a research team led by K. Vijayan Asari, an electrical and computer engineering professor at the University of Dayton in Ohio, has developed an automated surveillance technique that can be used to protect underground pipeline infrastructures (Figure 4).

The framework consists of three parts. The first part removes imagery that are not considered to be a threat to the pipeline. The method extracts a set of features that precisely represent the shape, structure and texture of various backgrounds, such as trees, buildings, roads and farmland, using a cascade of classifiers to eliminate the insignificant regions. The second part of the framework is a part-based object detection model for searching specific targets that are considered to be threat objects. The third part of the framework assesses the severity of pipeline threats by calculating the location and the temperature information of threat objects, such as construction equipment or drilling gear. "With our approach we can take into account the constraints associated with aerial imagery, such as low resolution, lower frame rate, large variations in illumination, and motion blurs," says Asari, who is also the director of the University of Dayton Vision Lab.

A major challenge to accurate threat detection are objects of interest that are partially occluded by shrubs, trees, buildings and other terrestrial elements. In the part-based model, an object is partitioned into a specific number of parts; the size of each part depends on the size of the object. "We then use local phase information to extract informative attributes for describing the individual parts," Asari says. "The next step is to group the object parts into several clusters. "In this process, we group similar parts into the same cluster and a histogram of oriented phase is used to describe the specific pattern of the parts," Asari continues. "This is to group similar parts of different vehicles, or similar parts in different images of the same vehicle, into the same cluster or category to find the presence of such categories in an occluded image to detect it as a threat object."

A MAJOR CHALLENGE TO ACCURATE THREAT DETECTION ARE OBJECTS OF INTEREST THAT ARE PARTIALLY OCCLUDED BY SHRUBS, TREES, BUILDINGS AND OTHER TERRESTRIAL ELEMENTS.

The output of the part-based object detection technique is the pixel location of the threat object in the input image. In real-world applications, however, a system user must also know the exact geographic location of a potentially threatening object. A registration process that links the acquired images to a geographical map provides this capability. Additionally, some detected threats may be far away from a pipeline, or have some other type of low threat probability. "Considering these issues, we have designed an additional framework that can automatically analyze the geolocation and temperature information of a detected object, and can assign a risk level to any given threat—high, medium, or low," Asari says. "A high temperature indicates that the vehicle is active and it may be moving to the pipeline right-of-way; a low temperature indicates that the vehicle is stationary and is of low risk."

"We have reached over 85% accuracy for machinery threat detection in tests," Asari says. "We are confident that our method can be used as a practical approach for wide-area surveillance and to protect pipeline infrastructures."

Asari says that he and his team are currently focusing on hardware-software integration and performance acceleration aspects. "We are looking to enable real-time processing in an onboard flight environment," he remarks.

AUTHOR

John Edwards (jedwards@johnedwardsmedia.com) is a technology writer based in the Phoenix, Arizona, area.

[SP]

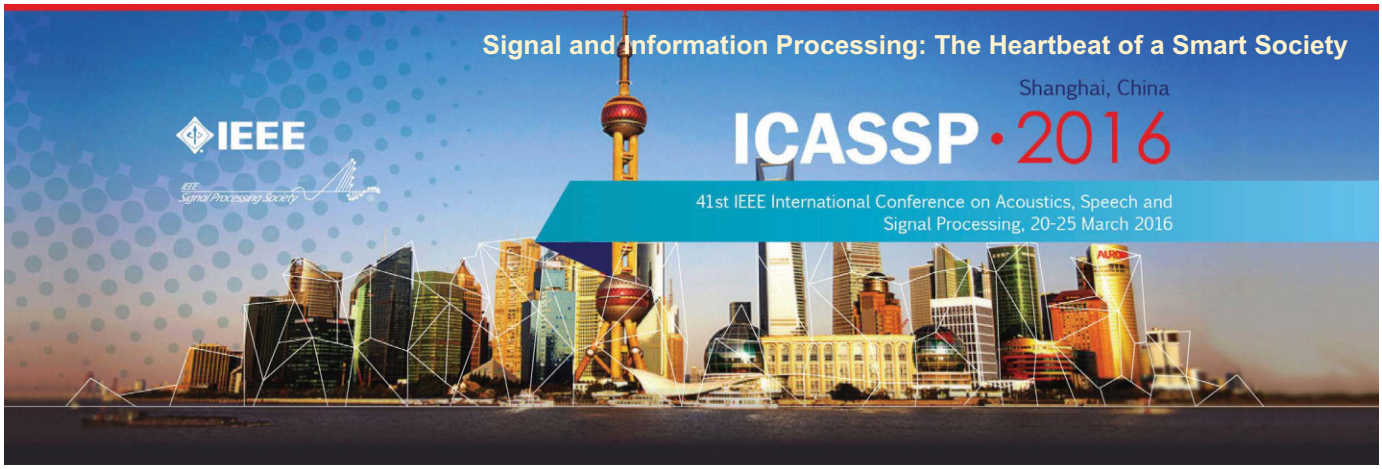

General Chairs
Zhi Ding, Univ. of California, Davis, USA
Zhi-Quan Luo, Univ. of Minnesota, USA
Wenjun Zhang, Shanghai Jiao Tong Univ., China

Technical Program Chairs
P. C. Ching, Chinese Univ. of Hong Kong, HK
Dominic K.C. Ho, Univ. of Missouri, USA

Finance Chairs
Shuguang Cui, Texas A&M Univ., USA
Rong Xie, Shanghai Jiao Tong Univ., China

Plenaries Chairs
Zhi-Pei Liang, UIUC, USA
Björn Ottersten, Univ. of Luxembourg, Luxembourg

Special Sessions Chairs
Tim Davidson, McMaster Univ., Canada
Jianguo Huang, Northwestern Polytech. Univ., China

Tutorials Chairs
Jian Li, Univ. of Florida, USA
Jose Principe, Univ. of Florida, USA

Student Session Chair
Wei Zhang, Univ. of New South Wales, AU

Registration Chairs
Tongtong Li, Michigan State Univ., USA
Xiaojun Yuan, ShanghaiTech Univ., China

Publicity Chairs
Xiaokang Yang, Shanghai Jiao Tong Univ., China
Mounir Ghogho, Leeds Univ., UK
Ignacio Santamaria, Univ. of Cantabria, Spain

Publication Chairs
Min Dong, Univ. of Ontario Inst. of Tech., Canada
Thomas Fang Zheng, Tsinghua Univ., China

Industrial & Exhibit Chairs
Li Deng, Microsoft, USA
Jinyu Li, Microsoft, USA
Cathy Wicks, Texas Instruments, USA

Local Arrangement Chairs
Ning Liu, Shanghai Jiao Tong Univ., China
Meixia Tao, Shanghai Jiao Tong Univ., China

Webmaster
Yi Xu, Shanghai Jiao Tong Univ., China

Workshop Chairs
Jianguo Huang, Northwestern Polytech. Univ., China
Jiwu Huang, Sun Yat-sen Univ., China

ICASSP2016: Signal and information processing is the driving heartbeat in the development of technologies that enrich our lives and advance our society. The 41st International Conference on Acoustics, Speech, and Signal Processing (ICASSP) will be held in the Shanghai International Convention Center, Shanghai, China between March 20 and 25, 2016. The conference provides, both for researchers and developers, an engaging forum to exchange ideas and propel new developments in this field. The 2016 conference will showcase world-class presentations by internationally renowned speakers and will facilitate a fantastic opportunity to network with like-minded professionals from around the world. Topics include but are not limited to:

- Audio and acoustic signal processing
- Bio-imaging and biomedical signal processing
- Signal processing education
- Speech processing
- Industry technology tracks
- Information forensics and security
- Machine learning for signal processing
- Signal processing for Big Data
- Multimedia signal processing
- Sensor array & multichannel signal processing
- Design & implementation of signal processing systems
- Signal processing for communications & networking
- Image, video & multidimensional signal processing
- Signal processing theory & methods
- Spoken language processing
- Signal processing for the Internet of Things

Shanghai: Shanghai is the most populous city in China and one of the most populous cities in the world. A global city, Shanghai exerts influence over global commerce, finance, culture, art, fashion, research and entertainment. The city is located in the middle portion of the Chinese coast, and sits at the mouth of the Yangtze River. The city is a tourist destination renowned for its historical landmarks, such as the Bund and City God Temple, and its modern and ever-expanding Pudong skyline including the Oriental Pearl Tower. Today, Shanghai is the largest center of commerce and finance in mainland China, and has been described as the "showpiece" of the world's fastest-growing major economy.

Submission of Papers: Prospective authors are invited to submit full-length papers, with up to four pages for technical content including figures and possible references, and with one additional optional 5th page containing only references. A selection of best student papers will be made by the ICASSP 2016 committee upon recommendations from the Technical Committees.

Tutorial and Special Session Proposals: Tutorials will be held on March 20 and 21, 2016. Tutorial proposals must include title, outline, contact information, biography and selected publications for the presenter(s), and a description of the tutorial and the material to be distributed to participants. Special session proposals must include a topical title, rationale, session outline, contact information, and a list of invited speakers. Additional information can be found at the ICASSP 2016 website.

Signal Processing Letters: Authors of IEEE Signal Processing Letters (SPL) papers will be given the opportunity to present their work at ICASSP 2016, subject to space availability and approval by the ICASSP Technical Program Chairs. SPL papers published between January 1, 2015 and December 31, 2015 are eligible for presentation at ICASSP 2016.

Show and Tell: S&T offers a perfect stage to showcase innovative ideas in all technical areas of interest at ICASSP. S&T sessions contain demos that are highly interactive and visible. Please refer to the ICASSP 2016 website for additional information regarding demo submission.

Important Deadlines:

Special session & tutorial proposals	August 3, 2015
Notification of special session & tutorial acceptance	September 11, 2015
Submission of regular papers	September 25, 2015
Signal processing letters	December 16, 2015
Notification of paper acceptance	December 21, 2015
Revised paper upload	January 22, 2016
Author registration	January 22, 2016

Digital Object Identifier 10.1109/MSP.2015.2411565

Abdenour Hadid, Nicholas Evans, Sébastien Marcel, and Julian Fierrez

Biometrics Systems Under Spoofing Attack

An evaluation methodology and lessons learned

©ISTOCKPHOTO.COM/GREYFEBRUARY

Biometrics Security and Privacy Protection

Biometrics already form a significant component of current and emerging identification technologies. Biometrics systems aim to determine or verify the identity of an individual from their behavioral and/or biological characteristics. Despite significant progress, some biometric systems fail to meet the multitude of stringent security and robustness requirements to support their deployment in some practical scenarios. Among current concerns are vulnerabilities to spoofing—persons who masquerade as others to gain illegitimate accesses to protected data, services, or facilities. While the study of spoofing, or rather antispoofing, has attracted growing interest in recent years, the problem is far from being solved and will require far greater attention in the coming years. This tutorial article presents an introduction to spoofing and antispoofing research. It describes the vulnerabilities, presents an evaluation methodology for the assessment of spoofing and countermeasures, and outlines research priorities for the future.

Digital Object Identifier 10.1109/MSP.2015.2437652
Date of publication: 13 August 2015

INTRODUCTION

The provision of security can entail the protection of sensitive data, services, or facilities by ensuring that only authorized

1053-5888/15©2015IEEE

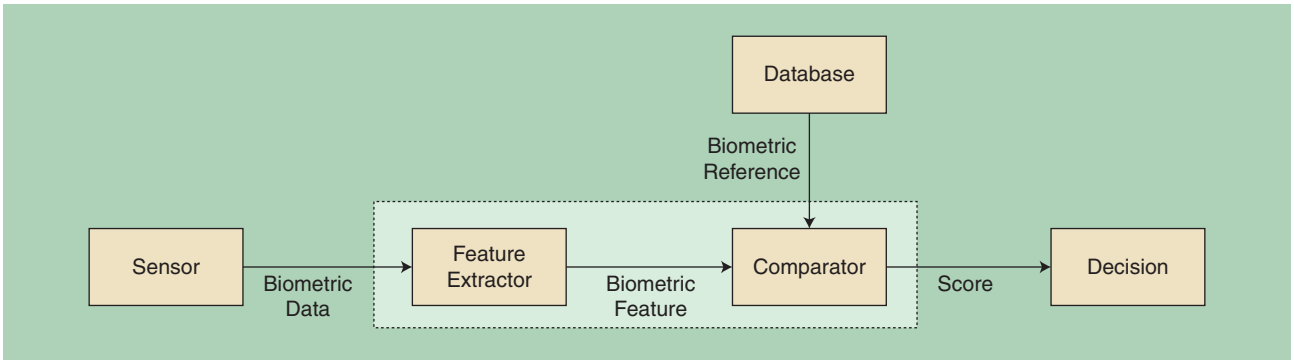

[FIG1] **A generic biometric system.**

persons have access. Though passwords provide some protection against illegitimate access, they are often so simple that they can be guessed or easily cracked. While offering improved security, complex passwords can be difficult to remember and consequently often "stored" via less secure means. Furthermore, the same password is often used across multiple applications or platforms meaning a cracked password can enable a fraudster to access multiple resources.

An attractive alternative to passwords involves biometric recognition. Biometrics refer to a person's behavioral and biological characteristics such as their face, fingerprint, iris, voice, hand geometry, and gait. Biometric traits can be highly discriminative yet less easily lost or stolen [1]. Despite their appeal, however, biometric systems are vulnerable to malicious attacks [2]. Among them are spoofing attacks, also called *presentation attacks*, which refer to persons masquerading as others to gain illegitimate access to sensitive or protected resources. As an example, a fraudster could fool or spoof a face-recognition system using a photograph, a video, or a three-dimensional (3-D) mask bearing resemblance to a legitimate individual.

Even though the threat of spoofing is now well recognized, the problem is far from being solved, thus antispoofing research warrants far greater attention in the future. This tutorial article introduces the problem of spoofing and related research to develop antispoofing solutions. The focus is an evaluation methodology for assessing both the effect of spoofing and the performance of spoofing countermeasures. A case study in face recognition is included to illustrate the application of the evaluation methodology in practice. Finally, the article also includes a summary of the lessons learned through our own research and outlines a number of research priorities for the future. Most of the material is based upon antispoofing research performed in the scope of the European TABULA RASA research project (http://www.tabularasa-euproject.org), which was identified as a success story by the European Commission (http://europa.eu/rapid/press-release_MEMO-13-924_en.htm).

The presentation is self-contained and aimed at both the generally knowledgeable and nonspecialist. The article aims to provide an overview of the research problem, not a comprehensive survey of the plethora of antispoofing techniques in the literature; such surveys can be found elsewhere, e.g., for fingerprint recognition [3], face recognition [4], and speaker recognition [5]. The intention is also to stimulate further work, particularly the development of standard metrics, protocols, and data sets for evaluating progress.

BIOMETRICS

The term *biometrics* is derived from the Greek words *bio* (life) and *metric* (to measure). The goal of a biometric recognition system is to determine or verify the identity of an individual from his/her behavioral and/or biological characteristics. Applications include criminal identification, airport checking, computer or mobile device log-in, building and critical infrastructure access control, digital multimedia rights control, transaction authentication, voice mail, and secure teleworking. Various biometrics have been investigated, from the most conventional including fingerprint, iris, face, and voice, to more emerging modalities such as gait, hand-grip, ear, and electroencephalograms. Each modality has its own strengths and weaknesses [1]. For example, face recognition is among the most socially accepted biometric; face recognition is a natural method of identification used everyday by humans. In contrast, while fingerprint and iris recognition may be more reliable, they are also more intrusive. In practice, the choice of biometric modality depends on the application.

Biometric systems typically function in one of two distinct modes: 1) verification (or authentication) and 2) identification. An authentication system aims to confirm or deny a claimed identity (one-to-one matching), whereas an identification system aims to identify a specific individual (one-to-many matching). Although there are some differences between the two modes, their most basic operation, namely that of feature-to-reference comparison, is identical and consists of the following steps illustrated in Figure 1.

First, a biometric sample (e.g., a face image) is acquired from a sensor (e.g., a digital camera). Biometric features (e.g., facial intensity, color, or texture) are then extracted from the sample. These can be a set of parameters (or coefficients) that provide a compact representation of the biometric sample, which is more discriminative and amenable to pattern recognition. Biometric features should minimize variations due to acquisition or environmental factors (e.g., facial expression, pose, and illumination)

while discriminating between the biometrics collected from different individuals.

To determine or verify the identity corresponding to a given biometric sample, the features are compared to a single (verification) or set of (identification) biometric references acquired previously during an enrollment phase. These comparisons are made by a comparator that produces a score reflecting the similarity between features and references. The decision is an acceptance or rejection in the case of verification, or the identity of the closest match in the case of identification.

(a)

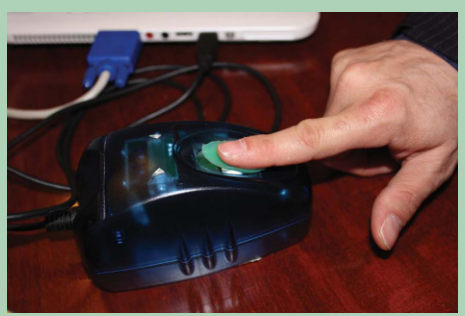

(b)

[FIG2] Example spoofing attacks: (a) face-recognition spoofing using a 3-D face mask; (b) fingerprint-recognition spoofing using a fake fingerprint.

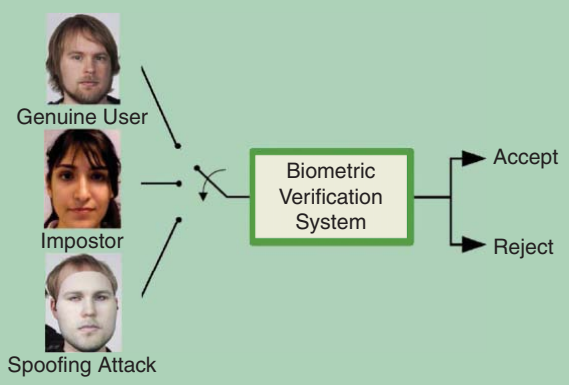

[FIG3] A biometric verification system should accept genuine users but reject both zero-effort impostors and concerted-effort spoofing attacks.

SPOOFING ATTACKS

It is now widely acknowledged that biometric systems are vulnerable to manipulation. There are two broad forms of attack: direct and indirect. Direct attacks, also referred to as *spoofing* or *presentation attacks*, are performed at the sensor level, outside the digital limits of the biometric system. Indirect attacks, however, are performed within the digital limits by intruders such as cybercriminal hackers. These attacks may attempt to bypass the feature extractor or the comparator, to manipulate biometric references, or to exploit vulnerabilities in the communications channels.

Whereas traditional digital protection mechanisms such as encryption can be deployed to prevent indirect attacks, they cannot prevent direct attacks. Furthermore, indirect attacks require specialist expertise or equipment, whereas direct attacks such as those illustrated in Figure 2 can be implemented by the layman. Direct attacks are therefore of significant concern.

Vulnerabilities to spoofing are a barrier to the successful exploitation of biometrics technology. Unfortunately, some vulnerabilities have been exposed and well publicized in the international media. One of the earliest examples was demonstrated at Black Hat 2009, the world's premier technical security conference. Researchers from the Security and Vulnerability Research Team at the University of Hanoi (Vietnam) showed how the face-recognition user authentication systems introduced by three different laptop computer manufacturers could be spoofed or bypassed using photographs of legitimate users. This vulnerability is now listed in the National Vulnerability Database maintained by the National Institute of Standards and Technology (NIST) in the United States. More recently, the Chaos Computer Club, a German hacking collective, showed how an artificial finger could be used to spoof a fingerprint-user authentication system developed by one of the world's most popular smartphone manufacturers.

The typical countermeasure deployed to detect spoofing attacks involves liveness detection: systems that aim to detect signs of life. Since it is inherently more difficult to spoof multiple modalities and systems simultaneously [6], multimodal biometric systems have also been investigated as a solution to spoofing. Even so, spoofing remains a serious cause for concern, with many biometric systems remaining vulnerable even to the simplest forms of spoofing attack. Unfortunately, research to develop spoofing countermeasures is still in its infancy and warrants considerably greater attention in the future.

EVALUATION METHODOLOGY: VULNERABILITIES TO SPOOFING

This section describes the first component of the proposed methodology for the evaluation of biometric systems and its vulnerabilities to spoofing. For simplicity, the following considers only biometric verification (one-to-one matching). Traditionally, biometric systems are evaluated using large data sets of representative biometric samples with which performance is assessed according to a standard protocol of genuine and impostor trials. An evaluation essentially measures the proportion of these trials

that the system misclassifies, i.e., genuine trials classified as impostor trials and vice versa. This approach, involving only casual impostors, equates to performance in the face of zero-effort spoofing attacks, i.e., where impostors make no effort to replicate the biometric traits of another, legitimate individual.

Spoofing is accomplished using fake biometric samples expressly synthesized or manipulated to provoke artificially high comparator scores. Biometric systems should be robust to both zero-effort impostor trials and concerted-effort spoofing attacks. This ternary scenario is illustrated in Figure 3. The biometric system should produce high scores in the case of genuine trials, but low scores in the case of impostor and spoofed access attempts.

Assessment requires biometric data of three categories comprising genuine, zero-effort impostor, and spoofed trials. Biometric data are typically further divided into three nonoverlapping subsets: a training set, a development set, and a test set, the purpose of which is as follows.

- The training set is used to train the biometric system (e.g., the learning of background models) or to develop spoofing countermeasures.
- The development set is used for decision threshold optimization and the a posteriori performance at a given operating point. For each identity, the development set is partitioned into two subsets:
  - an enrollment subset used to create biometric references or models for each identity
  - a probe subset used for biometric comparisons (genuine, zero-effort impostor and spoofing attack trials).
- The test set is used to compute the a priori performance given the threshold determined from the development set. The test set is similarly partitioned into enrollment and probe subsets.

Two assessment scenarios can then be defined:

- a licit scenario for assessing baseline performance using genuine and zero-effort impostor trials
- a spoofing scenario for assessing vulnerabilities to spoofing.

Figure 4 illustrates example score distributions for a licit scenario (i.e., genuine and impostor trials with no spoofing attacks). With only a little overlap between the two distributions, there is high potential to distinguish between genuine and impostor trials. Nonetheless, the system will still make mistakes, leading to two different error rates:

- the false rejection rate (FRR), which reflects the proportion of genuine trials misclassified as zero-effort impostor trials
- the false acceptance rate (FAR), which reflects the proportion of zero-effort impostors trials misclassified as genuine trials.

Both the FAR and FRR are dependent upon a decision threshold τ, an example of which is illustrated by the vertical, dashed line in Figure 4. Impostor trails corresponding to a score higher than this threshold will be misclassified as genuine trials, whereas genuine trials with a score lower than this threshold will be misclassified as impostor trials.

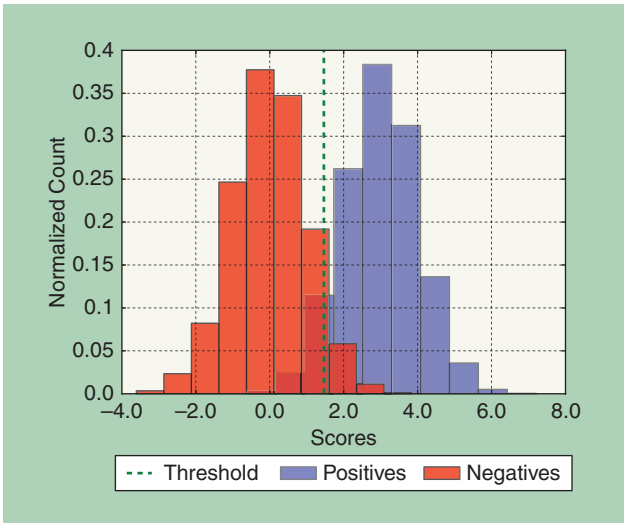

[FIG4] Score distribution of genuine users (positives) and zero-effort impostors (negatives) under a licit scenario. A decision threshold τ is illustrated by the vertical, dashed line.

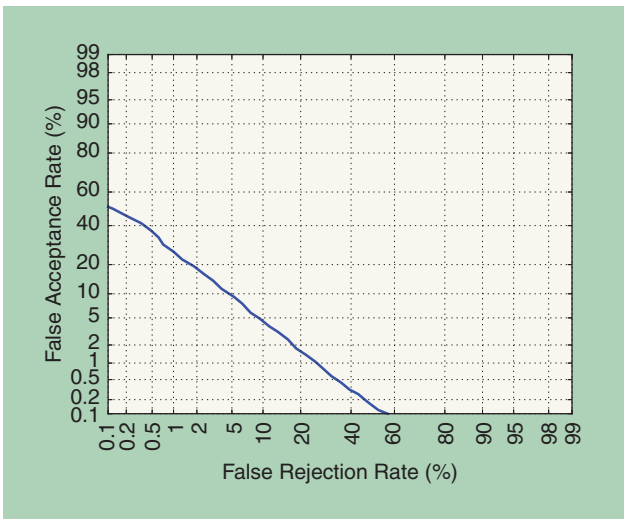

[FIG5] An example of a detection error tradeoff (DET) plot showing an EER in the order of 6%.

Since the two errors are inversely related, it is often desirable to illustrate performance as a function of the threshold τ. For a specific data set $\mathcal{D}$, one such measure is the half total error rate (HTER):

$$\mathrm{HTER}(\tau, \mathcal{D}) = \frac{\mathrm{FAR}(\tau, \mathcal{D}) + \mathrm{FRR}(\tau, \mathcal{D})}{2}. \quad (1)$$

Performance can also be illustrated with detection-error tradeoff (DET) profiles, an example of which is illustrated in Figure 5. DET profiles illustrate the behavior of a biometric system for varying decision thresholds, τ, and show the tradeoff between the FAR and the FRR.

In practice, the decision threshold is chosen to minimize an a posteriori performance criteria, which is normally application

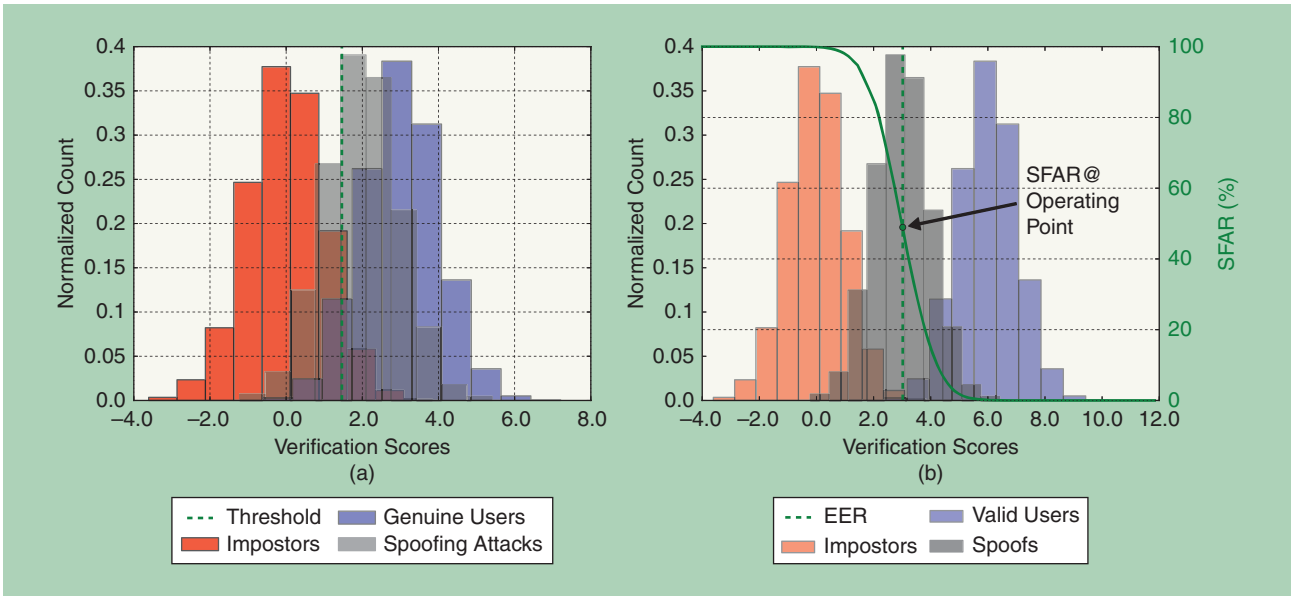

[FIG6] Example score distributions: (a) genuine users, zero-effort impostors, and (b) spoofing attacks with the SFAR as a function of the decision threshold τ.

dependent. Without focusing on a specific application, the equal error rate (EER) is also a common reference metric. For a development data set $\mathcal{D}_{\text{dev}}$, the EER is given by

$$\tau^*_{\text{EER}} = \arg\min_{\tau} |\,\text{FAR}(\tau, \mathcal{D}_{\text{dev}}) - \text{FRR}(\tau, \mathcal{D}_{\text{dev}})\,|, \qquad (2)$$

where the threshold τ^*_{EER} is set to equalize the FAR and FRR. The most reliable estimate of a priori performance is then determined using a test set $\mathcal{D}_{\text{test}}$ using the predetermined threshold. For example, the HTER is determined using the decision threshold τ^*_{EER} according to

$$\text{HTER}(\tau^*_{\text{EER}}, \mathcal{D}_{\text{test}}) = \frac{\text{FAR}(\tau^*_{\text{EER}}, \mathcal{D}_{\text{test}}) + \text{FRR}(\tau^*_{\text{EER}}, \mathcal{D}_{\text{test}})}{2}. \qquad (3)$$

Performance in the face of spoofing is assessed with the spoofing scenario, i.e., by replacing the subset of zero-effort impostor trials with spoofed trials. Figure 6(a) illustrates example score distributions for a population of genuine, zero-effort impostor and spoofed trials. The shift between the impostor and spoofed trial distributions illustrates the likely effect of spoofing on biometric recognition performance; the overlap between the score distributions for genuine and spoofed trials is significantly greater than that between genuine and impostor distributions.

A quantitative measure of system vulnerability can be expressed in terms of the spoof FAR (SFAR) which reflects the percentage of spoofed trials misclassified as genuine trials given a decision threshold τ. An example SFAR profile is illustrated as a function of the threshold τ in Figure 6(b).

The vulnerability to spoofing is thus reflected in the difference between the SFAR (spoof scenario) and the FAR (licit scenario), again as a function of τ. These differences are illustrated in the example DET plots for both licit (FAR versus FRR) and spoof (SFAR versus FRR) scenarios in Figure 7. Finally, it is often of interest to express system vulnerability for a specific operating point defined by a given FRR (e.g., the EER). System vulnerability is then given by the difference between the FAR and the SFAR for the same FRR (vertical, dashed line in Figure 7). In this illustrative example for an FRR in the order of 6%, an FAR of under 10% increases to an SFAR of over 60%. This system would misclassify approximately two in three spoofed trials as genuine accesses.

EVALUATION METHODOLOGY: SPOOFING COUNTERMEASURES

This section describes the extension of the evaluation methodology to assess spoofing countermeasures. Figure 8(a) illustrates the deployment of spoofing countermeasures as independent subsystems. Just like the biometric system, the countermeasure is a two-class classifier, its role being to distinguish between genuine and spoofed trials. Again, just like the biometric system, it will also make mistakes, leading to two new error rates referred to as the false living rate (FLR) and the false fake rate (FFR). The FLR reflects the percentage of spoofed trials misclassified as genuine trials, whereas the FFR reflects the percentage of genuine trials misclassified as spoofed trials.

Even as independent subsystems, spoofing countermeasures impact directly on the performance of the biometric system as a whole; while aiming to reduce vulnerabilities to spoofing, they also have potential to reject genuine trials [7]. Thus, while countermeasure subsystems can be assessed independently, it is often of greater interest to assess the performance of the system as a whole, for example, using a score fusion strategy illustrated in Figure 8(b).

Countermeasure subsystems may alternatively be integrated, not in parallel but in series, with the biometric system as illustrated in Figure 8(c). Assessment involves four different system configurations, for each of which a different DET profile is produced. Examples are illustrated in Figure 9. The first

configuration (blue profile) illustrates the performance of the baseline biometric system (no spoofing, no countermeasures). The second configuration (black profile) illustrates the performance of the same system when subjected to spoofing attacks. The third configuration (green profile) illustrates the improvement in performance with active countermeasures. While seemingly complete, a fourth configuration (red profile) is still needed to determine the performance of the integrated biometric and countermeasure systems in the absence of spoofing. This final configuration is essential to gauge the effect of the countermeasure on genuine trials that may be misclassified as spoofing attacks. To reiterate, all four configurations are needed to properly observe the complex impact of spoofing and countermeasures on integrated systems.

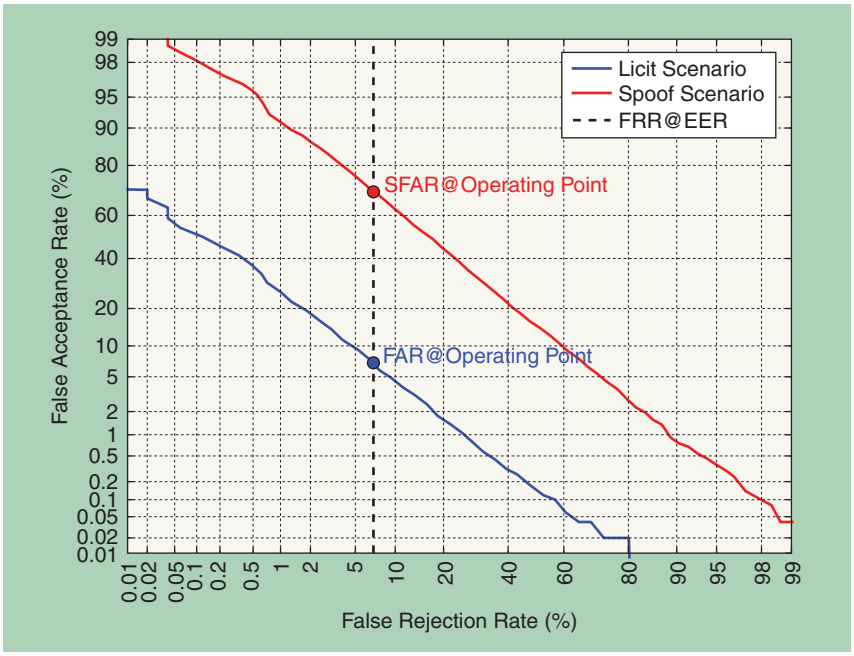

[FIG7] A DET profile illustrating the SFAR. FRR@EER refers to the FRR on the test set for a decision threshold τ for which the FAR and FRR are equal (determined from the development set).

Even if it is the effect on overall recognition performance that is of greatest interest, there will always be an interest to first assess countermeasure performance independently. This configuration might be limited to system development. The second, fused approach is simple and straightforward. However, unless separate decisions are applied to the countermeasure and recognition subsystem scores, it does not support the explicit detection of spoofed trials. The third approach is therefore the most appealing in practice, allowing for explicit spoof detection and an evaluation of countermeasure impacts on overall system performance.

Finally, whatever the approach to integration, performance is dependent on a separate countermeasure decision threshold τ^{CM}. Since it impacts upon overall performance, the FAR/FRR/SFAR can be determined for a range of different operating points specified, for example, by an application-dependent FFR. So as to reflect the likely performance for a variety of different application scenarios, the TABULA RASA project considered values of FFR $= \{1, 5, 10\}\%$. Even so, countermeasure integration and assessment is not a solved problem and very much a topic of ongoing research. This is discussed further in the section "Lessons Learned."

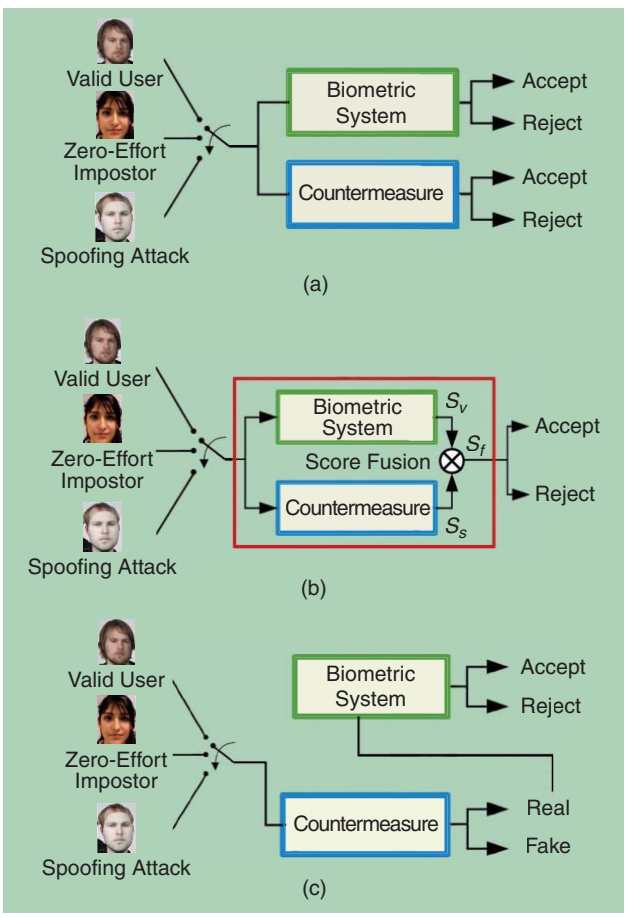

[FIG8] Integrated biometric systems and spoofing countermeasures. (a) Two independent components. (b) Fused subsystems. (c) A two-step process.

EVALUATION METHODOLOGY: A CASE STUDY IN 2-D FACE VERIFICATION

As an illustration of the evaluation methodology in practice, the following presents a case study in two-dimensional (2-D) face verification.

PREVIOUS WORK

The vast majority of past research in 2-D face recognition has focused on maximizing the discrimination between the faces of different persons [8]. Liveness detection has received considerably less attention, even though 2-D face-recognition systems have been known to be vulnerable to spoofing for some time.

Many face-recognition systems can be spoofed by presenting to the camera photographs, videos, or 3-D masks of enrolled persons [2]. While makeup or plastic surgery are also viable spoofing attacks, photographs are among the most easily implemented, most effective, and therefore the most likely in practice. Video attacks can be especially effective since they are dynamic signals that more closely resemble a genuine trial. Furthermore, it is natural to assume that systems that are vulnerable to photo attacks will also be vulnerable to video attacks.

SPOOFING IS ACCOMPLISHED USING FAKE BIOMETRIC SAMPLES EXPRESSLY SYNTHESIZED OR MANIPULATED TO PROVOKE ARTIFICIALLY HIGH COMPARATOR SCORES.

As is the case for the spectrum of other modalities, the typical countermeasure to prevent the spoofing of 2-D face-recognition systems involves liveness detection. Here, liveness indicators include eye blinking, changes in facial expression, mouth movements, estimates of skin texture, structure and motion analysis, and depth information, etc. Multispectral and reflectance analysis has also been used successfully to differentiate between living faces and lifeless fakes [2]. Alternatively, face recognition can be combined with other biometric modalities such voice or gait recognition. A survey of face-spoofing detection approaches can be found in [4].

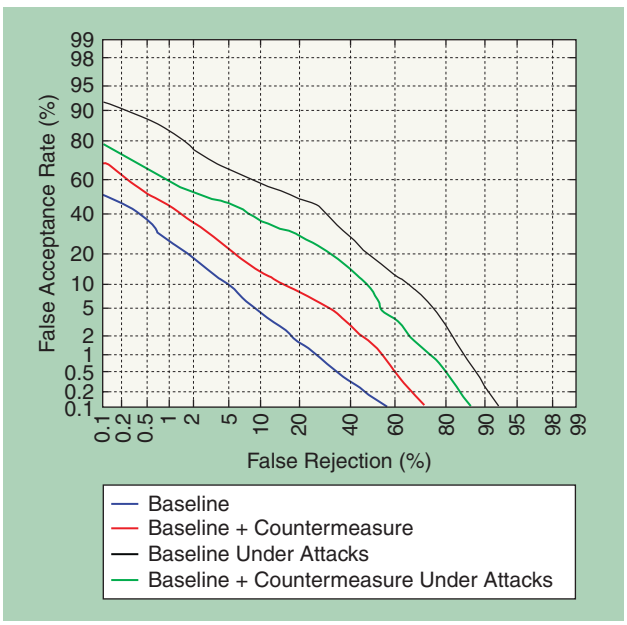

[FIG9] The FAR and FRR of a biometric system with and without spoofing attacks, and with and without countermeasures.

BASELINE FACE BIOMETRIC SYSTEM

The system combines a part-based face representation with Gaussian mixture models (GMMs) [9]. The system divides the face into blocks and treats each block as a separate observation of the same underlying signal (the face). A feature vector is thus obtained from each block by applying the discrete cosine transform (DCT). The distribution of the feature vectors is then modeled using GMMs.

For feature extraction, the face is normalized, registered, and cropped. The cropped and normalized face is divided into blocks (parts) from each of which a feature vector is extracted. Each feature vector is treated as a separate observation of the same underlying signal and the distribution of the feature vectors is modeled using GMMs. The feature vectors from each block are obtained by applying the DCT. Once the feature vectors are calculated, feature distribution modeling is achieved by performing background model adaptation of GMMs. Background model adaptation first involves the training of a world (background) model Ω_{world} from a large data set of faces. Client models Ω^i_{client} for client i are learned through the adaptation of the world model toward the observations of the client.

Verification is achieved by scoring an observation x against both the client (Ω^i_{client}) and world (Ω_{model}) models. They produce log-likelihood scores, which are combined to give the single log-likelihood ratio (LLR). The LLR is used to assign the observation

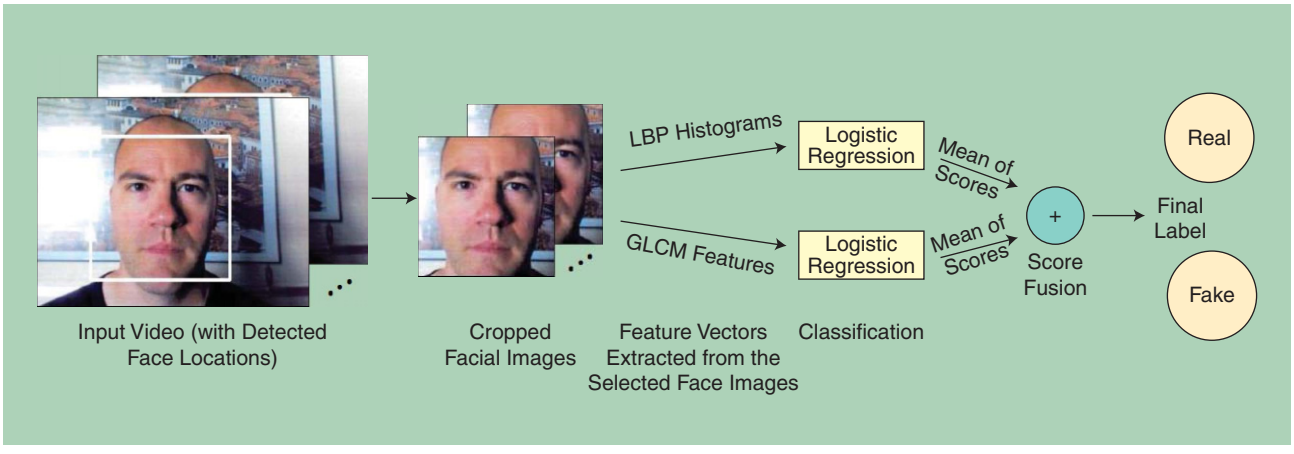

[FIG10] A block diagram of the texture-based countermeasure.

[FIG11] The setup and sample images from the REPLAY-ATTACK database. Spoofed data collection using (a) printed photographs and (b) examples of genuine accesses and spoofing attacks. In (b), the columns from left to right show examples of real accesses, printed photographs, mobile phone, and tablet attacks.

to the world class (not the client) or the client class based on a predefined threshold τ.

SPOOFING COUNTERMEASURE

Genuine and spoofed face images exhibit differences at the texture level. Spoofed face images are captured or acquired twice, first from the live client and second from the spoofing medium at recognition time. Accordingly, spoofed face images tend to exhibit degraded facial texture due to reproduction or printing on the spoofing medium (a photograph) and the two imaging systems [10]. As a result, estimates of facial texture quality, here obtained using logistic regression [11], serve to identify spoofing attacks. The method is computationally efficient and does not require any additional user cooperation; it is nonintrusive.

AS IS THE CASE FOR THE SPECTRUM OF OTHER MODALITIES, THE TYPICAL COUNTERMEASURE TO PREVENT THE SPOOFING OF 2-D FACE-RECOGNITION SYSTEMS INVOLVES LIVENESS DETECTION.

Figure 10 illustrates a block diagram of the texture-based spoofing countermeasure. The approach analyzes the texture of single facial images using two different texture descriptors: local binary patterns (LBPs), which encodes the microtexture patterns into a feature histogram and features computed from the gray-level co-occurrence matrices (GLCMs), which describe the distribution of different combinations of gray-level pairs in an image block. The combination of these two measurements provides an effective representation of the overall facial texture quality. The facial texture descriptions are fed into logistic regression classifiers. Score-level fusion of the individual classifier outputs determine whether the facial image corresponds to a living person or a spoofed reproduction. As the database consists of video sequences, several cropped face images from each video are used for feature extraction at intervals of 0.6 seconds. The final score for the two individual texture representations is determined by averaging the scores of each face image. The outputs of each classifier are then fused using the sum rule with min-max score normalization [6] to obtain the final label for each video sequence.

THE FACE SPOOFING ATTACK DATABASE

Experiments were performed with the REPLAY-ATTACK face spoofing database [12]. It consists of 1,300 samples comprising genuine trials, photo, and video spoofing attacks. The database contains samples collected from 50 persons under varying lighting conditions. The data is split into four subgroups comprising enrollment, training, development, and test data. While the enrollment set includes samples collected from all clients, there is no client overlap in the training, development, and test data sets.

All videos are generated by either having a (real) client trying to access a laptop through a built-in webcam or by displaying a photo or a video recording of the same client for at least nine seconds. In total, 20 attack videos were recorded for each client, whereas six videos were captured for real accesses. The enrollment set contains 100 videos (two per client) and is used exclusively to evaluate baseline performance. The training set contains 60 real accesses and 300 attacks. The development set contains 60 real accesses and 300 attacks, whereas the test set contains 80 real accesses and 400 attacks. Examples of real accesses and spoofing attacks from the REPLAY-ATTACK database are shown in Figure 11. A complete description of the database and associated protocol can be found in [12].

EXPERIMENTAL RESULTS

Vulnerabilities to spoofing and countermeasure impacts were assessed according to the evaluation methodology introduced in this article. Results in the form of DET profiles are illustrated in Figure 12: (a)–(c) represent overall system performance under a spoofing attack for three different operating points where countermeasures are tuned to produce FFRs of 1%, 5%, and 10%, respectively. For completeness, (d) is included to illustrate countermeasure performance in independence from biometric recognition. The EER of the countermeasures is in the order of 5%.

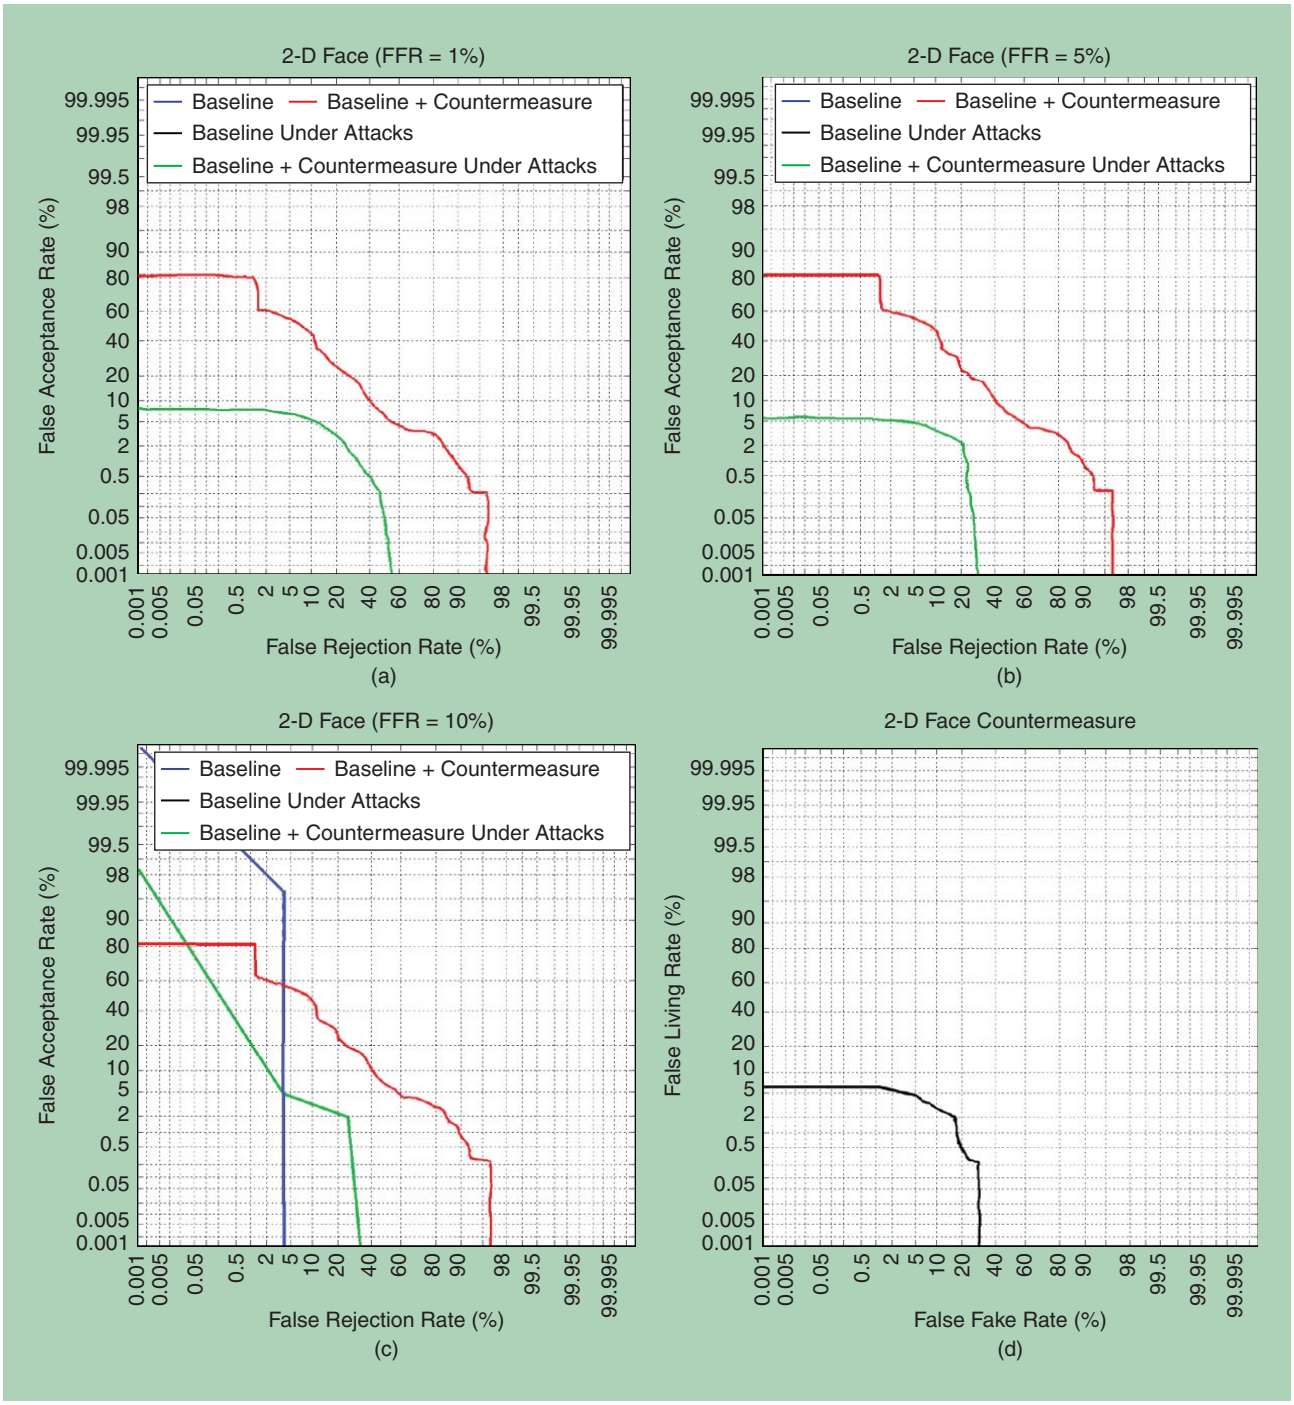

[FIG12] DET profiles for the 2-D face verification systems with/without spoofing and with/without countermeasures at different countermeasures operating points [FRRs of (a) 1%, (b) 5%, and (c) 10%]. Missing profiles indicate EER = FRR = FAR = 0%. Also illustrated in (d) is a DET profile for the independent countermeasure.

The texture-based countermeasure delivers significant improvements in robustness to spoofing at almost every operating point. For example, when the countermeasure is tuned to an FFR of between 1 and 5% as illustrated in Figure 12(a) and (b), and for a fixed FRR of 1%, then the FAR under a spoofing attack drops from 80% to just over 5%. However, if the countermeasure is tuned to a higher FFR such as 10% as illustrated in Figure 12(c), then the FAR increases to approximately 20%. This is high-security configuration. Even if almost no spoofing trials are misclassified as genuine trials, the countermeasure misclassifies 10% of genuine trials as spoofed trials.

In this case, the countermeasure degrades usability; this effect is illustrated by the blue profile in Figure 12(c). There is an inherent tradeoff between the security and usability, which is heavily

dependent on the modality, recognition system, countermeasure, and database. For this particular configuration, a sensible compromise is achieved when the countermeasure FFR is tuned to 5%. Here, 5% of genuine trials will be misclassified as spoofed trials, whereas for all but the highest FRRs, only 5% of spoofing attacks will succeed. Between 40 and 80% of attacks would otherwise be successful without the texture-based countermeasure.

WHILE POTENTIALLY MORE INTRUSIVE, CHALLENGE-RESPONSE COUNTERMEASURES CAN BE COMPLEMENTARY TO THE MORE TRADITIONAL APPROACHES TO LIVENESS DETECTION.

LESSONS LEARNED

The evaluation methodology presented in this article is based upon work funded through the European TABULA RASA project. Through a collaboration involving a large team of researchers, the study included a wide range of biometric modalities including face (2-D, 3-D, and multispectral), voice, gait, fingerprint, iris, vein, and electrophysiological signals. The study established state-of-the-art authentication technologies for each modality, investigated the vulnerabilities of each to spoofing, and proposed novel countermeasures integrated and evaluated according to the methodology outlined in the sections "Evaluation Methodology: Vulnerabilities to Spoofing" and "Evaluation Methodology: Spoofing Countermeasures." The findings described in the face-recognition case study are illustrative of the general trend across all modalities considered in TABULA RASA. The following discusses some of the lessons learned and the most pressing directions for future research.

VULNERABILITIES

Unless they are equipped with suitable countermeasures, all biometric systems were shown to be vulnerable to spoofing. Even so, some modalities (e.g., gait) are more robust than others (e.g., fingerprint), however, this should not be interpreted as meaning they are more reliable; in the absence of spoofing, fingerprint recognition generally outperforms gait recognition. Multimodal biometric systems are also vulnerable and can be overcome by the spoofing of only a single modality [13].

METRICS AND PROTOCOLS

Spoofing and countermeasure assessment is considerably more complex than might first appear. Countermeasures tend to be developed and evaluated independently from recognition systems and, although this methodology supports the comparison of different antispoofing approaches, it is the influence of countermeasures on overall system performance that is of greatest interest. New, standard metrics and protocols reflecting the robustness of integrated biometric and countermeasure systems should be adopted in the future.

USABILITY

While countermeasures are successful in reducing the vulnerability of biometric systems to spoofing, increased robustness comes at the expense of increased FRR. This impact on usability is generally manageable; increases in the FRR are usually negligible in comparison to increases in the FAR caused by spoofing. It is stressed, however, that the tradeoff is dependent on the modality and application.

GENERALIZATION

A large number of antispoofing studies have been reported in the literature, and encouraging results have been detailed for a small number of standard databases, e.g., [14]–[16]. Spoofing attacks are, however, varying and unpredictable in nature and it is difficult to predict how countermeasures will generalize to spoofing attacks "in the wild," where their true nature can never be known with certainty. As the field evolves, and as new forms of spoofing emerge, it will be necessary to collect databases with increasingly challenging and diverse spoofing trials, calling for new and increasingly generalized countermeasures. An alternative or complementary strategy involves one-class classification approaches [17]. As a form of anomaly detection that relies only on models of genuine data, they offer greater potential for generalization.

COUNTERMEASURE FUSION

The growing interest in spoofing and countermeasures and the number of diverse countermeasure strategies and algorithms in the literature lends support to research in fused countermeasures. While not increasing robustness to specific attacks, fused countermeasures offer a flexible antispoofing framework with which newly emerging vulnerabilities can be quickly patched with adaptable fusion strategies or the addition of new countermeasures.

CHALLENGE-RESPONSE COUNTERMEASURES

Interactive, challenge-response countermeasures require a user to perform a randomized, specific action such as smiling or blinking. The correct response to the given challenge infers liveness. While potentially more intrusive, challenge-response countermeasures can be complementary to the more traditional approaches to liveness detection. Further work is required to validate their potential.

COMBINED HARDWARE- AND SOFTWARE-BASED COUNTERMEASURES

The majority of previous work, including all that within TABULA RASA, has investigated software-based countermeasures. Hardware-based countermeasures and new sensors have great potential to detect more elaborate spoofing attacks, and a new generation of countermeasures that combine hardware- and software-based approaches may be needed.

OUTLOOK

Until relatively recently, the main focus of biometrics research has centered on the discrimination between genuine and zero-effort impostor trials. In the face of well-acknowledged vulnerabilities, greater effort to develop countermeasures will be required in the future to defend against concerted-effort spoofing attacks. The

growing body of related literature and number of competitive evaluations is testament to the increasing interest and importance of this research.

The past work is characterized by the lack of a standard evaluation methodology, which is needed to assess the influence of countermeasures on biometric system performance. This article presents what is, to the best of our knowledge, the first proposal for a formal evaluation methodology. Needless to say, however, further work is required to extend and adapt the methodology in view of the lessons learned through recent work.

While the performance of spoofing countermeasures proposed so far gives cause for optimism, their generalization to previously unseen spoofing attacks remains unclear. This leads to a number of research directions for the future, including networks of independent, fused countermeasures and one-class classification strategies. As the field evolves, new and more challenging databases will be essential for biometric system developers to stay one step ahead of the fraudsters. While extremely difficult, it will be critical to estimate the reliability of countermeasures in practical application scenarios including, not just their ability to detect spoofing, but also their impact on system usability.

ACKNOWLEDGMENTS

The following support is gratefully acknowledged: the European Commission (TABULA RASA: grant agreement number 257289, BEAT: grant agreement number 284989); the Academy of Finland; and the Spanish MICINN (BIO-SHIELD: grant agreement number TEC2012–34881).

AUTHORS

Abdenour Hadid (hadid@ee.oulu.fi) received the doctor of science in technology degree in electrical and information engineering from the University of Oulu, Finland, in 2005. He has been an adjunct professor and academy research fellow in the Center for Machine Vision Research of the University of Oulu since 2010. His research focuses on computer vision and pattern recognition with a particular interest in face analysis and biometrics. He made significant contributions to the state of the art, and his work is gaining increasing interest in the scientific community. According to Google Scholar (as of May 2015), his h-index is 24 and his work has been cited more than 5,600 times.

Nicholas Evans (evans@eurecom.fr) is an assistant professor at EURECOM, in Sophia Antipolis, France. He received the M.Eng. (1999) and Ph.D. (2003) degrees from the University of Wales, Swansea, United Kingdom. He was the lead guest editor of *IEEE Transactions on Information Forensics and Security* (special issue on biometric spoofing and countermeasures) as well as this special issue of *IEEE Signal Processing Magazine* on biometric security and privacy. He is an associate editor of *EURASIP Journal on Audio, Speech, and Music Processing* and is a member of the IEEE Speech and Language Technical Committee.

Sébastien Marcel (marcel@idiap.ch) is a researcher at the Idiap Research Institute, Switzerland, where he heads the biometrics group. He is also a lecturer at the Ecole Polytechnique Fédérale de Lausanne. He is an associate editor of *IEEE Transactions on Information Forensics and Security* (*T-IFS*), coeditor of the *Handbook of Biometric Anti-Spoofing*, and was a guest editor of *T-IFS* (special issue on biometric spoofing and countermeasures) and of this special issue of *IEEE Signal Processing Magazine* on biometric security and privacy.

Julian Fierrez (julian.fierrez@uam.es) received the M.Sc. and Ph.D. degrees from Universidad Politécnica de Madrid, Spain, in 2001 and 2006, respectively. Since 2004 he has been affiliated with Universidad Autonoma de Madrid, where he is currently an associate professor. From 2007 to 2009, he was a Marie Curie postdoctoral researcher at Michigan State University. His research interests include image processing, pattern recognition, authentication using biometrics such as fingerprints and handwritten signatures, and security of person authentication systems. He is actively involved in multiple EU projects and has received multiple research distinctions including the EURASIP Best Ph.D. Award in 2012.

REFERENCES

[1] A. Jain, A. Ross, and S. Pankati, "Biometrics: A tool for information security," *IEEE Trans. Inform. Forensics Security*, vol. 1, no. 2, pp. 125–143, June 2006.

[2] S. Marcel, M. Nixon, and S. Z. Li, *Handbook of Biometric Anti-Spoofing*. New York: Springer, 2014.

[3] E. Marasco and A. Ross, "A survey on antispoofing schemes for fingerprint recognition systems," *ACM Comput. Surv.*, vol. 47, no. 2, pp. 28:1–28:36, Nov. 2014.

[4] J. Galbally, S. Marcel, and J. Fierrez, "Biometric anti-spoofing methods: A survey in face recognition," *IEEE Access*, vol. 2, pp. 1–23, Dec. 2014.

[5] Z. Wu, N. Evans, T. Kinnunen, J. Yamagishi, F. Alegre, and H. Li, "Spoofing and countermeasures for speaker verification: A survey," *Speech Commun.*, vol. 66, no. 0, pp. 130–153, 2015.

[6] J. Fierrez, "Adapted fusion schemes for multimodal biometric authentication," Ph.D. dissertation, Universidad Politecnica de Madrid, May 2006.

[7] I. Chingovska, A. Anjos, and S. Marcel, "Biometrics evaluation under spoofing attacks," *IEEE Trans. Inform. Forensics Security*, vol. 9, no. 12, pp. 2264–2276, Dec. 2014.

[8] S. Z. Li and A. K. Jain, Eds., *Handbook of Face Recognition*, 2nd ed. New York: Springer, 2011.

[9] L. El-Shafey, C. McCool, R. Wallace, and S. Marcel, "A scalable formulation of probabilistic linear discriminant analysis: Applied to face recognition," *IEEE Trans. Pattern Anal. Mach. Intel.*, vol. 35, no. 7, pp. 1788–1894, July 2013.

[10] J. Galbally, S. Marcel, and J. Fierrez, "Image quality assessment for fake biometric detection: Application to iris, fingerprint and face recognition," *IEEE Trans. Image Processing*, vol. 23, no. 2, pp. 710–724, Feb. 2014.

[11] J. Määttä, A. Hadid, and M. Pietikäinen, "Face spoofing detection from single images using texture and local shape analysis," *IET Biometrics*, vol. 1, no. 1, pp. 3–10, 2012.

[12] T. de Freitas Pereira, J. Komulainen, A. Anjos, J. M. D. Martino, A. Hadid, M. Pietikäinen, and S. Marcel, "Face liveness detection using dynamic texture," *EURASIP J. Image Video Process.*, vol. 2014, no. 2, Jan. 2014.

[13] G. L. Marcialis, G. Fumera, and B. Biggio, "Anti-spoofing: Multimodal," *Encyclopedia of Biometrics*. New York: Springer, 2014.

[14] L. Ghiani, D. Yambay, V. Mura, S. Tocco, G. L. Marcialis, F. Roli, and S. Schuckers, "Livdet 2013—Fingerprint liveness detection competition," in *Proc. IEEE/IAPR Int. Conf. Biometrics*. IEEE Press, June 2013, pp. 1–6.

[15] I. Chingovska, J. Yang, Z. Lei, D. Yi, S. Z. Li, O. Kahm, C. Glaser, N. Damer, et al., "The 2nd competition on counter measures to 2D face spoofing attacks," in *Proc. IEEE/IAPR Int. Conf. Biometrics*. IEEE Press, June 2013, pp. 1–6.

[16] Z. Wu, T. Kinnunen, N. Evans, J. Yamagishi, C. Hanilci, M. Sahidullah, and A. Sizov, "ASVspoof 2015: The first automatic speaker verification spoofing and countermeasures challenge," in *Proc. INTERSPEECH*, 2015.

[17] F. Alegre, A. Amehraye, and N. Evans, "A one-class classification approach to generalised speaker verification spoofing countermeasures using local binary patterns," in *Proc. IEEE 6th Int. Conf. Biometrics: Theory, Applications and Systems (BTAS 2013)*, 2013, pp. 1–8.

[SP]

Battista Biggio, Giorgio Fumera, Paolo Russu, Luca Didaci, and Fabio Roli

Adversarial Biometric Recognition

A review on biometric system security from the adversarial machine-learning perspective

©ISTOCKPHOTO.COM/GREYFEBRUARY

Biometrics Security and Privacy Protection

In this article, we review previous work on biometric security under a recent framework proposed in the field of adversarial machine learning. This allows us to highlight novel insights on the security of biometric systems when operating in the presence of intelligent and adaptive attackers that manipulate data to compromise normal system operation. We show how this framework enables the categorization of known and novel vulnerabilities of biometric recognition systems, along with the corresponding attacks, countermeasures, and defense mechanisms. We report two application examples, respectively showing how to fabricate a more effective face spoofing attack, and how to counter an attack that exploits an unknown vulnerability of an adaptive face-recognition system to compromise its face templates.

Digital Object Identifier 10.1109/MSP.2015.2426728
Date of publication: 13 August 2015

INTRODUCTION

Adversarial machine learning is a novel research field that was born in response to the increasing use of pattern recognition and machine-learning techniques, including signal processing ones, in security-related applications such as biometric identity recognition,

spam, and malware detection. In these applications, intelligent and adaptive adversaries are interested in subverting system operation; e.g., nonauthorized users may aim to gain access to a resource secured by a biometric identity recognition system. Despite the fact that pattern recognition and machine-learning algorithms have enabled the development of more effective recognition systems, they have not been originally designed to operate in adversarial settings. In particular, their underlying assumption of data stationarity (i.e., that training and testing data follow the same distribution) is likely to be violated in adversarial environments. As a consequence, these algorithms can introduce additional, specific vulnerabilities that can be exploited by carefully crafted attacks to cause different security violations, including denial of service and missed detection of intrusive attempts. This may eventually compromise the whole system security. Even if countermeasures and novel algorithms have been proposed to improve security against these sophisticated attacks, they will not stop adversaries from developing novel ways of misleading such defense systems, engendering a long-lasting arms race.

To date, research efforts in adversarial machine learning have focused on identifying different kinds of potential attacks against machine-learning and pattern recognition algorithms, and developing the corresponding countermeasures to improve robustness in adversarial settings [1]–[4]. According to the security-by-design paradigm, ongoing work is also addressing the issue of extending learning theory and methods to explicitly account for the presence of malicious adversaries that can undermine algorithm operation; e.g., in [2] and [3], traditional performance evaluation methods have been extended to allow for a systematic evaluation of the security of pattern classifiers in adversarial settings, providing a better understanding of the system performance both in the absence and in the presence of well-crafted attacks. The relevance of these issues is also witnessed by an increasing number of publications and events, e.g., the Neural Information Processing Systems Workshop on Machine Learning in Adversarial Environments for Computer Security [5] and the more recent Dagstuhl Perspectives Workshop on Machine-Learning Methods for Computer Security [6].

Biometric identity recognition is a clear example of a widespread and still growing application field in which security is a key issue and pattern recognition techniques play a major role. Different vulnerabilities of biometric systems, specific attacks that can exploit them, and corresponding countermeasures have been analyzed in the literature [7], [8] and in research projects. For instance, the recent EU FP7 Tabula Rasa project has carried out an extensive analysis on spoofing attacks (i.e., attacks involving the submission of a fake biometric trait, like a gummy finger, to impersonate an authorized user), and on the development of possible countermeasures, like liveness detection techniques. Moreover, several approaches for the analysis and assessment of biometric systems security have been proposed; see, e.g., the ISO/IEC 19792:2009 implementation specifics for the security evaluation of biometrics, and the National Institute of Standards and Technology Common Criteria for Information Technology Security Evaluation. However, all existing efforts disregarded the potential, specific vulnerabilities introduced by pattern recognition algorithms used in biometric systems, and thus the investigation of the corresponding attacks and countermeasures. We argue that looking at biometric system security from the perspective of adversarial machine learning not only provides an original categorization of existing attacks against such systems, but it also allows us to consider more sophisticated attacks targeting vulnerabilities of the learning algorithms used in these systems, along with the countermeasures already proposed in the field of adversarial machine learning.

Based on the above motivations, in the following we first provide a concise overview of adversarial machine learning, to introduce kindly the readers to this recent research field; we use popular attacks against biometric systems, such as spoofing attacks, as running examples to make our explanation clearer. We then review the security of biometric identity recognition systems by showing how recent theoretical results and systematization efforts from this field enable: 1) the definition of a more complete taxonomy of attacks against biometric systems, based on a formal attacker's model explicitly accounting for her knowledge and capability, which allows one to identify novel attack scenarios associated to specific vulnerabilities of machine-learning and pattern recognition algorithms, besides encompassing known attacks; and 2) the design of the corresponding countermeasures, building on solutions proposed in adversarial machine learning, which can give rise to the design of novel, secure-by-design algorithms capable of improving adversarial biometric identity recognition. We finally discuss two application examples of the possible aforementioned achievements. We first show how a skilled attacker may fabricate more effective face-spoofing attacks, and then highlight a new vulnerability of adaptive biometric systems, devising the corresponding attack and a possible countermeasure.

The main goal of this article is to provide the readers of this magazine, and researchers in biometrics, a gentle introduction to adversarial machine learning, and a well-structured review of the state of the art on biometric security in light of the most recent findings in the area of adversarial machine learning.

ADVERSARIAL MACHINE LEARNING: AN OVERVIEW

During the last several decades, the increasing variability and sophistication of attack threats, in response to the growing complexity and amount of vulnerable attack points in security systems, has favored the adoption of machine-learning and pattern recognition techniques to timely detect variants of known and never-before-seen attacks. These techniques can, however, exhibit intrinsic vulnerabilities that can be exploited by skilled attackers, perpetuating their arms race against system designers. Adversarial machine learning aims at countering this phenomenon by focusing on vulnerabilities of learning algorithms. It attempts to anticipate the adversary's strategy by identifying novel threats and devising the corresponding countermeasures before system deployment. In practice, it follows a proactive rather than a reactive approach. The first step toward the aforementioned goal has been the proposal of a taxonomy categorizing attacks against learning algorithms along three axes [1], [2]: 1) the attack influence, which can be exploratory, if the adversary can only manipulate the testing data, or causative, if she can modify also the training data; 2) the attack specificity, which ranges from targeted to indiscriminate, depending on whether the classification of a set of specific samples

or any of them is affected by the attack; 3) the security violation, which can be an integrity violation, if the adversary is allowed to access a restricted service or resource (e.g., an impostor gaining access to a genuine client's account [3], [9]); an availability violation, if legitimate users are denied access or normal system operation is compromised (e.g., misclassifying legitimate e-mails as spam); and a privacy violation, if the adversary is able to exploit confidential information about the system (e.g., the clients' templates in a biometric recognition system [10]–[12]).

To date, several vulnerabilities and attacks against different learning algorithms (e.g., support vector machines and neural networks) have been investigated, along with the proposal of possible countermeasures [13]–[15]. We will summarize the main existing countermeasures in the remainder of this article, discussing how they can be exploited to improve the security of machine-learning algorithms used in biometric systems. In particular, to anticipate the adversary's strategy, the existing work in adversarial learning simulates attacks, based on more or less explicit models of the adversary. We recently formalized this approach within a general framework, proposing a formal model to characterize the adversary's behavior [3], [4]. We summarize our model next and exploit it in the remainder of this article to characterize and understand security of biometric systems under an adversarial machine-learning perspective.

Our model generalizes and encompasses other models proposed in the area of adversarial machine learning [1], [2], making explicit assumptions on the attacker's goal, knowledge of the targeted system, and capabilities of manipulating the input data or the system's components.

The goal has to be defined according to the desired security violation and attack specificity. Knowledge of system components (e.g., the kind of decision function and its parameters, or how a component operates) can be perfect or limited, and feedback on the classifier's decisions can also be exploited [13], [16], [17]. The capability is defined as the attack influence, based on how the adversary can affect training and testing data (e.g., which features can be manipulated and how, according to application-specific constraints). An attack strategy can be then defined based on the previous elements, to implement the attack. In formal terms, we assume the attacker's goal to be quantified by a function $g(a, \theta)$ measuring the extent to which an attack strategy a from a feasible strategy set $\mathcal{A}$ fulfills the attacker's goal. The feasible set $\mathcal{A}$ has to be defined according to the attacker's capability of manipulating the input data and system components, while the attacker's knowledge is encoded by the parameter vector $\theta \in \Theta$. Under this setting, the optimal attack strategy corresponds to the solution of the following optimization problem:

$$\max_{a \in \mathcal{A}} g(a; \theta). \tag{1}$$

Although this formulation may seem rather abstract at this stage, it enables us to consider trivial and sophisticated attacks under a consistent view, as shown in the remainder of this article.

BIOMETRIC RECOGNITION SYSTEMS UNDER ATTACK

Biometric recognition systems operate either in enrollment or in recognition mode [8]. During enrollment, each client provides his biometric traits and identity, in the presence of a human supervisor. A set of reference templates for each client is then stored in the template database along with the corresponding identity. During recognition, the biometric system is expected to recognize a previously enrolled client by comparing the submitted traits with those stored in the template database. Biometric systems may operate either in a verification or in an identification setting. In verification settings, biometric systems are often used to control access to protected resources, including confidential information or services. A user aiming to access them has to provide his/her biometric trait and claim an identity. The system then verifies whether the claim is genuine (i.e., the user's identity is the claimed one) or not (i.e., the user is an impostor trying to impersonate another client), and allows access only in the former case. This procedure is illustrated in Figure 1, which is general enough to also account for multibiometric systems; in this case, the fusion

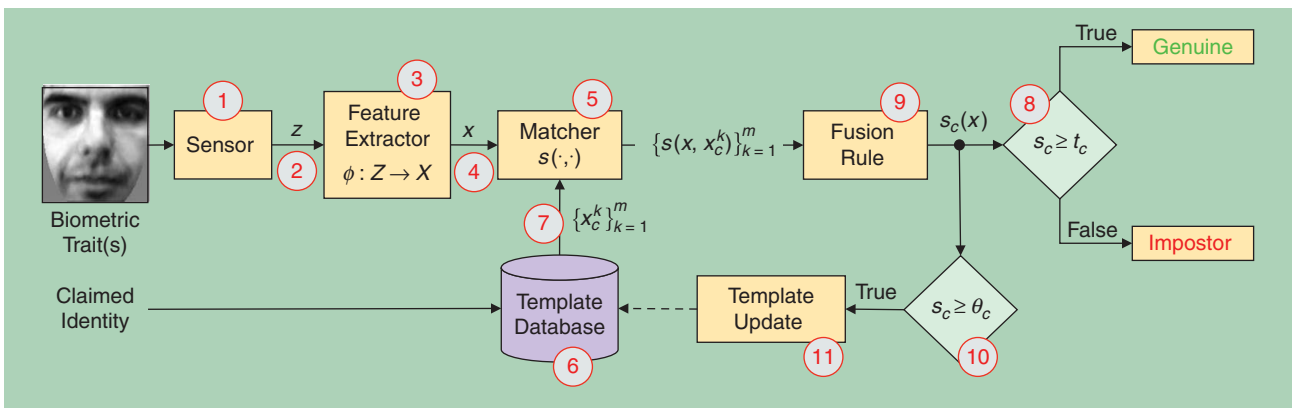

[FIG1] The architecture of a biometric verification system and corresponding attack points, highlighted with red-circled numbers. During verification, the image $z \in \mathcal{Z}$ (e.g., a face image) acquired by the sensor is processed by a feature extractor $\phi: \mathcal{Z} \mapsto \mathcal{X}$ to obtain a compact representation $x \in \mathcal{X}$ (e.g., a graph). The templates $\{x_c^k\}_{k=1}^m$ of the claimed identity c are retrieved from the template database, and compared to x using a matching algorithm $s: \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$. The resulting scores $\{s(x, x_c^k)\}_{k=1}^m$ are combined by a fusion rule, producing an aggregated score $s_c(x)$ that expresses the degree to which x is likely to belong to c. The score $s_c(x)$ is then compared with a decision threshold t_c to decide whether the claim is genuine or impostor. If a template self-update is implemented, and $s_c(x)$ is not lower than a self-update threshold θ_c, one of the templates in $\{x_c^k\}_{k=1}^m$ is updated depending on x, according to a given policy.

rule s_c should aggregate the matching scores coming from all biometric traits. In identification settings, instead, no identity claim is made: a user provides only the requested biometric trait $\boldsymbol{x}$, and the system is expected to correctly recognize the corresponding identity among those in the template database, by matching $\boldsymbol{x}$ against all the known clients' templates; the corresponding scores $s_c(\boldsymbol{x})$ are then sorted in descending order to provide a list of the most likely candidate identities.

To account for natural changes of biometric traits over time (i.e., biometric aging), and changes in the environmental or acquisition conditions during verification, adaptive biometric systems have been proposed. They enable an update of the stored templates automatically during verification [18], [19]. A popular technique is template self-update: if the submitted trait is sufficiently similar to the reference templates of the claimed identity, one of such templates is updated by exploiting information coming from the submitted trait, according to a given policy. A simple update policy is called *nearest-out self-update*, as it replaces the most similar template to the submitted trait with the latter [20], [21].

THE ATTACK SURFACE

Previous work has identified the main attack points and vulnerabilities of biometric recognition systems, along with the corresponding attacks [7], [8]. First, any system is subject to intrinsic failures not produced by adversarial attempts, i.e., rejected genuine claims and accepted zero-effort impostors (i.e., impostors that do not exert any special effort to intrude). Besides this, a number of adversarial attacks have been also considered in early work, leading to the identification of eight potentially vulnerable attack points highlighted by the red-circled numbers 1–8 in Figure 1 [7]. Additionally, we consider points 9–11: they correspond to vulnerabilities of adaptive biometric systems that update clients' templates during operation, which we recently exploited to implement a template poisoning attack [20], [21]. The set of all attack points defines the attack surface of a biometric recognition system. The corresponding attacks can be categorized into four main groups according to the targeted system component [8]: 1) attacks to the sensor (point 1), 2) to interfaces and channels connecting different modules (points 2, 4, and 7), 3) to processing modules and algorithms (points 3, 5, and 8–11), and 4) to the template database (point 6). We discuss them below, along with the corresponding countermeasures proposed so far, which are also summarized in Table 1. It is also worth mentioning here a special category of attacks, known as *insider attacks*, where the attacker is colluded with a system administrator or exercises coercion to escalate privileges [8].

Spoofing attacks consist of fabricating a fake biometric trait to impersonate an enrolled client. They target the sensor (point 1), so they are also referred to as direct attacks. Current defenses are based on liveness detection methods, which aim to verify whether the submitted trait is "alive" or "fake" by looking at specific patterns (e.g., perspiration patterns during fingerprint acquisition, or eye blinking during face verification). Multibiometric systems have been also proposed as a defense; however, to avoid spoofing them by only using a single fake trait, the matching scores coming from the different traits should be properly combined, using a secure score-level fusion rule [9], [22].

Replay attacks can be staged at interfaces between modules by replaying a stolen image of the biometric trait of the targeted client to the feature extractor (point 2), or directly the corresponding feature values to the matcher (point 4). An attacker may even replay a signal to replace the features of a given template of the claimed identity (point 7). This attack can be clearly staged if the corresponding communication channels are insecure, but also over encrypted channels, as the encrypted signal can be stolen and replayed into the channel directly. This can be avoided by encrypting a time stamp into the signal, or using challenge-response mechanisms. Another possible countermeasure is physical isolation, to avoid sending data over insecure channels (e.g., the Internet) subject to man-in-the-middle attacks. A popular example of physical isolation is the use of smart cards performing match-on-card operations. However, this technique has its own disadvantages, including limitations in terms of computational resources and memory, and the fact that the user should always use the smart card to be authenticated [8].

Hill-climbing attacks, similarly to replay ones, affect insecure communication channels between modules, and, in particular, points 2 and 4 in Figure 1. Their goal is to reconstruct a template image by iteratively sending a bunch of slightly perturbed images to the feature extractor (point 2), or their features to the matcher (point 4), and retaining the one that maximizes the matching score $s_c(\boldsymbol{x})$, where $\boldsymbol{x}$ is the current image (or set of features) submitted by the attacker. In practice, it is a gradient-ascent technique that approximates the gradient of $s_c(\boldsymbol{x})$ numerically. In this case, the attacker is assumed to be able to observe $s_c(\boldsymbol{x})$ for any queried image, which may only be feasible if the system provides (or leaks) such information. Besides the aforementioned channel protection schemes, an additional defense mechanism consists of quantizing

[TABLE 1] THE CATEGORIZATION OF ATTACKS AND COUNTERMEASURES FOR BIOMETRIC SYSTEMS. FOR EACH ATTACK TECHNIQUE, WE ALSO REPORT THE TARGETED COMPONENT (ATTACK LOCATION) AND THE ATTACK POINT(S), ACCORDING TO FIGURE 1.

ATTACK TECHNIQUE	ATTACK LOCATION	ATTACK POINT(S)	DEFENSE
SPOOFING	SENSOR	1	LIVENESS DETECTION, MULTIBIOMETRICS (SECURE FUSION)
REPLAY	INTERFACES/CHANNELS	2, 4, 7	ENCRYPTED CHANNEL, TIME STAMP, CHALLENGE-RESPONSE, PHYSICAL ISOLATION
HILL CLIMBING	INTERFACES/CHANNELS	2, 4	ENCRYPTED CHANNEL, TIME STAMP, CHALLENGE-RESPONSE, PHYSICAL ISOLATION, SCORE QUANTIZATION
MALWARE INFECTION	MODULES/ALGORITHMS	3, 5, 8–11	SECURE CODE, SPECIALIZED HARDWARE, ALGORITHMIC INTEGRITY
TEMPLATE THEFT, SUBSTITUTION, AND DELETION	TEMPLATE DATABASE	6	TEMPLATE ENCRYPTION, CANCELABLE/REVOKABLE TEMPLATES

the matching score to provide less accurate information to the attacker. However, attacks based on more sophisticated black-box optimization techniques, suited to quantized objective functions, can also be considered to make these countermeasures ineffective [10].

The algorithmic implementations of the software modules (points 3, 5, and 8–11) may exhibit vulnerabilities that can be exploited by a skilled attacker through well-known hacking techniques (e.g., buffer overflow), to install malicious software, i.e., malware, including worms, trojan horses, etc. This threat can be avoided or mitigated by exploiting well-known programming practices, like secure code programming, or using specialized hardware to perform some critical operations [8]. A secure programming practice is to check algorithmic integrity, i.e., that each algorithm and function correctly handles any input parameter and never shows any unexpected behavior. For instance, if the matching algorithm expects a vector $\boldsymbol{x} \in \mathbb{R}^d$ as input, and instead receives an input with a different format, is it going to crash or provide an output anyway? In the latter case, how is such an output handled by the subsequent modules? Does it lead to accepting by error the given claim as genuine or not?

Template theft, substitution, and deletion attacks target the template database (point 6). If templates are not protected properly, one may be able to steal them, and use them to create a spoof (i.e., a fake template), to perform a replay attack, or to impersonate the targeted client on a different system and perform other operations, e.g., searching on protected databases (function creep) [8]. Another possibility is to replace a template to impersonate a client without requiring any sophisticated attack as spoofing or replay; e.g., an attacker may add his own fingerprint template to the set of templates belonging to another client. Additionally, templates of a given client can be deleted to cause a denial of service, i.e., to avoid the targeted client to be recognized successfully. Countermeasures include template encryption, and also the use of cancelable/revokable templates, which can be used only on a specific system and reissued if stolen. The idea is to encode the templates using a key or pin code that can be changed to re-enroll the user and create a novel, different encrypted template [8].

ADVERSARIAL BIOMETRIC RECOGNITION

Here we analyze biometric system security in terms of our previously discussed framework by making assumptions on the adversary's goal, knowledge, capability, and attack strategy that are suited to biometric applications. Our aim is threefold: 1) to provide a well-structured categorization of the vulnerabilities of biometric systems and of the corresponding attacks, also through the definition of different, pertinent attack scenarios; 2) to provide a formal characterization of existing attacks within our framework and envision more sophisticated and effective attack strategies; and 3) to identify suitable countermeasures and defenses inspired by previous work on adversarial machine learning.

- *Adversary's goal*: It is defined in terms of security violation and attack specificity. Biometric system security can be violated by an attacker that aims at impersonating a genuine user (integrity violation); at compromising the template galleries of genuine users to deny them access to the system, causing a denial of service (availability violation); or at violating the privacy of genuine users, e.g., by inferring their templates through a hill-climbing attack (privacy violation). The attack specificity can be targeted, if the attack targets a specific set of clients, or indiscriminate, if any client may be affected.
- *Adversary's knowledge*: It is defined by leveraging on the definition of the attack surface of biometric systems given in the previous section by making specific assumptions on what the attacker knows of the system components and how they work. According to Figure 1 and Table 1, the attacker may know 1) the kind of sensor used (point 1), e.g., an optical or capacitive fingerprint sensor; 2) which interfaces/channels are used to implement connections (points 2, 4, and 7), e.g., if an insecure channel over the Internet is used to send the acquired images to the feature extractor (point 2); 3) how the modules/algorithms work, and whether they are vulnerable or not (points 3, 5, and 8–11), in particular, the feature mapping ϕ (point 3), the matching algorithm s (point 5), the decision threshold t_c (point 8), the fusion rule s_c (point 9), and, if the template update is implemented, the self-update threshold θ_c (point 10) and the template update policy (point 11); and 4) some of the templates stored in the template database (point 6). The attacker may also be able to collect images of the same biometric traits using other techniques; e.g., acquiring latent fingerprints, or collecting face images of the targeted clients from social networks. From a machine-learning perspective, this amounts to having different levels of knowledge of the classifier's training data. In practice, it is worth noting that attackers typically have limited knowledge of the sensors and algorithms used, of the users' templates, and of any other system components (e.g., communication channels, template encryption schemes, etc.). Several previous works have, however, considered vulnerabilities of biometric systems without clearly pointing out the underlying assumptions on the adversary's knowledge required to perform the corresponding attack. Under our framework, such assumptions become clearly explicit.
- *Adversary's capability*: This can also be defined in terms of the attack location: 1) the sensor (point 1); 2) interfaces/channels (points 2, 4, and 7); 3) the internals or even the output of modules and algorithms (points 3, 5, and 8–11), e.g., through malware infection attacks; and 4) the template database (point 6). In addition, one has to define the attack influence, i.e., the capability of manipulating the input data (e.g., using fake biometric traits), and how such data may be used to update the system (e.g., in adaptive biometric systems the attacker can produce spoofing attacks that can subsequently poison the clients' templates [20], [21]). Accordingly, the attack can influence only verification, or also enrollment/update.
- *Attack strategy*: According to the adversary's goal [generally expressed in terms of an objective function $g(a;\theta)$], knowledge (given in terms of the parameter vector $\theta \in \Theta$) and capability (which defines the feasible set of attack strategies $a \in \mathcal{A}$), an optimal attack strategy can be defined to implement the attack, as explained before; see (1). For instance, assume that the attacker aims to impersonate an enrolled client (integrity targeted attack), and she is only able to acquire a latent fingerprint

[FIG2] A conceptual representation of the adversary model and of the main attack scenarios (given in terms of the corresponding security violation and attack specificity) according to our framework.

of the client, without having any other knowledge of the system components and algorithms. Then, the corresponding optimal attack strategy amounts to fabricating a fake fingerprint that is as similar as possible to the latent one, and using it to perform a spoofing attack. In this case, $g(a;\theta)$ can be regarded as a measure of the similarity between the fake and the latent fingerprint, as θ only contains information related to the latent fingerprint, and a corresponds to the fake fingerprint. A more skilled attacker may, however, also know the matching algorithm s and the fusion rule s_c used by the system, and may be able to collect more than a single fingerprint image of the targeted client. As an application example of our framework, we will show that, under this setting, more sophisticated and effective spoofing attacks can be fabricated. As another application example, we will also consider a poisoning attack against an adaptive face verification system and propose a novel countermeasure based on the sanitization of the client's templates.

In the following, we define a set of representative attack scenarios to categorize known attacks according to our framework. The framework and the considered attack scenarios are also represented in Figure 2.

CATEGORIZATION OF BIOMETRIC ATTACK SCENARIOS

Previous work has categorized attacks to biometric systems and countermeasures simply in terms of the attack points of Figure 1 (e.g., spoofing attacks to the sensor, countered by liveness detection techniques, and attacks to a compromised channel, countered by channel encryption). Instead, looking at them from the broader perspective opened by our framework allows us to identify three main attack scenarios, described in the following, in which the attacks and countermeasures discussed in the previous sections play different roles. A few examples of known attack and defenses are categorized in Table 2 in terms of these attack scenarios. This also clears the way both to identify novel, more sophisticated attacks against biometric systems, and to adapt the corresponding countermeasures from the adversarial learning field.

EVASION

The goal of this attack scenario is to impersonate a client (integrity, targeted/indiscriminate attack). To this end, knowledge of the client's biometric trait is required, e.g., to create a fake trait or to carry out a replay attack. Attacks exploiting perfect knowledge of the targeted client's biometric trait include the so-called consensual method (in which the targeted client voluntarily provides the required biometric trait to the attacker), and template stealing. Conversely, exploiting a latent fingerprint is an example of limited knowledge, since the attacker may only partially know or observe the required biometric trait. A limited knowledge about the rest of the biometric system can also be sufficient. In this scenario, the capability of the attacker consists of manipulating data during the verification step, whereas no influence on the enrollment/update step is assumed (in particular, she can not access the template database). Most frequently, the attack

[TABLE 2] EXAMPLES OF CATEGORIZATION OF PREVIOUS WORK ON BIOMETRIC SECURITY ACCORDING TO THE THREE MAIN ATTACK SCENARIOS DEFINED IN ADVERSARIAL MACHINE LEARNING: EVASION, POISONING, AND PRIVACY ATTACKS.

	GOAL		KNOWLEDGE	CAPABILITY		ATTACK STRATEGY
	VIOLATION	SPECIFICITY		INFLUENCE	LOCATION	
EVASION ATTACKS						
MATSUMOTO et al. [23], RODRIGUES et al. [9], JOHNSON et al. [24]	INTEGRITY	TARGETED, INDISCRIMINATE	PERFECT (CONSENSUAL FAKE)	VERIFICATION	SENSOR	SPOOFING
POISONING ATTACKS						
BIGGIO et al. [21, 20]	INTEGRITY, AVAILABILITY	TARGETED, INDISCRIMINATE	PERFECT, LIMITED (UNKNOWN TEMPLATES)	ENROLLMENT/ UPDATE	SENSOR	SPOOFING (FACE)
PRIVACY ATTACKS						
ADLER [10], GALBALLY et al. [11], MARTINEZ et al. [12]	PRIVACY	TARGETED, INDISCRIMINATE	LIMITED (UNKNOWN TEMPLATES)	VERIFICATION	MATCHER	HILL CLIMBING

strategy corresponding to an evasion attack consists of submitting a fake trait (spoof) to the sensor (point 1), or of replaying the acquired image (point 2) into the system. In rarer cases (disregarded here), the biometric system can be infected by malware, potentially allowing the attacker to arbitrarily manipulate the functionality or the output of any system component.

POISONING

This nontrivial attack scenario has been originally defined in the context of adaptive biometric systems in our previous work [20], [21], inspired by our adversarial learning framework, and by work on poisoning learning algorithms [1]–[4]. The goal of poisoning attacks can be either an integrity or availability security violation; it can be either targeted to a specific client, or indiscriminate (see the section "Poisoning Biometric Systems that Learn from Examples"). The adversary's knowledge can be perfect or limited, depending on whether each of the system's components is exactly known to the attacker. More precisely, the attacker may have perfect (or limited) knowledge of each of the components discussed in Figure 1, including the targeted clients' templates, the matching algorithm, the template update algorithm, and the decision and update thresholds. The attacker's capability consists of modifying the template database, either by directly manipulating it (e.g., through malware infection), or, more realistically, by submitting fake traits that are erroneously used to update the template gallery of a given client. In terms of security violation, an integrity violation thus amounts to replacing a victim's template with an attacker's template or to adding an attacker's template in the victim's gallery. This indeed allows the attacker to impersonate the victim without using any further spoofing or replay attack, but directly using her own biometric trait. The goal of an availability violation is to cause a denial of service, instead, by replacing or compromising the majority of templates in the victim's gallery. This will indeed deny the victim access to the system. Under this setting, the attack strategy amounts to compromising the template gallery either by introducing an attacker's template in the victim's gallery (i.e., integrity violation), or by compromising the maximum number of victim's templates (i.e., availability violation). If the template database can not be compromised directly, the attacker can produce a well-crafted sequence of fake traits to gradually drift the victim's template gallery toward the desired set of templates, while minimizing the number of fake traits required to complete the attack. An example of such an attack is given next.

PRIVACY

In this case, the goal is to retrieve confidential information (i.e., one or more templates) about either a given set of clients (targeted attack) or about any client (indiscriminate attack). This is typically a preliminary step before performing another kind of attack (evasion or poisoning), when no simpler way to retrieve information on the victims' templates exists (e.g., acquiring a face image through a social network, or a latent fingerprint). To this end, the attacker can gain knowledge from the system's feedback, e.g., the outcome of the verification decision (either accept or reject), or the score value $s_c(\boldsymbol{x})$ (as in hill-climbing attacks). In common settings (i.e., disregarding cases like malware infection), the capability consists of sending a number of query images through a remote channel and observing the available feedback; e.g., if the sensor and the matcher are remotely operating, and interconnected through the Internet, an attacker may perform a man-in-the-middle attack and send replayed images through the channel. It is clear that the attack strategy in this case corresponds to an hill-climbing attack.

SECURE-BY-DESIGN BIOMETRIC SYSTEMS

The considered adversary model can be exploited not only to provide a different categorization of known defense mechanisms for biometric recognition systems but also to identify novel countermeasures among those proposed in adversarial learning, which can help countering attacks against machine-learning and pattern recognition algorithms used in biometric systems. We discuss them below, with reference to the three aforementioned attack scenarios.

COUNTERING EVASION

In this scenario, the main attack strategies involve spoofing and replay. As reported in Table 1, the pertinent defenses are: liveness detection, multibiometric systems with secure score-level fusion rules, encrypted channels and time-stamp/challenge-response schemes, and physical isolation (e.g., match-on-smart-cards). Novel defense mechanisms can also be devised, inspired by the adversarial machine-learning field. In particular, to counter evasion attacks, one can consider secure learning techniques. They consist of modifying existing learning algorithms (and developing novel ones) that explicitly take into account a specific kind of adversarial data manipulation. They follow the paradigm of security by design, which advocates that a system should be designed from the ground up to be secure. In the context of biometric systems, secure learning techniques can be exploited to design trainable score-level fusion rules, such as those based on game theory, or on the framework of learning with invariances [14], [25]–[27]. Investigating this issue would be an interesting research direction for future work.

COUNTERING POISONING

Spoofing and replay are the main attack strategies that are also under this scenario and, thus, the same defenses listed in Table 1 can be also exploited in this case. In addition, other countermeasures can be considered, among those proposed in adversarial learning, to improve the security of the training phase in the presence of poisoning, which may occur when the system is retrained on data collected during operation [15], [28]. These include secure learning (similar to countermeasures to evasion) and data sanitization. In a biometric system, the latter consists of detecting outlying template updates that may compromise the template gallery of a given client, e.g., by adding an impostor's template to the targeted client's gallery, or by replacing some of the client's templates. We will give a concrete example of a novel defense based on template sanitization in the next section. We point out that these additional defenses can be considered complementary to those listed in Table 1, like liveness detection and channel encryption.

PRESERVING PRIVACY

Known defenses that can be exploited against attacks targeting the template database are mainly based on template encryption

schemes [8] (Table 1). Score quantization has been also proposed to counter hill-climbing attacks, but it has already been shown to be ineffective [10]. Moreover, attacks proposed in adversarial learning have already been capable of reverse engineering the classifier by only exploiting only feedback on its decisions [16], [17]; thus, even by only looking at genuine or impostor classifications, an attacker may be able to successfully perform a hill-climbing attack. Among the proposed countermeasures that have not yet been considered for biometric systems, it would be worth investigating in future work the ones based on randomization and disinformation. They follow the paradigm of security by obscurity, aiming to improve system security by hiding information to the attacker. They have been suggested in adversarial learning to counter reverse-engineering attacks. This can be achieved by denying access to the actual classifier or training data, and randomizing the classifier's output to give imperfect feedback to the attacker [1]–[4], [29].

APPLICATION EXAMPLES

Here we consider two application examples of our framework related to the development of sophisticated spoofing and poisoning attacks against face-verification systems, respectively.

IMPROVED FACE SPOOFING FROM MULTIPLE FACES

Let us assume we are given a face verification system that authenticates clients by matching the acquired face image against the template gallery of the claimed identity (consisting of n images acquired during enrollment), and then thresholding the corresponding average score. According to the architecture depicted in Figure 1, we assume that our system maps the submitted face image $z \in \mathcal{Z}$ onto a reduced vector space $\mathcal{X}$ using principal component analysis (PCA), and computes the matching score for client c as $s_c(x) = (1/n)\sum_{i=1}^{n} s(x, x_i)$, where $s(x, x_i) = \exp\{-\|x - x_i\|\}$, and x_i is the ith template of the claimed identity. We further assume the attack scenario detailed below.

- *Adversary's goal*: The attacker aims to impersonate a targeted client (integrity, targeted attack).
- *Adversary's knowledge*: She is assumed to know 1) the feature extraction algorithm, 2) the matching algorithm s, 3) the fusion rule s_c, and 4) a set of n face images $\{\widehat{x}_j\}_{j=1}^{n}$ of the targeted client, different from those in the client's template gallery (e.g., potentially collected from a social network).
- *Adversary's capability*: She can only submit printed photos of faces to the sensor during verification.
- *Adversary's strategy*: Under these assumptions, the attacker can approximate the score $s_c(x)$ computed by the targeted system for the claimed identity c using the collected face images of the victim, i.e., she can compute an estimate $\hat{s}_c(x) = (1/n)\sum_{j=1}^{n} s(x, \widehat{x}_j)$. Accordingly, the optimal attack strategy is given by:

$$x^* = \arg\max_x \widehat{s}_c(x) = \frac{1}{n}\sum_{j=1}^{n} s(x, \widehat{x}_j), \tag{2}$$

where $\widehat{s}_c$ is the attacker's goal function $g(a, \theta)$ (1), and x^* is the attack sample that maximizes the objective in the PCA-induced feature space. The above problem can be solved by a simple gradient-ascent algorithm, and the resulting attack sample x^* can then be projected back onto the space of face images $\mathcal{Z}$ (where each feature corresponds to the gray-level value of a pixel) by inverting the PCA-induced mapping. This is also possible if more sophisticated matching algorithms and feature representations are used, using well-crafted heuristics. Please refer to [21] for further details.

An example of this improved spoofing technique on a simple case involving $n = 3$ templates is shown in Figure 3, where we also report the values of s_c for each of the client's face images $\{\widehat{x}_j\}_{j=1}^{n}$. It can be seen that the final attack face x^* yields a higher probability (i.e., s_c value) of successfully impersonating the victim than any of the available images $\{\widehat{x}_j\}_{j=1}^{3}$.

POISONING BIOMETRIC SYSTEMS THAT LEARN FROM EXAMPLES

We now report a different application example focusing on a poisoning attack against an adaptive face recognition system, and on the development of a corresponding defense. In recent work [20], [21], we have shown that an attacker may exploit the system's adaptation mechanism to compromise the templates of a given client by presenting a well-crafted sequence of fake faces to the camera, with the goal of denying him access to the system. At the same time, if the attacker replaces the targeted client's templates with her templates, she may also impersonate the client without presenting any fake traits to the sensor.

The face verification system considered in this example, as in [20] and [21], authenticates clients based on matching the acquired face image with a stored average template, referred to as *centroid*. For each client, the centroid is updated using the self-update algorithm: if the submitted face image is similar enough to the centroid, the

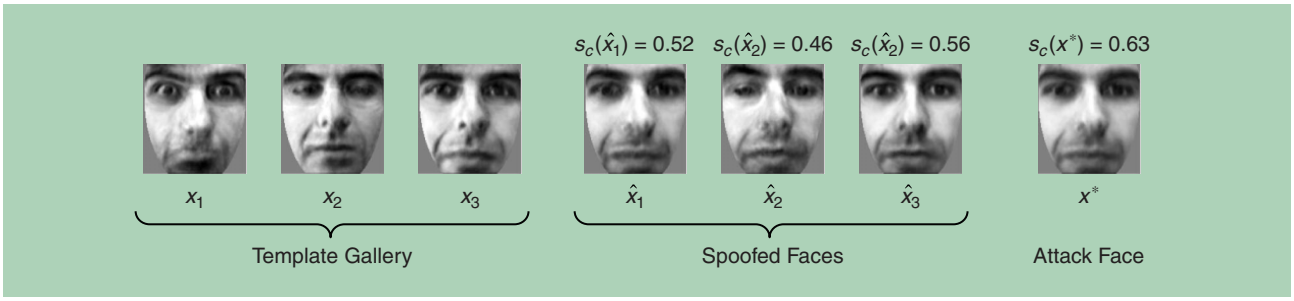

[FIG3] Face spoofing from multiple images. The client's templates $\{x_i\}_{i=1}^{3}$, the spoofed faces $\{\widehat{x}_j\}_{j=1}^{3}$, and the final attack face x^* [obtained from solving (2)] are shown, along with the corresponding s_c values.

latter is updated incorporating the new image into the computation of the average face image. As in the previous case, we consider a PCA-based mapping to map face images from $\mathcal{Z}$ onto a reduced vector space $\mathcal{X}$. The matching score for client c is computed here as $s_c(\boldsymbol{x}) = \exp\{-\|\boldsymbol{x}-\boldsymbol{x}_c\|^2\}$, where $\boldsymbol{x}_c$ is the client's centroid. The centroid $\boldsymbol{x}_c$ is initially computed as the average of n templates and, when $s_c(\boldsymbol{x}) \geq \theta_c$, updated as $\boldsymbol{x}_c' = \boldsymbol{x}_c + (1/n)(\boldsymbol{x}-\boldsymbol{x}_c)$, i.e., slightly drifted toward $\boldsymbol{x}$. Accordingly, in the PCA-based feature space, $\boldsymbol{x}_c$ is updated if the acquired image $\boldsymbol{x}$ is within a hypersphere centered on $\boldsymbol{x}_c$, with radius d_c dependent on the update threshold θ_c. The complete attack scenario is given next.

- *Adversary's goal*: It is that of replacing the centroid $\boldsymbol{x}_c$ of a given client with an attacker's template $\boldsymbol{x}_a$, both to deny access to client c, and to allow the attacker to impersonate c using her own face (i.e., a targeted attack violating both system availability and integrity).
- *Adversary's knowledge*: The attacker is assumed to know 1) the feature extraction algorithm; 2) the matching algorithm; 3) the template update algorithm; and 4) the decision and self-update threshold. In the case of perfect knowledge, she also knows the centroid $\boldsymbol{x}_c$ of the targeted client c, while when limited knowledge is considered, only a good estimate of $\boldsymbol{x}_c$ is available to the attacker, e.g., a frontal face image of the victim collected from a social network.
- *Adversary's capability*: The attacker can modify the template database by presenting fake faces at the sensor that enable template self-update. The attack influence is thus over the enrollment/update phase.
- *Attack strategy*: Under these assumptions, the shortest sequence of fake traits required to replace the victim's centroid $\boldsymbol{x}_c$ with that of the attacker $\boldsymbol{x}_a$ can be found by solving the following optimization problem, for each sample $\boldsymbol{x}$ in the attack sequence: $\min_{\boldsymbol{x}} \|\boldsymbol{x}-\boldsymbol{x}_a\|^2$, subject to the update condition $\|\boldsymbol{x}-\boldsymbol{x}_c\| \leq d_c$ [see Figure 4(a)]. At each iteration, this amounts to finding the closest attack sample to $\boldsymbol{x}_a$ that enables an update of $\boldsymbol{x}_c$. The solution is simply given as $\boldsymbol{x} = \boldsymbol{x}_c + d_c\vec{a}$, where $\vec{a} = (\boldsymbol{x}_a - \boldsymbol{x}_c/\|\boldsymbol{x}_a - \boldsymbol{x}_c\|)$ is the so-called attack direction. In practice, each attack sample $\boldsymbol{x}$ is found at the intersection between the hypersphere corresponding to the update condition, and the line connecting $\boldsymbol{x}_a$ and $\boldsymbol{x}_c$. As in the previous example, the face images for the attack sequence can be obtained by projecting the attack samples from the PCA-induced space $\mathcal{X}$ onto the space of face images $\mathcal{Z}$. Then, the attacker can fabricate the corresponding fake faces (e.g., by printing them on paper), and present them in the right order to the sensor. An example of how the victim's centroid is gradually updated by the corresponding sequence of attack faces, under the assumption of perfect

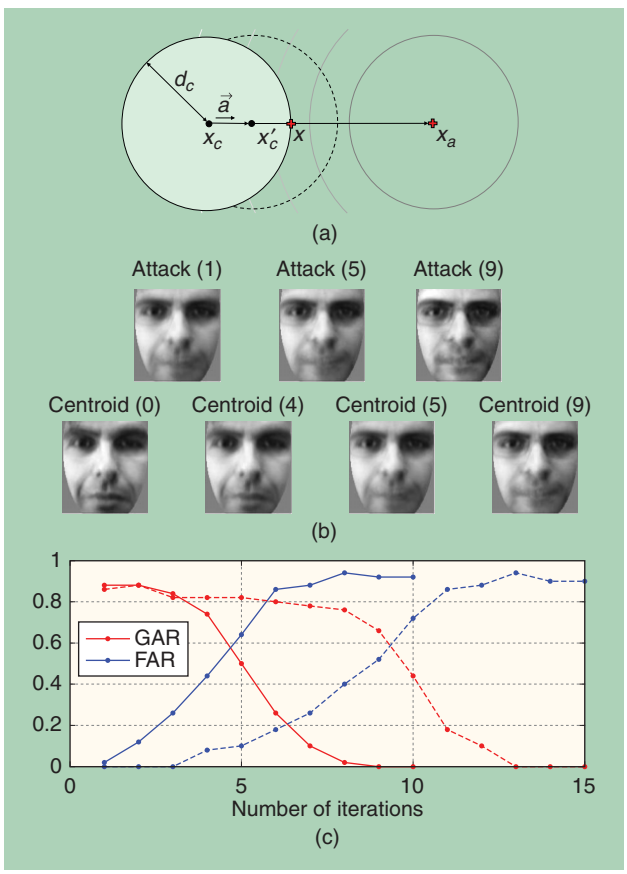

[FIG4] (a) A poisoning attack with perfect knowledge. The circles centered on $\mathbf{x}_a$ represent the objective function $\|\mathbf{x}-\mathbf{x}_a\|$, minimized by the attack point $\mathbf{x}$ on the feasible domain $\|\boldsymbol{x}-\boldsymbol{x}_c\| \leq d_c$. The updated centroid $\mathbf{x}_c'$ and the feasible domain for the next attack iteration are also shown. (b) Attack samples and the victim's centroids for poisoning with perfect knowledge at different iterations. (c) Genuine acceptance rate (GAR) and false acceptance rate (FAR) for poisoning with perfect (solid lines) and limited (dashed lines) knowledge at different iterations.

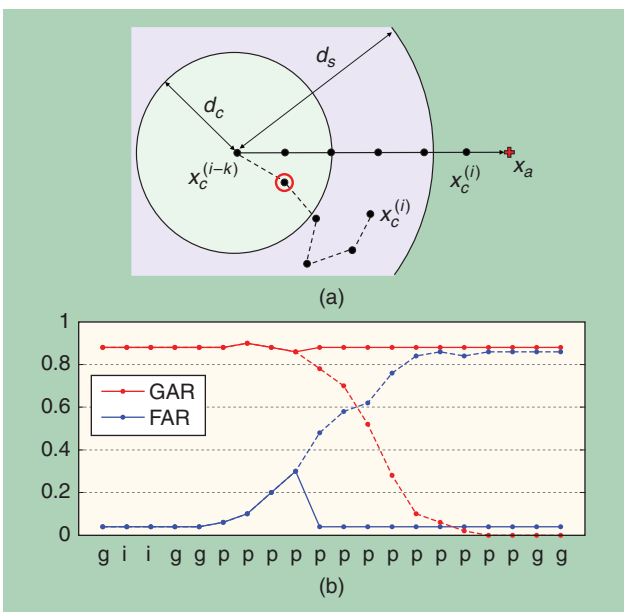

[FIG5] (a) Template sanitization. If the current centroid $x_c^{(i)}$ falls outside the sanitization hypersphere (dark gray area), as for the poisoning attack sequence (solid line), the centroid $x_c^{(i-k)}$ is restored; otherwise, as for the hypothesized genuine update sequence (dashed line), the center of the sanitization hypersphere is updated to $x_c^{(i-k+1)}$ (red circled point). (b) GAR and FAR values in the presence (solid lines) and in the absence (dashed lines) of template sanitization, after different centroid updates, including genuine ("g") and impostor ("i") attempts, and poisoning attacks ("p") with perfect knowledge.

knowledge, and for $n = 5$, is given in Figure 4(b). In (c), we show how the GAR and FAR vary as the attack proceeds, under perfect and limited knowledge of the victim's template. Note how the probability of authenticating the attacker (without presenting any fake face) as the victim (i.e., the FAR) increases, while the probability of correctly authenticating the victim as a genuine user (i.e., the GAR) decreases, since the victim's template is gradually morphed toward the attacker's face while the attack is in progress. Results for the perfect and limited knowledge attacks are similar, despite the latter case requires more iterations (i.e., submitting more fake faces) to compensate the lack of knowledge of the victim's template. The exact number of iterations required to complete both attacks can also be analytically computed [21].

- *Template sanitization*: Here we present a novel countermeasure based on the idea of sanitizing the template gallery (i.e., identifying anomalous template updates), inspired by the countermeasures proposed in adversarial machine learning against poisoning attacks [1]–[4], [28].

The underlying idea is to analyze whether the sequence of the most recent k updated centroids $\boldsymbol{x}_c^{(i-k)}, \ldots, \boldsymbol{x}_c^{(i)}$ (where i denotes the current iteration) falls within a given region of the feature space, called sanitization hypersphere [see Figure 5(a)]. If the current centroid $\boldsymbol{x}_c^{(i)}$ falls within the sanitization hypersphere, i.e., if $\| \boldsymbol{x}_c^{(i)} - \boldsymbol{x}_c^{(i-k)} \| \leq d_s$, then the center of the sanitization hypersphere is updated to the next centroid in the sequence, i.e., $\boldsymbol{x}_c^{(i-k+1)}$; otherwise, the current center of the sanitization hypersphere $\boldsymbol{x}_c^{(i-k)}$ is restored as the current centroid. In this case, an alert may be also (or alternatively) raised to the system administrator to report the anomalous update. The rationale of this approach is to identity sequences of centroid updates that consistently drift the centroid toward a given, biased direction, within a small number of iterations (as it happens in the presence of a poisoning attack), assuming that genuine updates exhibit a different (i.e., less biased and more random) behavior. The parameter k and the hypersphere radius d_s of the proposed approach should thus be chosen such that $d_s \leq k\,(d_c/n)$, otherwise poisoning attacks will not be detected, as they drift the centroid of an amount equal to (d_c/n) at each iteration.

In Figure 5(b), we report an example using the same attacker and victim pair considered in the previous case. Initially, we simulate a number of random accesses to the system, including genuine and impostor attempts, which do not significantly affect GAR and FAR. Then, the attacker launches a poisoning attack, and, as in the previous case, the GAR decreases and the FAR increases. If template sanitization is not implemented, the attack succeeds, and system integrity and availability are compromised. Conversely, in the presence of template sanitization, the attack is detected after four iterations, and a previous centroid is restored. This avoids the attack to succeed, preserving normal system operation and security. Although this simple countermeasure can be misled by a poisoning attack in which the attack samples are closer to the current centroid (instead of lying exactly at the boundary of the feasible domain), this would require the attacker to perform a significantly higher number of iterations to complete the attack. We can thus conclude that the proposed sanitization technique helps improving system security.

SUMMARY AND OPEN PROBLEMS

In the last decade, the use of electronic devices in our daily lives has become increasingly pervasive, providing several advantages in managing our tasks and communicating with other people. However, with such a huge number of powerful devices connected to the so-called Internet of Things (e.g., smartphones, smart TVs, etc.), the number of potential attack points and vulnerabilities has significantly increased, as well as chances for attackers to compromise the corresponding devices and systems. In addition, the level of sophistication of attacks has become increasingly higher over the years, witnessing the presence of very skilled attackers and strong economic incentives behind these activities. Biometrics can be considered a potential tool for improving the security of such systems in the digital era. However, besides having a strong deterrent effect, they should really be designed to be intrinsically secure, to successfully resist the very sophisticated attacks that may be incurred during operation.

In this article, we have provided an overview of the current state of the art on biometric system security from a perspective inspired by the field of adversarial machine learning. We have discussed how this novel perspective may not only inspire the simulation of more sophisticated attack scenarios, but also how, based on such scenarios, more effective countermeasures can be proactively developed. As concrete application examples, we have considered a sophisticated spoofing attack and a poisoning attack against an adaptive face verification system in which the attacker gradually compromises the template gallery of a given client by presenting a well-crafted sequence of fake faces. We have also proposed a novel countermeasure based on template sanitization. Another example of how the proposed perspective based on adversarial machine learning may depict novel attack scenarios and inspire potential countermeasures is related to recent work in adversarial learning, which has shown that clustering algorithms may be significantly vulnerable to well-planned attacks [30]. In the biometric setting, similar attacks may target systems that perform template selection and update exploiting clustering algorithms. Although investigating and developing secure clustering algorithms against adaptive and intelligent attackers is still an open issue in the adversarial machine-learning field, the corresponding results may also inspire countermeasures that can be adapted to improve the security of template selection and update procedures in biometric systems.

ACKNOWLEDGMENTS

This work has been partly supported by the TABULA RASA project, 7th Framework Research Programme of the European Union, grant agreement number 257289; and by the project CRP-18293, funded by Regione Autonoma della Sardegna, L.R. 7/2007, Bando 2009.

AUTHORS

Battista Biggio (battista.biggio@diee.unica.it) received the M.S. degree in electronic engineering (with honors) and the Ph.D. degree in electronic engineering and computer science, respectively, in 2006 and 2010, from the University of Cagliari, Italy,

where he has been with the Department of Electrical and Electronic Engineering of the same university since 2007 and now holds a postdoctoral position. His research interests currently include adversarial learning, multiple classifier systems, kernel methods, biometric authentication, spam, and malware detection. He serves as a reviewer for the main international conferences and journals in these fields. He is a member of the IEEE.

Giorgio Fumera (fumera@diee.unica.it) received the M.S. degree in electronic engineering (with honors) and the Ph.D. degree in electronic engineering and computer science, respectively in 1997 and 2002, from the University of Cagliari, Italy. Since February 2010, he has been an associate professor of computer engineering in the Department of Electrical and Electronic Engineering at the same university. His research interests are related to methodologies and applications of statistical pattern recognition, and include multiple classifier systems, classification with the reject option, adversarial classification, and document categorization. He acts as reviewer for the main international conferences and journals in these fields. He is a Member of the IEEE.

Paolo Russu (paolo.russu@diee.unica.it) received the B.S. degree and M.S. degree (with honors) in 2010 and 2013, respectively, from the University of Cagliari, Italy. Since May 2014 he has been a Ph.D. student at the same university. His research focuses on adversarial machine learning, biometrics, and computer security.

Luca Didaci (luca.didaci@diee.unica.it) received the M.S. degree in electronic engineering and the Ph.D. degree in electronic engineering and computer science, respectively, in 2001 and 2005, from the University of Cagliari, Italy. Since 2002 he has been an assistant professor of computer engineering in the Department of Electrical and Electronic Engineering at the same university. His research interests are in the fields of pattern recognition and applications. In particular, they rely on semisupervised multiple classifier systems, semisupervised biometric systems, and adversarial classification. He acts as reviewer for the main international conferences and journals in these fields. He is a Member of the IEEE.

Fabio Roli (roli@diee.unica.it) received his M.S. degree (with honors) and Ph.D. degree in electronic engineering from the University of Genoa, Italy. He was a member of the research group on Image Processing and Understanding of the University of Genoa, Italy, from 1988 to 1994. He was adjunct professor at the University of Trento, Italy, in 1993 and 1994. In 1995, he joined the Department of Electrical and Electronic Engineering of the University of Cagliari, Italy, where he is now professor of computer engineering and head of the research group on pattern recognition and applications. His research over the past 20 years addressed the design of pattern recognition systems in real applications. He is Fellow of the IEEE and of the International Association for Pattern Recognition.

REFERENCES

[1] M. Barreno, B. Nelson, R. Sears, A. D. Joseph, and J. D. Tygar, "Can machine learning be secure?" in *Proc. ACM Symp. Information, Computer and Communications Security (ASIACCS'06)*. New York: ACM, 2006, pp. 16–25.

[2] L. Huang, A. D. Joseph, B. Nelson, B. Rubinstein, and J. D. Tygar, "Adversarial machine learning," in *Proc. 4th ACM Workshop Artificial Intelligence and Security* Chicago, IL, 2011, pp. 43–57.

[3] B. Biggio, G. Fumera, and F. Roli, "Security evaluation of pattern classifiers under attack," *IEEE Trans. Knowledge Data Eng.*, vol. 26, no. 4, pp. 984–996, 2014.

[4] B. Biggio, G. Fumera, and F. Roli, "Pattern recognition systems under attack: Design issues and research challenges," *Int. J. Pattern Recogn. Artif. Intell.*, vol. 28, no. 7, p. 1460002, 2014.

[5] P. Laskov and R. Lippmann, Eds., *NIPS Workshop on Machine Learning in Advances in Environments for Computer Security*, 2007.

[6] A. D. Joseph, P. Laskov, F. Roli, and D. Tygar, Eds., *Dagstuhl Perspectives Workshop on Machine Learning Methods for Computer Security*, 2012.

[7] N. K. Ratha, J. H. Connell, and R. M. Bolle, "An analysis of minutiae matching strength," in *AVBPA*, ser. LNCS, J. Bigün and F. Smeraldi, Eds., vol. 2091. New York: Springer, 2001, pp. 223–228.

[8] A. K. Jain, K. Nandakumar, and A. Nagar, "Biometric template security," *J. Adv. Signal Process.*, vol. 2008, New York: Hindawi, pp. 1–17, 2008.

[9] R. N. Rodrigues, L. L. Ling, and V. Govindaraju, "Robustness of multimodal biometric fusion methods against spoof attacks," *J. Vis. Lang. Comput.*, vol. 20, no. 3, pp. 169–179, 2009.

[10] A. Adler, "Vulnerabilities in biometric encryption systems," in *Proc. 5th Int. Conf. Audio- and Video-Based Biometric Person Authentication*, ser. LNCS, T. Kanad, A. Jain, N. K. Ratha, Eds., vol. 3546. Rye, NY: Springer, 2005, pp. 1100–1109.

[11] J. Galbally, C. McCool, J. Fierrez, S. Marcel, and J. Ortega-Garcia, "On the vulnerability of face verification systems to hill-climbing attacks," *Pattern Recogn.*, vol. 43, no. 3, pp. 1027–1038, 2010.

[12] M. Martinez-Diaz, J. Fierrez, J. Galbally, and J. Ortega-Garcia, "An evaluation of indirect attacks and countermeasures in fingerprint verification systems," *Pattern Rec. Lett.*, vol. 32, no. 12, pp. 1643–1651, 2011.

[13] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, and F. Roli, "Evasion attacks against machine learning at test time," in *European Conf. Machine Learning and Principles and Practice of Knowledge Discovery in Databases, Part III*, ser. LNCS, H. Blockeel, K. Kersting, S. Nijssen, and F. Železný Eds., vol. 8190. Berlin: Springer, 2013, pp. 387–402.

[14] A. Globerson and S. T. Roweis, "Nightmare at test time: Robust learning by feature deletion," in *Proc. 23rd Int. Conf. Machine Learning*, W. W. Cohen and A. Moore, Eds., vol. 148. New York: ACM, 2006, pp. 353–360.

[15] B. I. Rubinstein, B. Nelson, L. Huang, A. D. Joseph, S.-h. Lau, S. Rao, N. Taft, and J. D. Tygar, "Antidote: Understanding and defending against poisoning of anomaly detectors," in *Proc. 9th ACM SIGCOMM Internet Measurement Conf., IMC'09*. New York: ACM, 2009, pp. 1–14.

[16] D. Lowd and C. Meek, "Adversarial learning," in *Proc. 11th ACM SIGKDD Int. Conf. Knowledge Discovery Data Mining*. Chicago, IL: ACM, 2005, pp. 641–647.

[17] B. Nelson, B. I. Rubinstein, L. Huang, A. D. Joseph, S. J. Lee, S. Rao, and J. D. Tygar, "Query strategies for evading convex-inducing classifiers," *J. Mach. Learn. Res.*, vol. 13, pp. 1293–1332, May 2012.

[18] U. Uludag, A. Ross, and A. K. Jain, "Biometric template selection and update: A case study in fingerprints," *Pattern Recogn.*, vol. 37, no. 7, pp. 1533–1542, 2004.

[19] C. Ryu, H. Kim, and A. K. Jain, "Template adaptation based fingerprint verification," in *Proc. 18th Int. Conf. Pattern Recognition, ICPR'06*, vol. 4, Washington, DC, 2006, pp. 582–585.

[20] B. Biggio, G. Fumera, F. Roli, and L. Didaci, "Poisoning adaptive biometric systems," in *Structural, Syntactic, and Statistical Pattern Recognition*, LNCS, G. Gimel'farb, E. Hancock, A. Imiya, A. Kuijper, M. Kudo, S. Omachi, T. Windeatt, and K. Yamada, Eds. New York: Springer, 2012, vol. 7626, pp. 417–425.

[21] B. Biggio, L. Didaci, G. Fumera, and F. Roli, "Poisoning attacks to compromise face templates," in *Proc. 6th IAPR Int. Conf. Biometrics*, Madrid, Spain, 2013, pp. 1–7.

[22] B. Biggio, Z. Akhtar, G. Fumera, G. L. Marcialis, and F. Roli, "Security evaluation of biometric authentication systems under real spoofing attacks," *IET Biomet.*, vol. 1, no. 1, pp. 11–24, 2012.

[23] T. Matsumoto, H. Matsumoto, K. Yamada, and S. Hoshino, "Impact of artificial "gummy" fingers on fingerprint systems," *Datenschutz und Datensicherheit*, vol. 26, no. 8, pp. 1–15, 2002.

[24] P. Johnson, B. Tan, and S. Schuckers, "Multimodal fusion vulnerability to non-zero effort (spoof) imposters," in *Proc. IEEE Int. Workshop on Information Forensics and Security*, 2010, pp. 1–5.

[25] N. Dalvi, P. Domingos, Mausam, S. Sanghai, and D. Verma, "Adversarial classification," in *Proc. 10th ACM SIGKDD Int. Conf. Knowledge Discovery Data Mining*, Seattle, 2004, pp. 99–108.

[26] M. Brückner, C. Kanzow, and T. Scheffer, "Static prediction games for adversarial learning problems," *J. Mach. Learn. Res.*, vol. 13, pp. 2617–2654, Sept. 2012.

[27] C. H. Teo, A. Globerson, S. Roweis, and A. Smola, "Convex learning with invariances," in *Proc. Neural Information Processing Systems 20*, J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, Eds. Cambridge, MA: MIT Press, 2008, pp. 1489–1496.

[28] G. F. Cretu, A. Stavrou, M. E. Locasto, S. J. Stolfo, and A. D. Keromytis, "Casting out demons: Sanitizing training data for anomaly sensors," in *Proc. IEEE Symp. Security and Privacy*. Los Alamitos, CA: 2008, pp. 81–95.

[29] A. D. Joseph, P. Laskov, F. Roli, J. D. Tygar, and B. Nelson, "Machine learning methods for computer security (Dagstuhl Perspectives Workshop 12371)," *Dagstuhl Manifestos*, vol. 3, no. 1, pp. 1–30, 2013.

[30] B. Biggio, I. Pillai, S. R. Bulò, D. Ariu, M. Pelillo, and F. Roli, "Is data clustering in adversarial settings secure?" in *Proc. ACM Workshop on Artificial Intelligence Security (AISec'13)*. New York: ACM, 2013, pp. 87–98.

[SP]

Gene Itkis, Venkat Chandar, Benjamin Fuller, Joseph P. Campbell, and Robert K. Cunningham

Iris Biometric Security Challenges and Possible Solutions

For your eyes only—Using the iris as a key

©ISTOCKPHOTO.COM/GREYFEBRUARY

Biometrics Security and Privacy Protection

Biometrics were originally developed for identification, such as for criminal investigations. More recently, biometrics have been also utilized for authentication. Most biometric authentication systems today match a user's biometric reading against a stored reference template generated during enrollment. If the reading and the template are sufficiently close, the authentication is considered successful and the user is authorized to access protected resources. This binary matching approach has major inherent vulnerabilities.

An alternative approach to biometric authentication proposes to use fuzzy extractors (also known as *biometric cryptosystems*), which derive cryptographic keys from noisy sources, such as biometrics. In theory, this approach is much more robust and can enable cryptographic authorization. Unfortunately, for many biometrics that provide high-quality identification, fuzzy extractors provide no security guarantees.

Digital Object Identifier 10.1109/MSP.2015.2439717
Date of publication: 13 August 2015

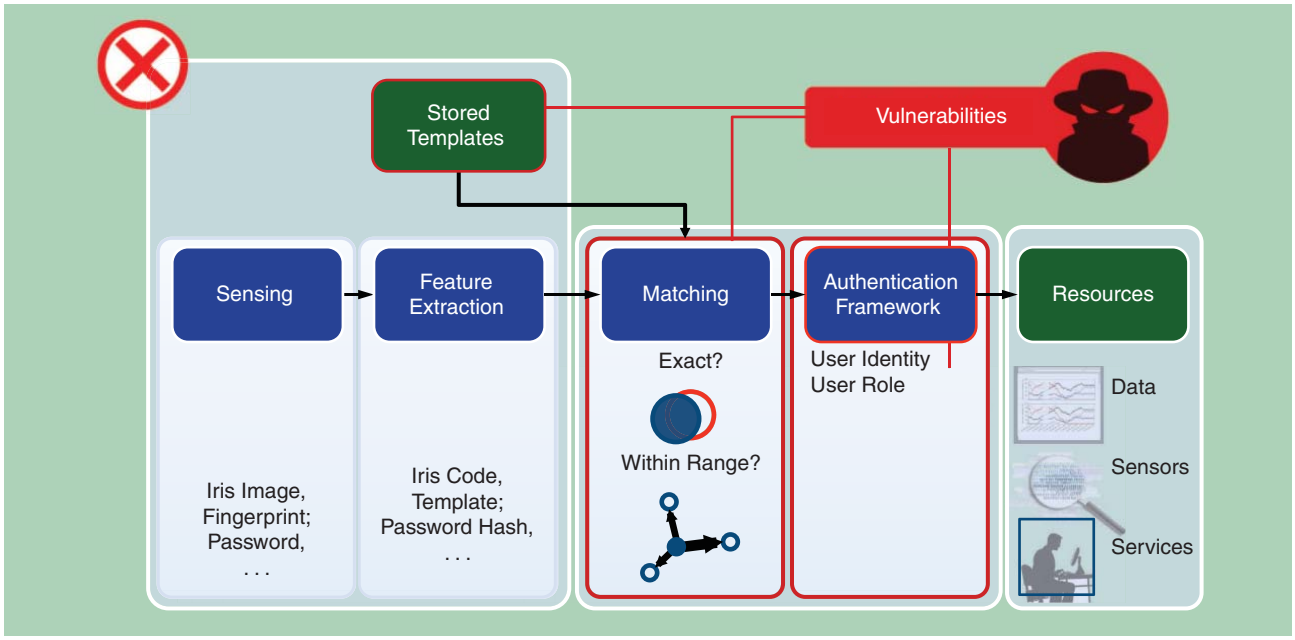

[FIG1] Binary authentication and authorization. First, authentication data, such as a password and/or biometric, is collected from the user and transformed into the appropriate canonical form, or template. Next, the acquired template is matched against the stored reference template. If the match is successful, the authorization framework grants the user access to the appropriate resources (e.g., via appropriate cryptographic keys). The three major vulnerable components of this approach are highlighted: matching (with its fragile binary decision), stored templates, and authorization framework (and its cryptographic keys).

This gap arises in part because of an objective mismatch. The quality of a biometric identification is typically measured using false match rate (FMR) versus false nonmatch rate (FNMR). As a result, biometrics have been extensively optimized for this metric. However, this metric says little about the suitability of a biometric for key derivation.

In this article, we illustrate a metric that can be used to optimize biometrics for authentication. Using iris biometrics as an example, we explore possible directions for improving processing and representation according to this metric. Finally, we discuss why strong biometric authentication remains a challenging problem and propose some possible future directions for addressing these challenges.

INTRODUCTION

AUTHENTICATION AND AUTHORIZATION

Security systems commonly include two components: 1) an authentication framework that validates a user's identity and 2) an authorization framework (sometimes called a *reference monitor*) that controls access to resources for the validated identity [1], [2]. Biometrics represent one important way to verify a user's identity [3]–[7]. An adversary able to impersonate you can do surprisingly bad things to you and in your name. Authentication is a crucial aspect of security, but it is surprisingly difficult to do in a way that is easy to deploy and use, as well as provide proper protection [8].

BINARY MATCHING AUTHENTICATION PARADIGM

In typical modern authentication systems, during the enrollment stage—when a user is added to the system—an original reading is collected from the user; transformed into a canonical form; and stored as a reference value, often referred to as a *stored template*.

Later, when the user authenticates to the system, a new reading is collected, transformed into the same canonical form, and matched against the stored reference template. The authentication is successful if the two values match. Then an authorization framework may grant the user access to protected resources, e.g., by providing cryptographic keys (called *content keys*) to the appropriate resources (see Figure 1).

For example, for passwords, a user enrolls by selecting a password, the canonical form is a cryptographic hash, and an authenticated match is required to be exact [9]. Biometrics are typically noisy: readings vary even for the same subject—hence, the match is approximate. We call this paradigm *binary matching* (sometimes using only one of these words) since the authentication results in just a single bit: match or not match.

INHERENT WEAKNESSES OF BINARY MATCHING PARADIGM

This paradigm has a number of crucial inherent weaknesses (see [10] for their manifestations in biometric systems):

- Binary authentication decisions in the matching step are fragile and can be skipped or flipped even by accidental errors.
- Matching requires the reference templates to be readily available during authentication, creating opportunities for an attacker to steal the templates, which in turn enables further attacks (see the section "Weaknesses of Binary Matching Paradigm: Details").

The binary nature of the decision also implies that the system must have access to all the resources that a properly authenticated

FUZZY EXTRACTORS

Fuzzy extractors are a pair of algorithms for deriving keys from a noisy source of entropy. The first algorithm, *generate* or Gen, is used at enrollment time. It takes an initial reading w, producing a key as well as public information P. The second algorithm, *reproduce* or Rep, is used at authentication time, taking w' (a nearby reading of an iris) and the public value P. The correctness guarantee is that Gen and Rep should give the same key if the distance between w and w' is at most some bounded parameter denoted t. To protect against the attacks described in the Introduction, the key should be strong even in the presence of P. The problem is trivial if P is private. A private P can store a key and the original reading. Then Rep outputs the key if and only if the new reading w' is close enough. This essentially reduces the problem to having a good biometric source.

Bennett, Brassard, and Robert identified two crucial tasks for deriving keys from noisy data [47]. The first, information-reconciliation, removes errors from w'. The second, privacy amplification, converts w to a uniform value. Traditionally, a fuzzy extractor uses two separate algorithms to accomplish these tasks. A secure sketch [48] performs information-reconciliation and a randomness extractor [49] performs information-reconciliation. A fuzzy extractor that separates information-reconciliation and privacy amplication is called the *sketch-and-extract construction*. See the work of Dodis et al. for formal definitions of the requirements of fuzzy extractors and secure sketches [48, Sec. 2.5–4.1]. Here we provide a brief review of standard constructions and recent advances. The goal of secure sketch is to map nearby w' back to the original w without revealing unnecessary information about w.

The simplest construction of a secure sketch uses the syndrome of an error correcting code. The public information P consists of applying a parity check matrix to the original reading w. This allows decoding of the original w from a nearby w' and P. The entropy of w conditioned on this secure sketch is at least the starting entropy minus the length of the syndrome. The length of a syndrome must increase as the error tolerance increases. This means the lower bound on the remaining entropy of w decreases as the error tolerance increases.

There are many coding-based constructions of secure sketches. The security analysis of these constructions usually considers the difference between the size of the metric space and the deficiency of the best code correcting enough errors in the metric space. (The measure can be relaxed considerably by allowing the secure sketch to occasionally output the wrong value.) This imposes a tradeoff between the remaining entropy of w and the noise that can be tolerated.

Correcting more errors decreases FNMR and decreases the length of the derived key. Standard constructions of fuzzy extractors work well when the source has a high entropy rate (nearly uniform). Recent work builds fuzzy extractors from the face biometric when the entropy rate is almost full, even for an error rate of nearly 30% [50]. However, standard fuzzy extractors are not known to be secure on sources with low entropy rates. This is not a limitation of a particular construction; there are probability distributions with the same entropy and error rate as irises where key derivation is impossible (see [40] and [48, Appendix C]).

Recent works have also built fuzzy extractors using properties of a distribution other than entropy and desired error tolerance [40], [51], [52]. Unfortunately, these constructions are not known to work for the iris distribution. These constructions assume properties of the physical source that irises do not appear to satisfy. Thus, authentication from the iris remains challenging.

user might need, thus forcing a violation of the principle of least privilege [11]. As a result, by compromising the system, an attacker can also obtain this access; this is what, in particular, makes privilege escalation attacks so attractive. These inherent weaknesses become especially apparent in settings where it is not clear who can be trusted to perform the matching and grant access, e.g., in the cloud.

THE BIOMETRIC CRYPTOSYSTEM PARADIGM

An alternative approach is to derive cryptographic keys from biometrics. These keys can then be used to access the resources. The challenge here is that the biometrics are inherently noisy. This approach originated with the work on fuzzy commitment and fuzzy vault [12]–[14], building on ideas from [15] and [16]. These ideas were formalized in a cryptographic object known as fuzzy extractors that reliably produce a uniform key from a noisy source of entropy (see "Fuzzy Extractors"). In particular, this assures that even if the entropy was distributed unevenly among the bits of the source, the output will have all bits random. This approach has also been investigated under the names of *biometric cryptosystems* and *biometric key generation* [17], [18]. We show the main stages in this approach in Figure 2.

We have implemented a full authentication system where authorization is implicit based on the knowledge of the proper cryptographic keys. Compared to the single bit produced by binary matching, this approach enables leveraging the full entropy (to the extent possible) that the user provides as part of the authentication. We call such an approach *cryptographic authorization* or *cryptographic access control*. The advantages of this approach according to the metrics of [8] are discussed in "Benefits of the General Full-Entropy Approach." During the development of this system, we found that fuzzy extractors often do not provide meaningful guarantees on the key strength (KS) of keys derived from biometrics. For example, the iris code [3], which is a representation of one of the best biometrics [19], produces a key with no provable security using standard fuzzy extractors [20, Sec. 5]. The goal of this work is to examine why this is the case and how to derive stronger keys from biometrics.

THE GAP BETWEEN BIOMETRICS AND FUZZY EXTRACTORS

Extensive work on fuzzy extractors and similar techniques may lead one to believe that deriving cryptographic keys from

biometrics is a solved problem. Unfortunately, authentication from many noisy sources remains challenging.

Biometric techniques have been developed and optimized for identification (these optimizations naturally carry over to binary matching). The metrics used for evaluating biometrics are typically variants of FMR versus false FNMR plots. But these characteristics say little about the cryptographic security of the keys derived from the biometric. A different metric must be identified for the cryptographic authentication task, and biometric techniques need to be optimized according to this metric.

In this work, we propose that biometric quality for authentication should be measured as FNMR versus the strength of the key (see the section "Metric for Cryptographic Authentication"). This will help focus the development of biometrics for authentication as their quality will be measured for that task.

Using the iris as an example, we provide initial optimizations of biometrics to this authentication metric. However, key derivation from the iris remains a challenging problem.

For the purposes of simplicity, we assume the goal is to derive a 128-bit key. This is the key length for the standard symmetric cipher AES-128 approved by the National Security Agency to protect information classified up to secret level, so this should be useful for the most commonly used systems and accounts [21].

THE CASE FOR KEY DERIVATION FROM BIOMETRICS

In this article, we consider the suitability of biometrics for strong authentication. We discuss authentication modalities in "Other Authentication Modalities."

The central problem that complicates the use of biometrics is noise: readings of the same biometric of the same user can differ significantly, even under the most controlled (and thus least flexible and convenient for the user) environments. For binary matching, this forces use of approximate comparisons (using an appropriate metric). When biometrics are used for cryptographic authorization, noise represents a larger challenge.

WEAKNESSES OF BINARY MATCHING PARADIGM: DETAILS

As discussed in the section "Inherent Weaknesses of Binary Matching Paradigm," attackers can exploit fragility of the matching step of the binary matching paradigm, obtaining access to a single session. Similarly, privilege escalation attacks are extremely powerful.

In the case of biometrics, the transformation must preserve locality (i.e., similar readings, such as those taken from the same user, must remain close according to an appropriate metric, even after the transformation into the canonical form). This makes it difficult to design truly one-way transformations. Indeed, for commonly used transformations, it turns out that the stored reference templates can be reverse-engineered to produce realistic biometrics that would match the corresponding templates. Adversaries can manufacture a natural looking iris biometric that will pass the identification test [22], [23]; the same is true for fingerprints [24].

CRYPTOMATCHING: GOING HALFWAY

It is, in principle, possible to use cryptographic key derivation to obtain a cryptographic key, but then revert back to a binary match paradigm—comparing the key to the stored reference. Since, in this case, a cryptographic key is matched against a stored reference, we call this approach a *cryptomatch* authentication.

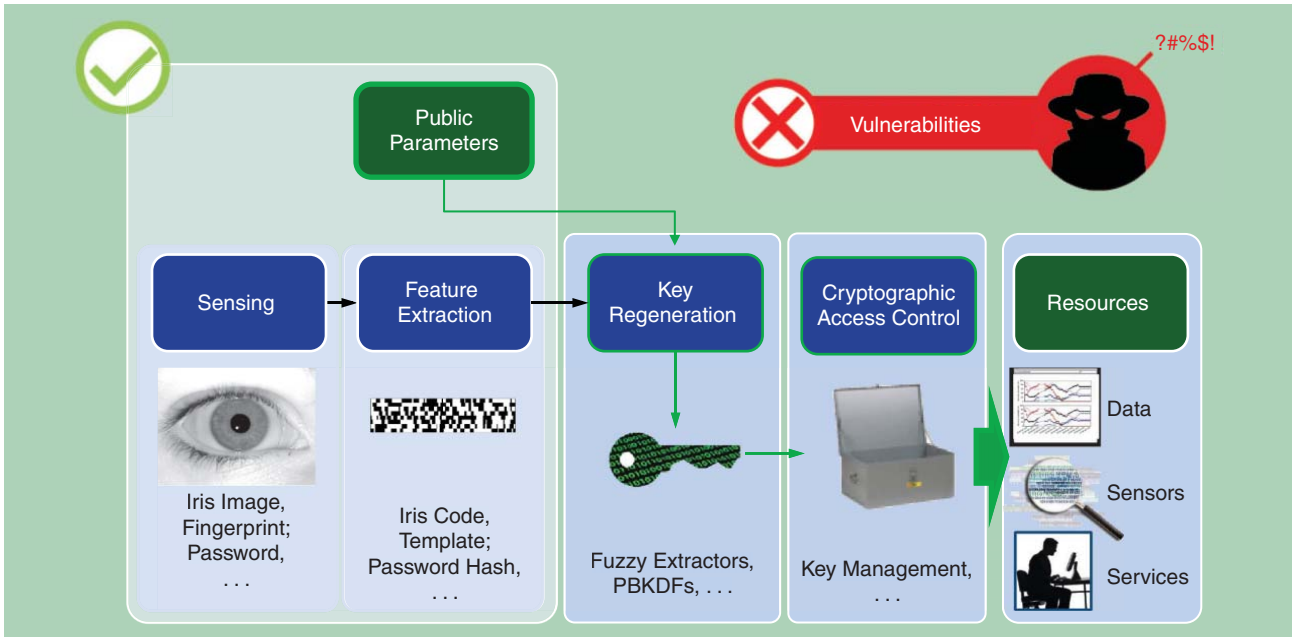

[FIG2] Cryptographic authentication and authorization: Authentication data is collected from the user and transformed into the appropriate canonical form or template. But instead of matching this acquired template against the stored reference template, the acquired template is used (with some nonsecret public parameters) to regenerate the cryptographic key. This key can then be used to obtain access to the appropriate resources via cryptographic access control and key management. The vulnerabilities highlighted in Figure 1 are no longer present.

BENEFITS OF THE GENERAL FULL-ENTROPY APPROACH

In a comprehensive study of two decades of general-purpose user authentication on the web, [8] proposes a broad set of usability, deployability, and security properties (constituting benefits, when satisfied), and uses these to compare different authentication techniques.

However, these properties had been formulated for the binary match approach. We believe that they should be revised in the context of the cryptographic authorization and cryptographic access control paradigms. Whenever it makes sense, we compare binary match to cryptomatch, rather than cryptographic authorization.

The change from binary match to cryptographic authorization has little or no effect on many properties listed in [8]: e.g., U1-U6 and D1-D2. For other properties, effect might be nonobvious. For example, Property U7: Infrequent-Errors aims to minimize the number of failed authentications. For the exact match (such as for passwords), this property would not be affected by the transition from binary matching to cryptomatching and cryptographic authorization. But for biometrics, the matching and key regeneration may involve somewhat different techniques, resulting in somewhat different error rates.

But furthermore, if proper cryptographic access control authorization is used—for any type of authentication modality, including passwords and biometrics—the failed authentication may result in a failure during access. Such access failures can be made easily detectable, eliminating the difference between binary match and cryptographic authorization approaches. But this failure detection can be exploited by an attacker as well—e.g., in an exhaustive search for the correct password. So, this represents a potential tradeoff between the usability and security (specifically, Property S4: Resilience-to-Unthrottled-Guessing). This example also illustrates the aforementioned need for revision of the properties in [8]: for example, inserting Property U7′: Easy-Recovery-from-Authentication-Failure.

Property U8: Easy-Recovery-from-Loss provides an example of a different tradeoff. A simple implementation, dogmatically following the least privilege principle, will result in a system where if a user loses the authentication data (whether public parameters or, say, the password), the data cannot be accessed by anyone at all and thus can be considered lost. However, a wiser implementation will build in various recovery mechanisms, which would trade off security and usability once again: making the data easier to recover but also easier to steal, or making the data more secure but also more difficult to recover. While passwords are typically considered to have the property U8 satisfied, this is typically because the choice in the security-usability trade off is made for the users, favoring usability at the cost of security (any administrator able to reset the password is also able to impersonate the user).

Property D3: Server-Compatibility requires the (authenticating) server to be compatible with the passwords. In our case, the servers only need to provide the right public information, and thus our approach completely eliminates the need for the server to do anything special for the authentication. D3 does not quite capture this benefit of cryptographic authorization. Cryptomatch, on the other hand, can be made server-compatible by using the key as the password.

The client in the cryptographic authorization and cryptomatch approaches needs to compute the cryptographic key. Currently,

For passwords, the traditional binary matching implementations can be seen as a specific implementation of exactly this cryptomatch principle. For biometrics, even this halfway approach can yield significant advantages compared to the binary matching the way it is typically practiced today.

As discussed previously, binary match biometrics suffer from an additional vulnerability: stealing a reference template enables an attacker to generate realistic biometrics. Some research on so-called cancelable biometrics and similar techniques helped to address this issue (e.g., see [18] for a survey). However, many of the considered methods do not result in strong security, say, compatible to provable security of common cryptographic tools (there are many notable exceptions, such as [25], [26], and others). In contrast, an ability of users to consistently regenerate cryptographic keys immediately results in a very strong version of cancelable biometrics. Once a key is derived, cryptomatch authentication can be easily implemented by deriving different (and easily replaceable) keys for each verifier, using the regenerated key as a master secret and unique random value (salt)—this is essentially similar to the way different keys are derived from a master secret in many secure protocols, e.g., such as transport-layer security [27].

BENEFITS OF CRYPTOGRAPHIC AUTHORIZATION

However, a cryptographic matching approach is still a binary decision. For high-security applications, authorization decisions should be based on knowledge of cryptographic keys derived from authentication material. Because there is no single-bit match/no match decision (one instead recovers a key), it is more difficult for an adversary to cause hardware to fail or software to branch in a way that will enable access.

Even in the most primitive case when the key regenerated in authentication stage is used as content key to access protected data directly, there is no information on the system that the attack can use to break security. The templates are replaced with public parameters and the content key is no longer stored, but rather recreated when needed from the biometric. This approach allows for strong authentication even against adversaries with physical access to the system when it is not actively used by a legitimate user. It is crucial for this approach that the public parameters produced by the fuzzy extractor do not compromise security of the biometric.

This approach also allows the creation of multiple keys from the same biometric in such a way that compromising one or a few of these keys (by leaking them to an adversary) in no way compromises the others. More sophisticated cryptographic access

our implementations run in Java, but not from a browser. It may be possible to implement similar tools as applets or run them from a browser in some other way, but until that happens, we need to consider our approach as having a negative impact on Property D4: Browser-Compatibility. While the change to cryptographic authorization has mixed effects on deployment and usability, it has positive impacts in the security area.

S1: Resilience-to-Physical-Observation, S5: Resilience-to-Internal-Observation, and S10: Requiring-Explicit-Consent are not directly affected by the shift from binary match to cryptographic authorization. These properties are typically provided only by special hard tokens (for which providing, or even better—encapsulating, so it can be used securely—a cryptographic key would typically be rather trivial). Similarly, S2: Resilience-to-Targeted-Impersonation (when personal knowledge of the target user person can be utilized in the attack) should not be affected by the shift, but it does eliminate some of the easiest pitfalls: while a user might use her mother's maiden name as a password, in cryptographic authorization, the user would not be prompted to use such information for security.

As discussed above, resilience to guessing (properties S3 and S4) can be strengthened at some usability cost. However in the throttled variant (S3), the verifier can limit the rate of attacker's guessing, but in our approach, there is no verifier. So, implementing throttling requires a different approach: e.g., a PBKDF2 function [53] can be used to increase the time to derive the key. Such an approach can also benefit even S4 (the unthrottled version, since the above mechanism requires no help from any servers). More sophisticated throttling mechanisms, however, can be designed for our approach, making it at least as resilient to guessing as the binary matching, and probably even more so.

Property S6: Resilience-to-Leaks-from-Other-Verifiers can be provided even by the cryptomatch version. For the cryptographic authorization, this property can be strengthened significantly (to be resilient) to leaks from any verifiers, since no verifiers are even present. The only other approach that assures this property is the zero-knowledge proofs of identity [54]. Such proofs are cryptographic protocols requiring significant computational power from both the prover and the verifier. Furthermore, just as in binary matching authentication, this protocol results in a single bit (pass or fail), and hence is vulnerable to the corresponding weaknesses.

Since cryptographic authorization requires no verifiers, it also has the other properties that rely on the verifiers not failing: S7: Resilient-to-Phishing and S11: Unlinkability. In fact, even the crypto-match approach can achieve unlinkability.

Cryptographic authorization also satisfies S9: No-Trusted-Third-Party. By necessity, the binary matching approach must exclude the verifier from the consideration. This is a direct consequence of the binary nature of the traditional approach. In contrast, cryptographic authorization has no verifiers and no (implicit or explicit) reference monitor acting upon the result of the binary authentication, it only relies on the client security not to steal the authentication data provided by the user (if hardware tokens or devices are an option, then even this vulnerability can be reduced or eliminated, depending on the specifics of the token). Cryptographic authorization improves security in a number of ways, with the improvements to S9 and S6 being the most significant.

control structures can be built as well, for example, traditional operating system permissions.

THE GAP BETWEEN BIOMETRICS AND FUZZY EXTRACTORS

Biometrics have been extensively used for the task of identification: discriminating effectively between two individuals. Rightfully, biometric systems were evaluated according to metrics for identification. The standard metric is a function between how reliably a single person's biometric can be recognized as such (FNMR) and how often two individuals are confused (FMR). Obviously, for a system always reporting a match, FNMR = 0 (since nonmatch, false or not, is never reported) and FMR = 1 (since different subjects are always falsely reported as matched). Conversely, never reporting a match makes FMR = 0 but FNMR = 1.

In practice, biometric systems allow adjusting their parameters to achieve some tradeoff between these characteristics depending on needs of specific applications. This tradeoff can be depicted as a function comparing FNMR versus FMR and is often used as a measure of quality of the biometric systems. We call this the *identification* metric. Current iris biometrics techniques produce a very strong biometric according to the identification metric [28].

There can be different ways to improve biometric systems according to this metric. For example, we can fuse iris codes from three readings by taking a majority for each bit in the iris code, as proposed in [15] and [29]. Then we use the result in matching, improving its quality according to the match metric (see Figure 3).

According to the identification metric, fusing provides an impressive improvement—see Figure 3, which depicts our experiments using the multispectral iris data set, licensed through the Scitor Corporation. But what this metric says is that we can get the FMR rate down to between 0.001–0.1%. In other words, by trying between 100,000 and 1,000 irises would give an attacker a good chance of impersonating a targeted user.

MEASURING BIOMETRIC SUITABILITY FOR CRYPTOGRAPHIC AUTHORIZATION

In this section, we show that the identification metric may be inappropriate for the cryptographic authorization task. Recall that the goal of this task is to create a strong cryptographic key from the same user, and note that the strength of the cryptographic key implies that different users rarely map to the same key.

Intuitively, FNMR corresponds to an authorized user failing an authentication attempt. Thus, FNMR can be viewed as,

OTHER AUTHENTICATION MODALITIES

It is common to organize all the authentication methodologies into three categories, according to the nature of the input provided by the user—we call them modalities: 1) what you have (e.g., hardware tokens, such as smart cards, etc.), 2) what you know (e.g., a password), and 3) what you are (biometrics).

Each of these categories comes with its own typical characteristics. Often multifactor authentication (using multiple modalities) is recommended to achieve higher security. But it is convenient to consider each modality, with its pros and cons, separately.

The strength of hardware tokens is that they can store plenty of entropy and perform complex computations far beyond human capacity. On the other hand, such tokens impose deployability limitations and can also be forgotten, lost or destroyed, or even stolen. Furthermore, in some cases (e.g., theft), the token might be used by an attacker. To defend against this, the user should authenticate to the token using another authentication modality [e.g., a short password, or personal identification number (PIN)]. Since token integrity can often be assumed, this authentication to the token is typically easier. Tokens are also subject to various hardware attacks and might also be used without explicit consent (e.g., if a PIN is cached after the first use of a smartcard—as done by many drivers—then malware can request smartcard to decrypt or authenticate data without the user knowing it).

Passwords—what you know—suffer from limitations of the human mind: since our memory is relatively weak, passwords have notoriously little entropy and are, hence, open to exhaustive search attack, such as password cracking. On the other hand, it is ultimately deployable, imposing the least amount of restrictions and only moderate inconvenience.

Finally, the biometric approach can be easier for the user, since there is nothing to remember or carry. Biometric reading can also be made very easy, requiring minimal effort from the user. But the easier the reading, the harder it is to enforce explicit consent. There may also be a separate tradeoff between the ease of reading and the reliability and the amount of entropy collected. This approach also requires some special equipment.

Biometric modality has another important feature: it is almost like a password that we wear literally on our face (or hands, in that case leaving its copies on everything we touch). In other words, biometrics can be easy to steal when the subject is present. So, the best use of biometrics is in remote authentication, where an attacker may not have physical access to the target user. Sometimes biometrics is used more as a test of physical presence, rather than the authentication. In that arena there is a constant arms race between biometric device manufacturing and the attackers, who use anything from jelly beans to cameras and special contact lenses. The most expert attackers tend to be leading in that race most of the time.

When a theft does occur, unlike passwords, biometrics cannot be easily replaced, no matter which authentication approach is used.

loosely speaking, a nuisance factor for authentication systems: e.g., it reflects how many attempts you would have to make before successfully logging into your own account. While this can have a serious impact on the system usability, it has no effect on security.

In contrast, FMR reflects probability that a user can be impersonated by someone else: e.g., if my iris is accepted when authenticating to log into your account. This is extremely important for security, and so FMR reflects insecurity of the system—the higher it is, the easier it is to trick the system into letting unauthorized users in.

FNMR versus FMR assumes an attacker that attempts authentication using random biometrics from a suitable population. Crucially, it says nothing about what is revealed by the authentication system. In a binary matching system, the authentication system writes down a template, revealing the original reading to the attacker. In a cryptographic authorization system, using fuzzy extractors, the system writes down some public information necessary to map nearby readings to the same cryptographic key. A dedicated attacker will use all information available; a metric for the authorization task must include this information.

As noted in [20, Sec. 5], there is no known fuzzy extractor for the iris. The core of the problem is that irises a relatively low entropy rate compared to their noise. For provable security, the known constructions of fuzzy constructors need the entropy rate to be significantly higher than the noise.

METRIC FOR CRYPTOGRAPHIC AUTHENTICATION

In an authorization system, the two important parameters are how often a user completes authentication (FNMR) and the strength of the resulting cryptographic key (in the presence of public parameters). Thus, when a biometric is used for authentication, the relevant metric is FNMR versus KS. It may be possible to generate a key whose strength and length are unequal. Since in our key derivations our goal is to produce a key indistinguishable from random, we can assume that key has strength proportional to $2^{-|\text{key}|}$.

As described in the section "Weaknesses of Binary Matching Paradigm: Details," any information written by the authentication system should be assumed to be available to the attacker. For this reason, we call any authentication information public parameters or P. Fuzzy extractors and other biometric cryptosystems write down error correcting information to ensure the same user reliably generates the same key. Therefore, the strength of a key should be measured relative to an adversary with access to public parameters P (such as the public information produced by fuzzy extractors in "Fuzzy Extractors").

We suggest measuring the quality of the biometric when used for cryptographic authorization as a function of FNMR and

the strength, KS, remaining after the key is derived (with public information known). We call this the *authorization metric*. The metric specifies how often for a given KS a legitimate user will be rejected. There are two ways of viewing this metric:

1) measuring the authentication strength possible from the source with the best known fuzzy extractor.
2) measuring the information theoretic capacity of the source for key derivation.

Both the identification and authentication metrics measure how frequently the legitimate user is granted access. The identification metric measures how frequently other users from the distribution are also granted access. When considering a determined attacker, this is not a sufficient threat model. For the authentication task, it is prudent to assume that the attacker has access to public parameters, P, and then measures the probability of recovering the key, given this information.

OPTIMIZING FOR THE ENTROPY METRIC: IRIS EXPERIENCE

THE IRIS CODE

The human iris is believed to be a strong biometric attribute [19]. The iris pattern develops in utero and is fairly stable throughout the lifetime of an individual [30]. Irises are diverse in small homogeneous populations (even right and left eyes of the same subjects appear to be independent) and are believed to be largely epigenetic (not dependent on genetic information), although some correlations may be observed [31], [32]. Typically, the near-infrared (NIR) images are used, although some research into using multispectral images has been undertaken [33].

For biometric applications, an iris is typically transformed into an iris code as follows (see Figure 4). First the image of the iris is segmented, locating the iris and the pupil. Then the iris image is unwrapped using polar coordinates, translating the iris from a two-dimensional (2-D) tor (doughnut) into a rectangle. Finally, the rectangle is subjected to various special processing (typically, filtering and quantization), producing a bit-vector called the *iris code*.

Modern transforms derive from the work of Daugman based on 2-D Gabor wavelets [3].The filtering phase uses 2-D Gabor wavelets [34] at various angular and radial coordinates. The image is divided into a polar grid of some angular and radial resolution and a wavelet is computed at each coordinate. We denote these values

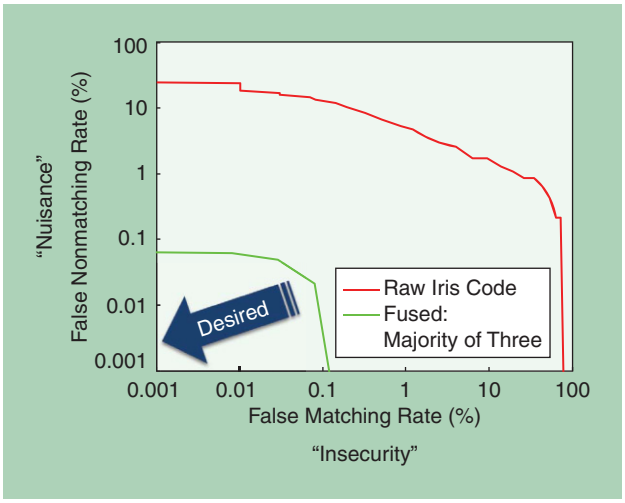

[FIG3] The match metric: FNMR versus FMR. The red plot is obtained using IrisCodes computed on iris images from the multispectral iris data set, licensed through the Scitor Corporation, using open source IrisCode software [35]. The green plot is obtained by faking the bit-wise majority of three repeated IrisCodes from the same user. As we discuss in the section "The Gap Between Biometrics and Fuzzy Extractors," while it is desired to reduce both FNMR and FMR, the latter corresponds to insecurity of the system, while the former reflects inconvenience for the users. For FNMR, the rates below 0.1% (eligible user failing to authenticate in less than one out of 1,000 attempts) can be considered quite acceptable. However, FMR range needs to be tens of orders of magnitude lower to be compatible with cryptographic security. This preference of FMR over FNMR is reflected by the angle of the "desired" arrow.

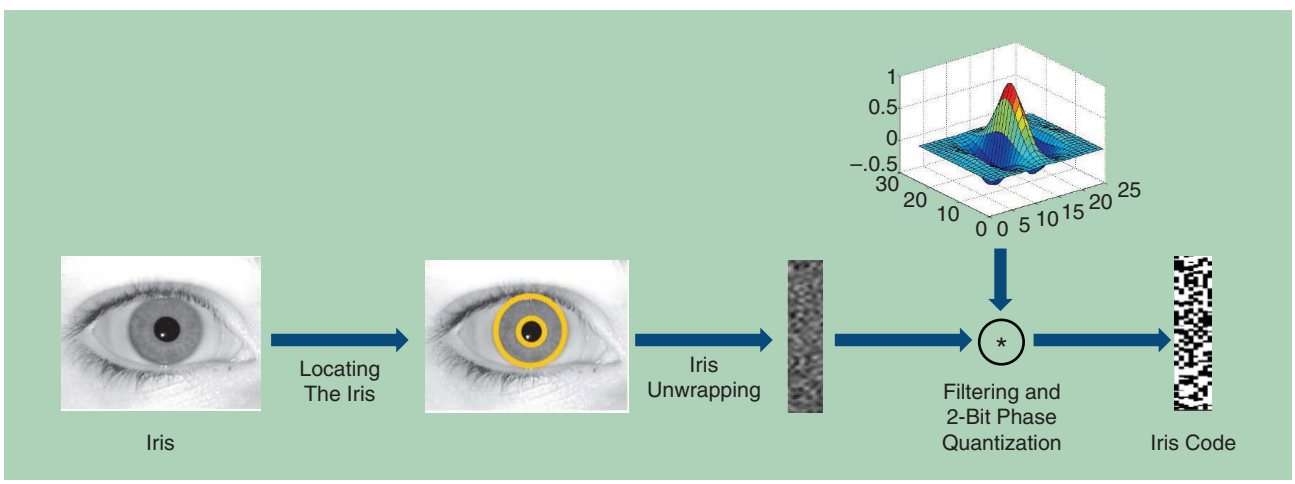

[FIG4] Deriving an iris code from the iris. The process starts with an image of an iris. The first step is segmentation: the iris is located in the image. Then the iris is unwrapped into a rectangle. The rectangular image is filtered using 2-D Gabor filters and quantized. The result is a bit array, called an iris code.

by ang_{res} and rad_{res}, respectively. The sign of the real and imaginary components yield 2-bit values for each location. The total length of the transform is $2 * \text{ang}_{\text{res}} * \text{rad}_{\text{res}}$. We then utilize the transform of Masek [35], which uses an angular resolution of $\text{ang}_{\text{res}} = 240$ and $\text{rad}_{\text{res}} = 20$ for an overall length of 9,600 bits. The transform of Daugman [3] provides superior performance but it is not publicly available. Performance is improved by applying simple automatic tests detecting some unsuccessful segmentations, to eliminate these from the statistical analysis, when working with various iris corpora.

We denote as w the iris code from an original reading collected from a user and stored as a reference template. During authentication, a new reading is then collected producing IrisCode w'. Typically, fractional Hamming distance (FHD)—the fraction of bits that differ between w and w'—is used as the distance between the two readings. Typical FHD between two readings of the same iris of the same person is between 10–30% (more careful tools can get it much lower). This is what is referred to as *in-class* FHD. For readings from different people, or different irises, FHD is within 40–60%. This is called *out-of-class* FHD.

In-class FHD can be increased by rotational distortions; e.g., when the image is taken at slightly different angles. Binary matching can compensate for this distortion by comparing the reference and authentication templates w, w' multiple times, applying a serious of small relative rotations, and picking the best (smallest) FHD. For the out-of-class FHD, this has a relatively negligible effect. Also, when reflections and occlusions occur, it is possible to simply ignore (mask) some portions of the image when the matching is performed. Both of these optimizations present additional challenges for cryptographic authorization. Next we consider the suitability of the iris for the cryptographic authorization task.

IRIS BIOMETRIC SUITABILITY FOR CRYPTOGRAPHIC AUTHORIZATION

In cryptographic authentication, we view users' inputs—in this case, irises—as sources of entropy.

IRIS CODE ENTROPY

Random strings have entropy essentially equal to their length. Unfortunately, most biometrics are not fully random. Estimating entropy of nontrivial distributions is a difficult problem [36], [37].

Assuming the existence of pseudorandom generators [38] (implied by one-way functions [39]), it is possible for distributions to appear to have significantly more entropy than they actually possess. This means it may be fundamentally impossible to estimate the entropy of distributions occurring through complex unknown processes.

Pseudorandom distributions are sophisticated and one may hope that they do not often appear in nature, or at least that biometric distributions appear to have the same entropy to all parties. There are several heuristics used to estimate entropy by comparison to well-known probability distributions. Daugman notes that the FHD between different individuals in an iris corpus fits a binomial distribution with mean $p = .5$ and $N = 249$ [3]. This yields an estimate of 249 bits of entropy.

As described in the section "Metric for Cryptographic Authentication," when measuring KS, we must include public parameters of the authentication system. This means we also need to consider the noise between repeated readings w and w'.

Rederiving the key essentially requires correcting w' to w. If we assume an error rate of ε, then the number of possible errors that might take place is $(n!/(n\varepsilon)!(n(1-\varepsilon))!)$. Hence, the error-correction requires $n(\varepsilon \lg \varepsilon + (1-\varepsilon)\lg(1-\varepsilon))$ bits of information, where n is the number of bits in w. For ε around 10%, this is approximately $n/2$ bits. For the transforms above, $n = 9,600$, this means that 4,800 bits of error correction information is necessary. Standard fuzzy extractor constructions may lose security proportional to this information. Since irises are estimated to have 249 bits of entropy, fuzzy extractors provide no guarantee on the resulting cryptographic key. The key challenge in deriving keys from irises is that the entropy rate and error rate are approximately the same. Fuzzy extractors provide good performance when the entropy rate is significantly higher than the error rate (see [40] for more details).

OPTIMIZING THE IRIS FOR THE AUTHENTICATION METRIC

As described in the previous section, a major obstacle to deriving keys from irises is the low entropy rate of current transforms. In this section, we discuss techniques for trying to improve the entropy rate of irises. We emphasize that none of these techniques currently seem sufficient to derive strong keys from irises. However, similar ideas will be necessary to derive strong keys from the iris.

SUBSAMPLING WAVELETS

Irises have a low entropy rate but each bit on its own behaves like a Bernoulli coin with $p = .5$. This indicates that each bit has full entropy and that the correlations in the iris exist between multiple bits. Thus, one approach to improve the iris transform is to try and find sets of bits that are uncorrelated with a similar error rate as the overall transform. Random subsampling preserves both entropy [41] and error rate. If iris bits are uncorrelated on large sets, then subsampling should produce an entropy rate higher than 10%. There has been work in the iris community on producing better transforms using structured subsampling. The entropy rate and error rate can be maintained while reducing length using random subsampling. If each bit of an iris is entropic on its own, random subsampling may be helpful. In particular, if any 249 bits can be used to reconstruct the iris, we can randomly subsample to 249 bits while maintaining all entropy and keeping error rate constant. The goal of structured subsampling techniques is to find better strategies.

The work of Gentile, Ratha, and Connell [42] introduced short-length iris codes, which were designed to improve processing speed of iris codes by reducing their length. Their work contains several observations: 1) the inside and outside of an iris tend

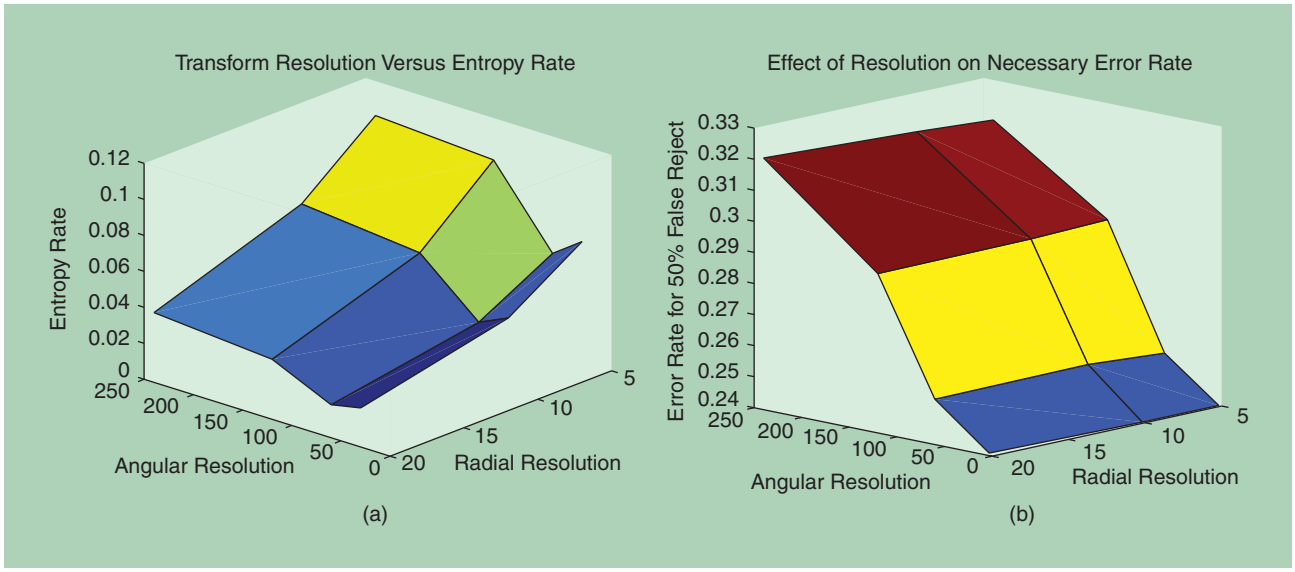

[FIG5] The effect of varying resolution on entropy and error rates. (a) Entropy rates and at various angular and radial resolutions. (b) Error rates at various angular and radial resolutions

to be less reliable due to increased deformations and occlusions, respectively, and 2) there is significant correlation between radially adjacent bits and little correlation between angularly adjacent bits. This leads them to create a transform that subsamples every tenth row of the iris code starting from the fifth row.

In Figure 5(a) and (b), we show the entropy rate and error rate at various resolutions. Observe the following:

- The entropy rate increases as the radial dimension is subsampled.
- The error rate remains constant as the radial dimension is subsampled.
- The entropy rate decreases slightly as the angular dimension is subsampled.
- The error rate decreases as the angular dimension is subsampled.

Figure 5 confirms the observations of Gentile et al. that there is significant redundancy in the radial dimension and this can safely be subsampled. We do note that although subsampling in the radial dimension improves the entropy rate, it does reduce the overall entropy. Careful analysis is needed to determine where the maximum strength key is possible.

FINDING THE BEST POSTTRANSFORM BITS

We will now look at subsampling in the bit domain (postwavelet transform). We start from a 9,600-bit transform. Each bit of the distribution follows a Bernoulli distribution with $p = 0.5$. However, it may be that some bits are more likely to contain errors and contribute more to the overall error rate. We now will try and find the bits that have the lowest error rate. This idea stems from the work of Hollingsworth et al. [43]. The results are shown in Figure 6; unfortunately, this graph is very flat, meaning there are not a large number of bits with lower error rate. Each bit has roughly the same error probability. This means subsampling at the bit level is unlikely to be helpful. We note that for a particular individual, there are bits that consistently have a lower error rate. Writing down such bits for a particular individual may reveal information about the original reading w. Thus, we only consider consistency of bits across the population.

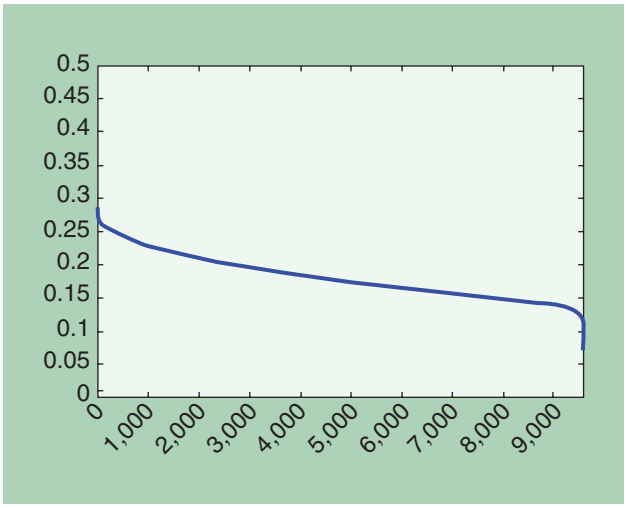

[FIG6] The best bits of an iris code.

DISCUSSION AND SUMMARY

For decades, authentication has relied on matching an original reading from a user against a previously captured and stored reference template. The binary outcome of the matching required a corresponding authorization mechanism to control access to the resources, granting the access based on the result of the matching. Implementations of this paradigm suffer from inherent weaknesses: fragility of the binary decisions, vulnerability of the stored reference templates, and the high value target of the authorization mechanisms.

Biometrics represent an important authentication source but with significant new challenges. Biometrics can be stolen and are not replaceable. Research on cancelable biometrics has tried to mitigate this issue (e.g., see [18] for a survey). These methods do not aim to replace the actual biometrics, of course, but rather they aim to replace a compromised (stolen) stored reference template with a different one. Most of these methods add certain distortions to the transformation to the canonical template, so that if a reference template is compromised for one distortion, a different one can be generated and used. However, only some (e.g., [26]) of these considered methods result in strong security, compatible to provable security of common cryptographic tools. This provides more motivation to move away from models where a template is stored, precisely because if that template is successfully attacked, then an attacker will be able to leverage that information remotely and at scale.

The main challenge to authentication using biometrics is their noisy nature: repeated readings can differ significantly. Current techniques for eliminating noise, such as fuzzy extractors, come at a significant entropy cost. However, we believe this approach has promise and that key derivation from noisy sources can be improved significantly. For noisy sources such as biometrics, existing processing algorithms have been optimized for identification, not authentication. Revisiting feature extraction for such sources with authentication in mind should reduce the entropy loss.

In this article, we show that the traditional FMR versus FNMR identification metric does not properly optimize for the authentication task. We instead propose using an authentication-specific metric, such as KS, rather than FMR versus FNMR. To illustrate the difference between these two approaches, we discuss attempts to optimize the iris biometric according to the authentication metric. Unfortunately, key derivation from the iris still remains a challenge. This article, and the works of our predecessors, lay the foundation for future progress in optimized key derivation from biometrics and their application to authentication systems.

ACKNOWLEDGMENTS

This research was enabled by access to the multispectral iris data set licensed through the Scitor Corporation. The multispectral iris data set was collected by Southern Methodist University (SMU) under the oversight of the SMU Institutional Review Board [44], [45] that enabled our research and experiments. We thank Jonas Borgstrom and Jason Thornton for their help in the early stages of our project, and Salil Prabhakar of Delta ID for providing us with an iris camera that we used for our prototypes. We are grateful to Shawn Campbell and Carolyn Greenberg for their assistance in preparing the manuscript and to Jessica Barragué for her assistance in the production of the article. This work is sponsored by Assistant Secretary of Defense for Research and Engineering under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are our own and are not necessarily endorsed by the United States Government.

AUTHORS

Gene Itkis (itkis@ll.mit.edu) received his Ph.D. degree in computer science from Boston University (BU) in 1996. From 1999 to 2009, he was on the faculty of the Computer Science Department at BU, where he founded the Applied Cryptography and e-Security group, the Industry Talks series, and served as the associate director for research at the Center for Reliable Information Systems and Cyber Security, National Security Agency National Center of Academic Excellence in Information Assurance Education. In 2009, he joined the Massachusetts Institute of Technology Lincoln Laboratory, where he is working on applied security projects including cloud security, key management, biometrics, and antitamper technologies, among others.

Venkat Chandar (chandarvenkat@verizon.net) received S.B. degrees in electrical engineering and computer sciences and mathematics in 2006, an M.Eng. degree in electrical engineering and computer sciences in 2006, and a Ph.D. degree in electrical engineering and computer sciences in 2010, all from the Massachusetts Institute of Technology (MIT). His current research interests include coding theory and algorithms, optical communications and quantum information theory, and more recently, finance theory. He was with MIT Lincoln Laboratory from 2010 until 2014, and is now a quantitative researcher at D.E. Shaw and Co.

Benjamin Fuller (bfuller@ll.mit.edu) joined the Massachusetts Institute of Technology (MIT) Lincoln Laboratory in 2007. His research focuses on cryptography and practical solutions to secure communication with a past focus on cryptographic key management and key derivation. He received the B.S. degree from Rensselaer Polytechnic Institute in 2006 and the M.A. and Ph.D. degrees from Boston University in 2011 and 2015, respectively. He completed his Ph.D. degree at Boston University under the direction of Prof. Leonid Reyzin, focusing on cryptography with imperfect and noisy randomness. His Ph.D. research focused on new approaches for the construction of fuzzy extractors.

Joseph P. Campbell (j.campbell@ieee.org) received the Ph.D. degree in electrical engineering from Oklahoma State University in 1992. Since 2001, he has been with the Massachusetts Institute of Technology's Lincoln Laboratory. He is currently the associate group leader of the Human Language Technology Group. He is a member of the IEEE Awards Planning and Policy Committee and chaired the Biometric Consortium and the IEEE Jack S. Kilby Signal Processing Medal Committee. He was a member of the IEEE Information Forensics Security Committee, the IEEE Signal Processing Technical Committee, and the IEEE Signal Processing Society's Board of Governors. He was an associate editor of *IEEE Transactions on Speech and Audio Processing*, the vice president of Technical Activities of the IEEE Biometrics Council, and a member of the IEEE Signal Processing Society Fellow Reference Committee. He was an IEEE Distinguished Lecturer and is an IEEE Fellow.

Robert K. Cunningham (rkc@ll.mit.edu) received the Ph.D. degree in computer engineering and cognitive and neural systems from Boston University in 1998. He is currently the leader of the Secure, Resilient Technology Group at the Massachusetts Institute

of Technology (MIT) Lincoln Laboratory. He won the 2015 MIT Excellence Award for Bringing Out the Best. He has published works on computer intrusion detection, computer worms, system protection, software development best practices, and on signal and image processing. He has served the IEEE as a program member for multiple conferences and workshops, and as a program and general chair for the IEEE Symposium on Technologies for Homeland Security and the IEEE Security and Privacy Symposium. He is a Senior Member of the IEEE.

REFERENCES

[1] J. H. Saltzer and M. D. Schroeder, "The protection of information in computer systems," *Proc. IEEE*, vol. 63, no. 9, pp. 1278–1308, 1975.

[2] H. M. Levy, *Capability-Based Computer Systems*. Newton, MA: Butterworth-Heinemann, 1984.

[3] J. Daugman, "How iris recognition works," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 14, no. 1, pp. 21–30, Jan. 2004.

[4] M. Zviran and W. J. Haga, "A comparison of password techniques for multilevel authentication mechanisms," *Comput. J.*, vol. 36, no. 3, pp. 227–237, 1993.

[5] S. Brostoff and M. Sasse, "Are passfaces more usable than passwords? A field trial investigation," *People Comput.*, pp. 405–424, 2000.

[6] C. Ellison, C. Hall, R. Milbert, and B. Schneier, "Protecting secret keys with personal entropy," *Future Generation Comput. Syst.*, vol. 16, no. 4, pp. 311–318, 2000.

[7] F. Monrose, M. K. Reiter, and S. Wetzel, "Password hardening based on keystroke dynamics," *Int. J. Inform. Security*, vol. 1, no. 2, pp. 69–83, 2002.

[8] J. Bonneau, C. Herley, P. C. v. Oorschot, and F. Stajano. (2012, Mar.). The quest to replace passwords: a framework for comparative evaluation of Web authentication schemes. Univ. Cambridge, Computer Lab., Tech. Rep. UCAM-CL-TR-817. [Online]. Available: http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-817.pdf

[9] R. Morris, R. Morris, K. Thompson, and K. Thompson, "Password security: A case history," *Commun. ACM*, vol. 22, pp. 594–597, Nov. 1979.

[10] N. K. Ratha, J. H. Connell, and R. M. Bolle, "Biometrics break-ins and band-aids," *Pattern Recognit. Lett.*, vol. 24, no. 13, pp. 2105–2113, 2003.

[11] J. Saltzer and M. Schroeder, "The protection of information in computer systems," *Proc. IEEE*, vol. 63, no. 9, pp. 1278–1308, Sept. 1975.

[12] A. Juels and M. Wattenberg, "A fuzzy commitment scheme," in *Proc. 6th ACM Conf. Computer and Communication Security.*, Nov. 1999, pp. 28–36.

[13] A. Juels and M. Sudan, "A fuzzy vault scheme," in *Proc. IEEE Int. Symp. Information Theory*, 2002, p. 408.

[14] A. Juels and M. Sudan. (2006). A fuzzy vault scheme. [Online]. *Designs, Codes Cryptogr.*, 38, 237–257. Available: http://dx.doi.org/10.1007/s10623-005-6343-z

[15] G. I. Davida, Y. Frankel, and B. J. Matt, "On enabling secure applications through off-line biometric identification," in *Proc. IEEE Symp. Security and Privacy.*, 1998, pp. 148–157.

[16] G. Davida, B. Matt, Y. Frankel, and R. Peralta, "On the relation of error correction and cryptography to an off line biometric based identification scheme," in Proc. *Workshop on Coding and Cryptography*, 1999, pp. 129–138.

[17] U. Uludag, S. Pankanti, S. Prabhakar, and A. K. Jain, "Biometric cryptosystems: issues and challenges," *Proc. IEEE*, vol. 92, no. 6, pp. 948–960, 2004.

[18] C. Rathgeb and A. Uhl, "A survey on biometric cryptosystems and cancelable biometrics," *EURASIP J. Inform. Security*, vol. 2011, no. 1, pp. 1–25, 2011.

[19] S. Prabhakar, S. Pankanti, and A. K. Jain, "Biometric recognition: Security and privacy concerns," *IEEE Security Privacy*, vol. 1, no. 2, pp. 33–42, 2003.

[20] M. Blanton and W. M. Hudelson, "Biometric-based non-transferable anonymous credentials," in *Information and Communications Security*. New York: Springer, 2009, pp. 165–180.

[21] E. Grosse and M. Upadhyay, "Authentication at scale," *IEEE Security Privacy*, vol. 11, no. 1, pp. 15–22, Jan./Feb. 2013.

[22] J. Galbally, A. Ross, M. Gomez-Barrero, J. Fierrez, and J. Ortega-Garcia, "From the iriscode to the iris: A new vulnerability of iris recognition systems," in *Black Hat Briefings USA*, 2012. [Online]. Available: https://www.blackhat.com/html/bh-us-12/bh-us-12-briefings.html#Galbally

[23] J. Galbally, A. Ross, M. Gomez-Barrero, J. Fierrez, and J. Ortega-Garcia, "Iris image reconstruction from binary templates: An efficient probabilistic approach based on genetic algorithms," *Comput. Vis. Image Understand.*, vol. 117, no. 10, pp. 1512–1525, Oct. 2013.

[24] A. Ross, J. Shah, and A. Jain, "From template to image: Reconstructing fingerprints from minutiae points," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 4, pp. 544–560, Apr. 2007.

[25] T. Ignatenko and F. M. Willems, "Biometric security from an information-theoretical perspective," in *Foundations and Trends in Communications and Information Theory*. Now Publishers, pp. 135-316, 2012.

[26] J. Bringer and H. Chabanne, "An application of the naccache-stern knapsack cryptosystem to biometric authentication," in *2007 IEEE Workshop on Automatic Identification Advanced Technologies*, June 2007, pp. 180–185.

[27] T. Dierks and E. Rescorla, "The transport layer security (TLS) protocol," *IETF RFC 5246*, 2008. [Online]. Available: https://tools.ietf.org/html/rfc5246

[28] J. Daugman, "Probing the uniqueness and randomness of iriscodes: Results from 200 billion iris pair comparisons," *Proc. IEEE*, vol. 94, no. 11, pp. 1927–1935, 2006.

[29] S. Ziauddin and M. N. Dailey, "Iris recognition performance enhancement using weighted majority voting," in *Proc. 15th IEEE Int. Conf. Image Processing (ICIP)*, 2008, pp. 277–280.

[30] S. P. Fenker and K. W. Bowyer, "Experimental evidence of a template aging effect in iris biometrics," in *2011 IEEE Workshop on Applications of Computer Vision (WACV)*, 2011, pp. 232–239.

[31] J. Howard and D. Etter, "The effect of ethnicity, gender, eye color and wavelength on the biometric menagerie," in *Proc. IEEE Int. Conf. Technologies for Homeland Security (HST)*, 2013, pp. 627–632.

[32] S. Lagree and K. W. Bowyer, "Predicting ethnicity and gender from iris texture," in *Proc. 2011 IEEE Int. Conf. Technologies for Homeland Security (HST)*, 2011, pp. 440–445.

[33] C. K. Boyce, "Multispectral iris recognition analysis: techniques and evaluation," M.S.E.E. thesis, West Virginia Univ., Dept. Comput. Sci. Electric. Eng., Citeseer, 2006.

[34] I. Daubechies, "The wavelet transform, time-frequency localization and signal analysis," *IEEE Trans. Inform. Theory*, vol. 36, no. 5, pp. 961–1005, 1990.

[35] L. Masek, "Recognition of human iris patterns for biometric identification," Bachelor's thesis, School Comp. Sci. Software Eng., Univ. Western Australia, 2003.

[36] J. Beirlant, E. J. Dudewicz, L. Györfi, and E. C. Van der Meulen, "Nonparametric entropy estimation: An overview," *Int. J. Math. Stat. Sci.*, vol. 6, no. 1, pp. 17–39, 1997.

[37] T. Schürmann, "Letter to the editor: Bias analysis in entropy estimation," *J. Phys. A*, vol. 27, pp. L295–L301, July 2004.

[38] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby, "A pseudorandom generator from any one-way function," *SIAM J. Comput.*, vol. 28, no. 4, pp. 1364–1396, 1999.

[39] L. A. Levin. (2000). The tale of one-way functions. [Online]. CoRR, vol. cs.CR/0012023. Available: http://arxiv.org/abs/cs.CR/0012023

[40] R. Canetti, B. Fuller, O. Paneth, L. Reyzin, and A. Smith. (2014). Key derivation from noisy sources with more errors than entropy. [Online]. Cryptology ePrint Archive, Report 2014/243. Available: http://eprint.iacr.org/

[41] S. P. Vadhan, "Constructing locally computable extractors and cryptosystems in the bounded-storage model," *J. Cryptol.*, vol. 17, no. 1, pp. 43–77, 2004.

[42] J. Gentile, N. Ratha, and J. Connell, "SLIC: Short-length iris codes," in *IEEE 3rd Int. Conf. Biometrics: Theory, Applications, and Systems (BTAS'09)*, 2009, pp. 1–5.

[43] K. P. Hollingsworth, K. W. Bowyer, and P. J. Flynn, "The best bits in an iris code," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 6, pp. 964–973, 2009.

[44] D. Etter, J. Webb, and J. Howard, "Collecting large biometric datasets: A case study in applying software best practices," *Immutable Laws Softw. Develop., CrossTalk: J. Defense Softw. Eng.*, pp. 4–8, May/June 2014.

[45] J. J. Howard, "Large scale pattern recognition models for identifying subject specific match probability across datasets with controlled variability," Ph.D. dissertation, Southern Methodist Univ., 2014.

[46] C. H. Bennett, G. Brassard, and J.-M. Robert, "Privacy amplification by public discussion," *SIAM J. Comput.*, vol. 17, no. 2, pp. 210–229, 1988.

[47] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith, "Fuzzy extractors: How to generate strong keys from biometrics and other noisy data," *SIAM J. Comput.*, vol. 38, no. 1, pp. 97–139, 2008.

[48] N. Nisan and D. Zuckerman, "Randomness is linear in space," *J. Comput. Syst. Sci.*, vol. 52, no. 1, pp. 43–52, Feb. 1996.

[49] O. Dunkelman, M. Osadchy, and M. Sharif, "Secure authentication from facial attributes with no privacy loss," in *Proc. ACM SIGSAC Conf. Computer & Communications Security.*, 2013, pp. 1403–1406.

[50] B. Fuller, X. Meng, and L. Reyzin, "Computational fuzzy extractors," in *Proc. Advances in Cryptology (ASIACRYPT)*, 2013, pp. 174–193.

[51] C. Herder, L. Ren, M. van Dijk, M.-D. M. Yu, and S. Devadas. (2014). Trapdoor computational fuzzy extractors. [Online]. Cryptology ePrint Archive, Rep. 2014/938. Available: http://eprint.iacr.org/

[52] B. Kaliski, "Pkcs# 5: Password-based cryptography specification version 2.0," in *IETF RFC* 2898. RFC Editor, 2000. [Online]. Available: http://dx.doi.org/10.17487/RFC2898

[53] U. Feige, A. Fiat, and A. Shamir, "Zero-knowledge proofs of identity," *J. Cryptol.*, vol. 1, no. 2, pp. 77–94, 1988.

[SP]

[Vishal M. Patel, Nalini K. Ratha, and Rama Chellappa]

Cancelable Biometrics

[A review]

©ISTOCKPHOTO.COM/GREYFEBRUARY

Biometrics Security and Privacy Protection

Recent years have seen an exponential growth in the use of various biometric technologies for trusted automatic recognition of humans. With the rapid adaptation of biometric systems, there is a growing concern that biometric technologies may compromise the privacy and anonymity of individuals. Unlike credit cards and passwords, which can be revoked and reissued when compromised, biometrics are permanently associated with a user and cannot be replaced. To prevent the theft of biometric patterns, it is desirable to modify them through revocable and noninvertible transformations to produce cancelable biometric templates. In this article, we provide an overview of various cancelable biometric schemes for biometric template protection. We discuss the merits and drawbacks of available cancelable biometric systems and identify promising avenues of research in this rapidly evolving field.

Digital Object Identifier 10.1109/MSP.2015.2434151
Date of publication: 13 August 2015

INTRODUCTION

Biometrics refers to the physiological or behavioral characteristics of an individual. Many physical characteristics, such as face, fingerprints, and iris and behavioral characteristics, such as voice, gait, and keystroke dynamics, are believed to be unique to an individual. Hence, biometric analysis offers a reliable solution to the problem of identity verification. Recent developments in sensing and computing technologies have made biometric systems more affordable and, as a result, they are easily embedded in a variety of smart consumer devices such as mobile phones and tablets. Despite the widespread deployment of biometric systems in

1053-5888/15©2015IEEE

various applications, the use of biometrics raises several security and privacy concerns as outlined below [1].

1) *Biometrics is not secret*: The knowledge-based authentication methods totally rely on secrecy. For instance, passwords and cryptographic keys are known only to the user and hence secrecy can be maintained. In contrast, biometrics such as voice, face, signature, and even fingerprints can be easily recorded and potentially misused without the user's consent. Face and voice biometrics are vulnerable to being captured without the user's explicit knowledge.

2) *Biometrics cannot be revoked or canceled*: If a biometric can be presented by a human being who is one of the enrolled users, many biometrics security issues will be different. For example, biometric-based authentication systems will not have to deal with spoofed biometrics and also replay attacks on biometric systems. If a hacker gets access to the biometrics samples and has the ability to present them to a system emulating a human presence, there will be no trust associated with the biometrics. In this scenario, we say that the biometrics have been compromised forever. Passwords, crypto-keys and PINs can be changed if compromised. When items such as credit cards and badges are stolen, they can be replaced. However, biometrics is permanently associated with the user and cannot be revoked or replaced if compromised.

3) *Cross application invariance and cross-matching*: It is highly encouraged to use different passwords and tokens in traditional authentication systems. However, biometrics-based authentication methods rely on the same biometrics. If a biometric template is exposed once, it is compromised forever. If a biometric template is compromised in one application, then the same method can be used to compromise all applications where the biometric is used. Furthermore, since the same biometrics is used across all applications and locations, the user can be potentially tracked if one or more organizations collude and share their respective biometric databases.

4) *Persistence*: While relative robustness over time is a boon for biometrics it can also be a big challenge from a privacy point of view when it needs to be changed. The uniqueness contained in them is still the same even though the signal as well the template can look different.

Regarding privacy violations, cross-matching and the inability to revoke a biometric are two major issues. A simple approach would be to use standard encryption techniques such as hash functions or encryption to enhance the privacy. Hash functions have been used to protect biometric templates in which one-way functions are used to compute a digest. Even though these functions are almost impossible to invert, they produce a significantly different digest even with minor changes in the input. In practice, all biometric templates change with environmental conditions. For instance, face and iris biometrics are significantly affected by illumination variations. Therefore, these functions cannot be used directly in practice despite being theoretically very strong as they apply only to exact data. Furthermore, when data are encrypted, they need to be decrypted to carry out matching. This creates a possible attack point to get access to the decrypted templates.

To overcome the vulnerabilities of biometric systems, both the biometrics and cryptoresearch communities have addressed some of the challenges. Several biometric template protection schemes have been proposed in the literature [2]–[8]. In particular, cancelable biometrics [3]–[5], [9] has gained a lot of interest in recent years. In this method, instead of storing the original biometric, it is transformed using a one-way function. The transformation can be applied either in the original domain or in the feature domain. It was shown that this way of constructing biometric templates has the desired properties of cancelable biometric templates [3]–[5]. In particular, it provides revocability since a compromised biometric can be re-enrolled using another transformation. It preserves privacy since it is computationally difficult to recover the original biometric from a transformed one. It prevents cross-matching between databases since each application uses a different transformation, and it does not degrade the accuracy of a matching algorithm as the statistical characteristics of features are approximately maintained after transformation. This allows one to use existing matching algorithms.

There are also some closely related but not equivalent biometric template protection schemes based on *cryptosystems* [10] that have been studied extensively. These methods combine cryptographic keys with transformed versions of the original biometric templates to generate secure templates. In these methods, some public information, known as *helper data*, is generated. Depending on how the helper data is used, biometric cryptosystems can be broadly classified into *key binding* and *key generation* systems. In the key generation systems, both the helper data and the key are directly generated from the biometric templates, while in the key binding systems, the helper data are obtained by combining the key with the biometric template. Examples of key binding systems include fuzzy commitment [11] and fuzzy vault [6]. Key generation schemes based on secure sketches [7] have also been proposed in the literature. In the biometric cryptosystems, the level of security depends on the amount of information revealed by the helper data. Other methods for biometric template protection include distributed source coding [12] and fuzzy extractors [13]. A review of biometric cryptosystems can be found in [8], [10], and [14].

Our goal in this article is to survey recent available approaches for designing cancelable biometric templates, discuss their advantages and limitations, and identify areas still open for exploration. Furthermore, we will discuss possible ways of attacking cancelable biometric systems. The development of cancelable schemes for biometric template protection is crucial as biometric systems are beginning to proliferate into the core physical and information infrastructure of our dynamic society.

CANCELABLE BIOMETRIC TEMPLATES

In this section, we review a number of recent strategies for generating cancelable biometric templates. In these methods, a function that is dependent on some parameter is used to generate protected biometric templates. The parameter of the function is used as the key. Figure 1 shows the basic concept of cancelable biometric template-based on noninvertible transformations.

NONINVERTIBLE GEOMETRIC TRANSFORMS

One of the earliest methods for generating cancelable biometric templates was based on noninvertible geometric transformations.

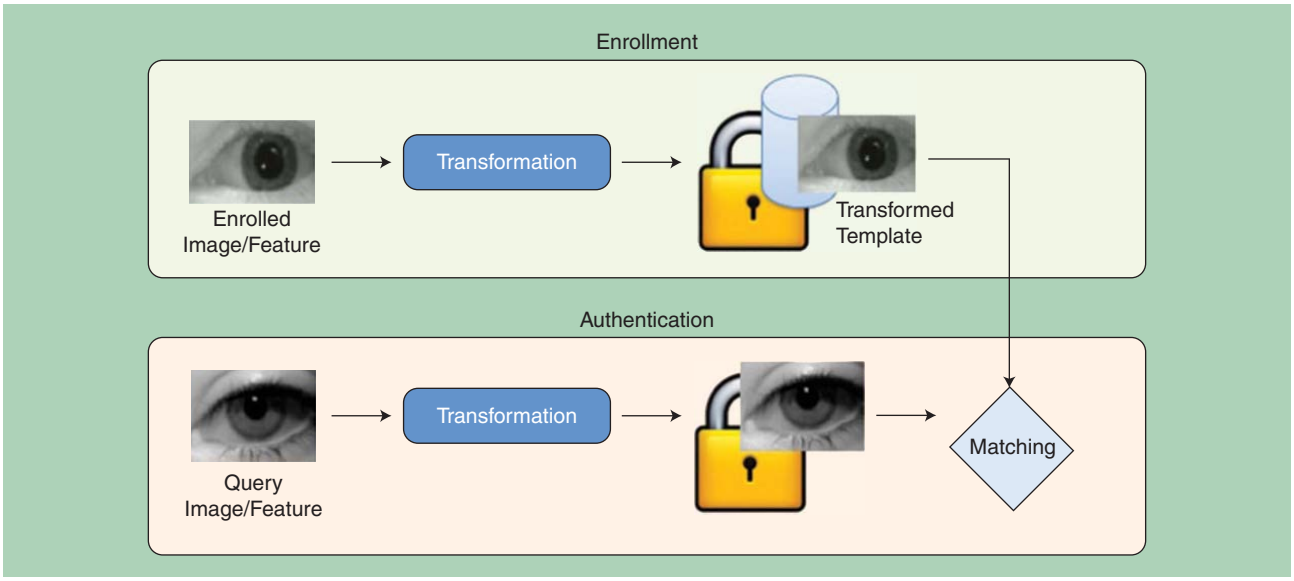

[FIG1] **A block diagram of a cancelable biometric system.**

The idea is to morph the original biometric templates by applying signal domain or feature domain transformations [3]–[5]. Figure 2(a) and (b) shows examples of these transformations applied in the signal domain and in the feature domain, respectively, for face and fingerprint biometrics. Three different transformations were proposed for the fingerprint biometric in [4] and [5]. These transformations are the Cartesian transformation, the polar transformation, and the functional transformation. Prior to applying these transformations, the images are registered by first estimating the position and orientation of the singular points (core and delta) and expressing the minutiae positions and angles with respect to these points. Registration is an integral part of this method.

For the Cartesian transformation, the minutiae positions are measured in rectangular coordinates with reference to the position of the singular point by aligning the x-axis with its orientation. The coordinate system is divided into cells of fixed size. The transformation consists of changing the cell positions [4], [5]. Note that this transformation is not a simple permutation as the condition of irreversibility requires that cells are mapped to the same cell.

In the polar transformation method, the minutiae positions are measured in the polar coordinate with reference to the core position. The angles are measured with respect to the core orientation. As a result, the coordinate space is divided into polar regions. The noninvertible transform consists of changing the polar wedge positions. The minutiae angles also change with differences in the wedge positions before and after transformation [4], [5].

One of the limitations of both polar and Cartesian transformations is that they are unstable in the sense that a small change in

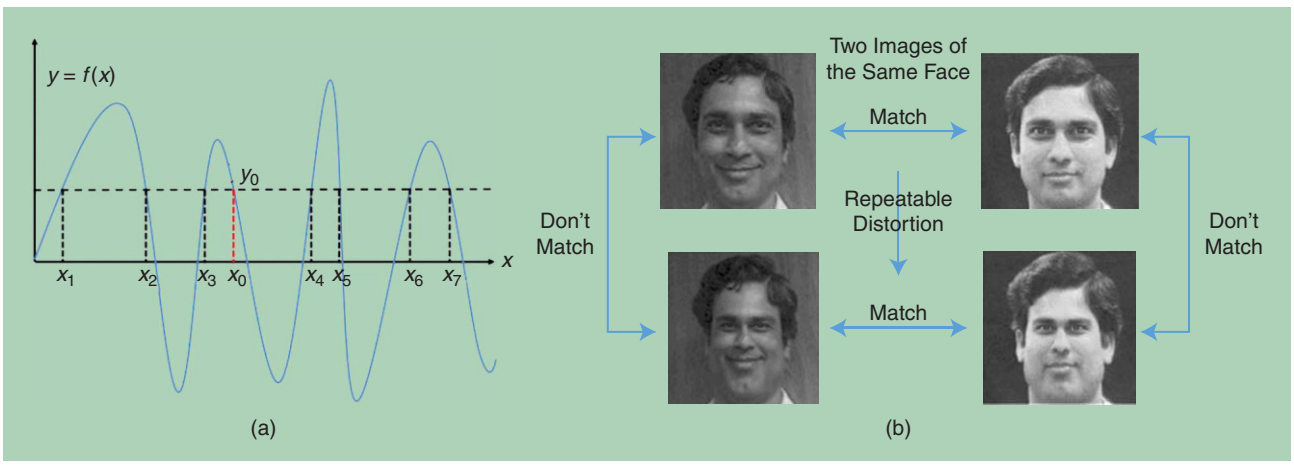

[FIG2] **An illustration of a nonlinear transformation applied to face and fingerprint biometrics [4]. (a) Feature-domain transformation for fingerprint biometrics. Each minutiae (feature) position is transformed using a noninvertible function $y = f(x)$. The minutiae position x_0 is mapped to $y_0 = f(x_0)$. If we know y_0, the inverse mapping is a many-to-one transformation. x_1, x_2, x_3, x_4, x_5, x_6, x_7 are all valid inverse mappings to y_0. (b) An illustration of cancelable biometrics for face recognition. The face is distorted in the original pixel (signal) domain prior to feature extraction. The distorted version does not match with the original face, while the two instances of distorted faces match among themselves.**

minutiae position in the original fingerprint can lead to a large change in minutiae position after transformation if the point crosses a sharp boundary [4]. As a result, various functions giving a locally smooth transformation of the minutiae positions were introduced in [4] and [5]. The transformation is modeled using a vector valued function $\vec{F}(x, y)$ whose phase determines the direction of translation and the extent of translation is given by the magnitude $|\vec{F}|$ or alternately another vector valued function $\vec{G}(x, y)$. One such function proposed in [4] and [5] is an electric potential field parameterized by a random distribution of charges. The magnitude and phase of this function are given by

$$|\vec{F}| = \left| \sum_{i=1}^{K} \frac{q_i (z - z_i)}{|(z - z_i)|^3} \right|,$$
$$\Phi(x, y) = \frac{1}{2} \arg\left(\sum_{i=1}^{K} \frac{q_i (z - z_i)}{|(z - z_i)|^3} \right),$$

where $z = x + iy$ is the position vector and the random key $K = [z_1, z_2, \cdots, z_K, q_1, q_2, \cdots, q_K]$ determines the position and magnitude of the charges. The transformation is given by

$$x' = x + K|\vec{G}(x, y)| + K\cos(\Phi_F(x, y))$$
$$y' = y + K|\vec{G}(x, y)| + K\sin(\Phi_F(x, y))$$
$$\theta' = \mathrm{mod}(\theta + \Phi_G(x, y) + \Phi_{\mathrm{rand}}, 2\pi).$$

See [4] for more examples of various transformations and their analysis in terms of noninvertibility and attack strength. This method was later extended in [15] so that it does not require the registration of images. However, the approach in [15] exhibits lower verification rates than [4].

In a related work, [16] proposes a mesh warping-based approach for generating cancelable iris templates. In this method, the iris texture is remapped according to a distorted grid mesh laid over it. Distortions are specified by a key that offsets each vertex in the original mesh by some unspecified amount. Specifically, a regular grid is placed over the texture in which the vertices are then randomly displaced using the key as seed to a random number generator.

RANDOM PROJECTIONS

Another noninvertible transformation that is widely used for generating cancelable biometric templates is based on random projections [17], [18]. In these methods, the extracted feature $\mathrm{x} \in \mathbb{R}^N$ from a biometric is projected onto a random subspace $\mathrm{A} \in \mathbb{R}^{n \times N}$ with $n < N$. Here, each entry $a_{i,j}$ of A is an independent realization of a random variable. This process is described as follows

$$\mathrm{y} = \mathrm{Ax}, \tag{1}$$

where y is the n-dimensional random projection vector. Since we are embedding N dimensional feature vectors in a space of a lower dimension n, for any biometric recognition to be effective, it is important that the relative distances between any two points in the feature space be preserved in the output random space. This is essentially characterized by the Johnson–Lindenstrauss (JL) lemma [19].

Lemma 1: For any $0 < \varepsilon < 1$ and any integer p, let n be a positive integer such that $n \geq (4\ln(p)/(\varepsilon^2/2 - \varepsilon^3/3))$. Then, for any set $\mathcal{S}$ of $p = |\mathcal{S}|$ data points in $\mathbb{R}^N$, there is a map $f: \mathbb{R}^N \to \mathbb{R}^n$ such that, for all x, y $\in \mathcal{S}$,

$$(1 - \varepsilon)\|\mathrm{x} - \mathrm{y}\|^2 \leq \|f(\mathrm{x}) - f(\mathrm{y})\|^2 \leq (1 + \varepsilon)\|\mathrm{x} - \mathrm{y}\|^2. \tag{2}$$

This lemma essentially states that a set $\mathcal{S}$ of points in $\mathbb{R}^N$ can be embedded into a lower-dimensional Euclidean space $\mathbb{R}^n$ such that the pairwise distance of any two points is approximately maintained. In fact, it can be shown that f can be taken as a linear mapping represented by an $n \times N$ matrix A whose entries are randomly drawn from certain probability distributions [19]. This in turn implies that it is possible to change the original form of the data and still preserve its statistical characteristics useful for recognition.

In recent years, various improvements in the proof and the statement of the JL lemma have been made (see [20] and [21] for more details). In fact, it has been shown that given any set of points $\mathcal{S}$, the following are some of the matrices that will satisfy (2) with high probability, provided n satisfies the condition of the Lemma 1 [21]:

- $n \times N$ random matrix A whose entries $a_{i,j}$ are independent realizations of Gaussian random variables $a_{i,j} \sim \mathcal{N}(0, (1/n))$.
- Independent realizations of ± 1 Bernoullie random variables

$$a_{i,j} = \begin{cases} +\frac{1}{\sqrt{n}} & \text{with probability } \frac{1}{2} \\ -\frac{1}{\sqrt{n}} & \text{with probability } \frac{1}{2}. \end{cases}$$

- Independent realizations of related distributions such as

$$a_{i,j} = \begin{cases} +\sqrt{\frac{3}{n}} & \text{with probability } \frac{1}{6} \\ 0 & \text{with probability } \frac{2}{3} \\ -\sqrt{\frac{3}{n}} & \text{with probability } \frac{1}{6}. \end{cases}$$

A random projection-based cancelable biometric method for iris recognition was proposed in [17]. Applying the random projections directly on the iris images usually degrades the performance due to the following reasons. First of all, in real iris images, despite good segmentation algorithms, there will still be some outliers due to specular reflections, eye lashes and eyelids. Also, different parts of the iris have different quality. By taking a linear transformation of the entire vector, one combines the good iris regions as well as the outliers and thereby corrupts the data. To deal with this, [17] proposes sectored random projections (SRPs) in which random projections are applied separately on each sector and the resulting transformed vectors are concatenated to form the cancelable template. As a result, outliers can corrupt only the corresponding sector and not the entire iris vector.

Figure 3 shows an overview of this method [17]. The enrollment system extracts the iris pattern of the user, computes the Gabor features, applies a different random projection for each application, and transfers the new pattern to the application database. Note that even if the transformed pattern and the key (i.e., the projection matrix) are stolen, the user's iris pattern cannot be generated from them due to the dimension reduction caused by the projection. Also, even if a hacker steals the user's iris pattern

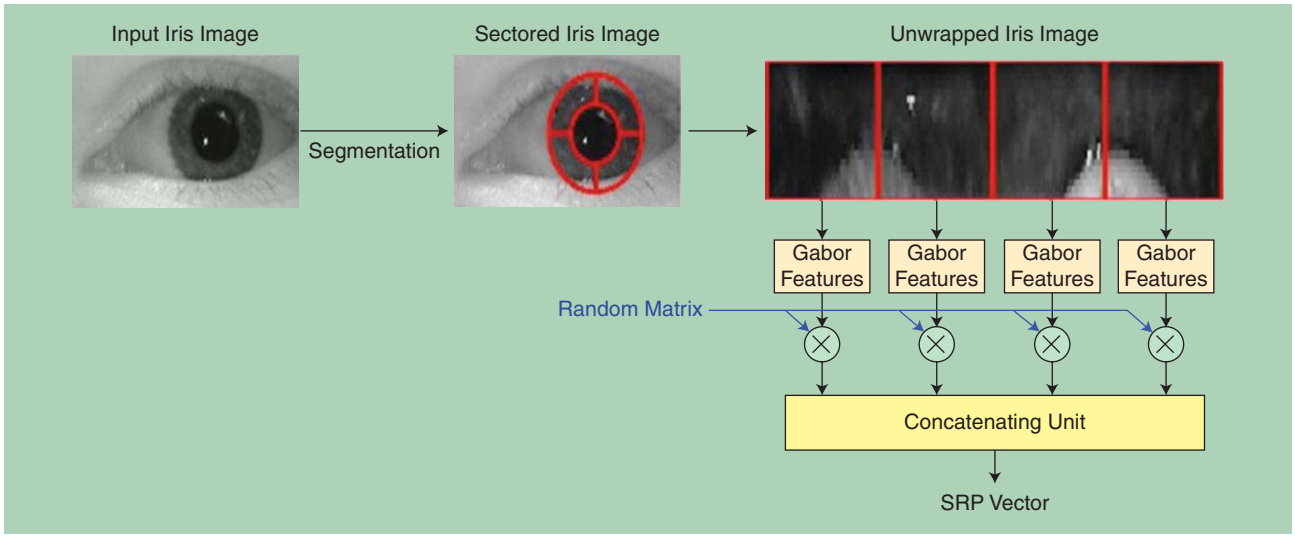

[FIG3] **An overview of the SRP method [17].**

either from the client system or using a hidden scanner, without knowing the random projection he/she cannot generate the transformed patterns required by the application. During the verification stage, the application obtains the iris image and the random projection matrix from the user, computes the transformed pattern, and compares it with the ones in its database. In case the random projection matrix or the transformed patterns are compromised, one can create a new random projection matrix and obtain a new transformed pattern that can be updated into the application database. Instead of the user providing the random matrix during verification, the application can generate and store it along with the cancelable template in its database. Though this will be an easier scheme for the user to operate, it is less secure as a hacker can get both the random projection matrices and the transformed patterns by breaking into the application database.

The approach of the SRP method [17] was later extended in [18] using sparse representation-based classification. It was shown that the sparsity patterns that one obtains before and after applying random projections are similar. As a result, cancelable biometric templates can be directly used for authentication rather than the original ones without degrading the performance of a sparse representation-based classification algorithm.

CANCELABLE BIOMETRIC FILTERS

Motivated by the success of the correlation filter-based methods in pattern recognition and computer vision applications [22], a random convolution method for generating cancelable biometric templates was proposed in [23]. The idea is to encrypt biometric templates using random user-specific convolution kernels. The training images are convolved with a random convolution kernel. The seed used to generate the random convolution kernel is used as the personal identification number (PIN). The convolved training images are then used to generate a minimum average correlation energy (MACE) biometric filter. This encrypted filter is stored and used for authentication. Figure 4(a) shows the enrollment stage using this method.

During the recognition stage, the user presents the PIN and the encrypted filter that is used to generate the convolution kernel. This random convolution kernel is convolved with the test face

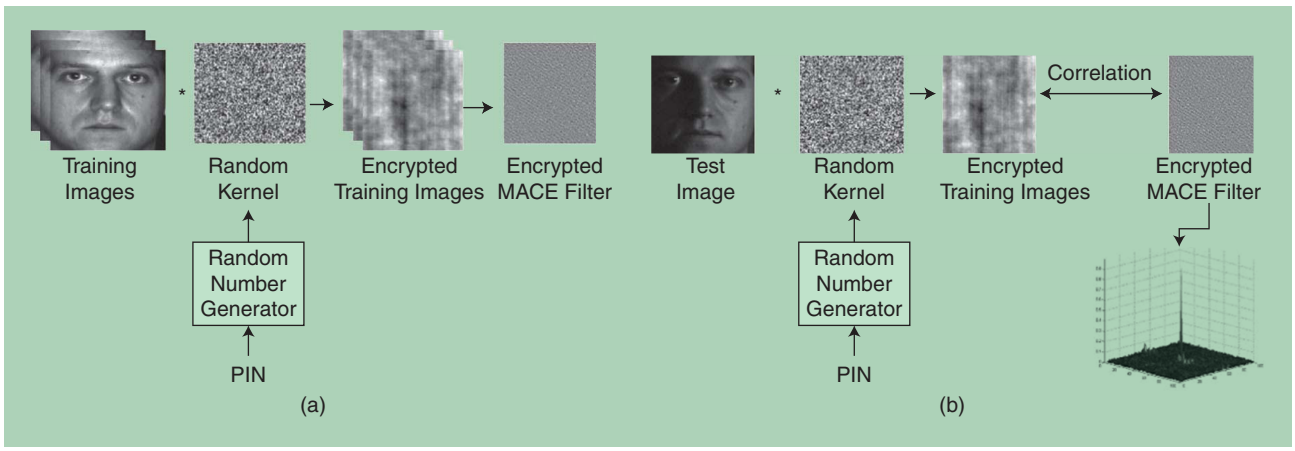

[FIG4] **A correlation filter-based approach to cancelable biometrics [23]. (a) The enrollment stage for encrypted filters. (b) The authentication stage using encrypted MACE filters.**

images presented by the user. The convolved test images are cross-correlated with the encrypted MACE filter and the resulting correlation outputs are used to authenticate the user. Figure 4(b) shows the authentication stage for this method.

It was shown that convolving the training images with any random convolution kernel prior to building the MACE filters used for biometric recognition does not change the resulting correlation output [23]. As a result, the recognition accuracy is maintained. Furthermore, different cancelable biometric templates can be generated from the same biometric by simply changing the convolution kernels.

Other correlation-based cancelable biometric methods include correlation invariant random filtering (CIRF) [24], [25], which was shown to have almost the same accuracy as the conventional fingerprint verification based on the chip matching algorithm.

BIOCONVOLVING

Another convolution-based approach for generating cancelable biometric templates was recently proposed in [26]. This method is applicable to any biometric whose template can be represented by a set of sequences. In this method, each transformed sequence $f_{(i)}[n], i = 1, \cdots, F,$ is obtained from the corresponding original sequence $r_{(i)}[n], i = 1, \cdots, F,$ which represents a generic discrete sequence of length N belonging to the original biometric template. In particular, a number $(W-1)$ of different integer values d_j between 1 and 99 are randomly selected, ordered in ascending order such that $d_j > d_{j-1,}\ j = 1, \ldots, W,$. These numbers are arranged in a vector $\mathbf{d} = [d_0, \cdots, d_W]^T$, where d_0 and d_W are set to 0 and 100, respectively. Here, the vector $\mathbf{d}$ represents the key of the transformation. The original sequence $r_{(i)}[n]$ is divided into W nonoverlapping segments $r_{(i)j,N_j}[n]$ of length $N_j = b_j - b_{j-1}$

$$r_{(i)j,N_j}[n] = r_{(i)}[n + b_{j-1}], n = 1, \cdots, N_j, j = 1, \cdots, W, \quad (3)$$

where

$$b_j = \left\lfloor \frac{d_j}{100} N \right\rfloor, \quad j = 1, \cdots, W. \quad (4)$$

A transformed sequence $f_{(i)}[n], n = 1, \cdots, K,$ is then obtained through the linear convolution of the sequences $r_{(i)j,N_j}[n]$, $j = 1, \cdots, W$ as

$$f_{(i)}[n] = r_{(i)1,N_1}[n] * \cdots * r_{(i)W,N_W}[n]. \quad (5)$$

Each original sequence $r_{(i)}[n], i = 1, \cdots, F$ undergoes the same decomposition before applying convolutions. As a result, the length of the transformed sequences is equal to $K = N - W + 1$. A normalization step is applied to make the transformed sequences zero mean and unit standard deviation. Different templates can be generated from the original biometric template by simply changing the size or the values of the parameter key $\mathbf{d}$. Figure 5 shows an example of a feature transformation where $W = 3$ [26]. See [26] for more details on different ways of generating transformed sequences, invertibility analysis, and their application in signature-based authentication.

BLOOM FILTERS

Recently, Bloom filter-based cancelable biometric template protection schemes were proposed in [27]–[29]. A Bloom filter is essentially a space-efficient probabilistic data structure representing a set to support membership queries. In particular, alignment-free cancelable iris biometric templates based on adaptive Bloom filters were introduced in [27] in which the generic adaptive Bloom filter-based transform is applied to binary feature vectors of different iris recognition algorithms. It was shown that such a method can enable template protection, compression of biometric data, and computationally efficient biometric identification. Furthermore, rotation-invariant Bloom filter-based transform can provide a high level of security while maintaining recognition accuracy [27].

KNOWLEDGE SIGNATURES

Voice-based cancelable biometric templates using knowledge signatures were proposed in [30]. The idea is based on a group signature scheme, which allows members of a group to sign messages on the group's behalf such that the resulting signature does not reveal their identity. They consider voiceprint as the knowledge of the user and the user's voiceprint transmitted to the template, which isn't the original feature but a signature of knowledge. Legitimate signatures cannot be generated without factorizing a large integer and the original feature. As a result, an individual's privacy can be protected. We refer readers to [30] and [31] for more details on knowledge signatures and their uses in generating cancelable biometric templates for voiceprints.

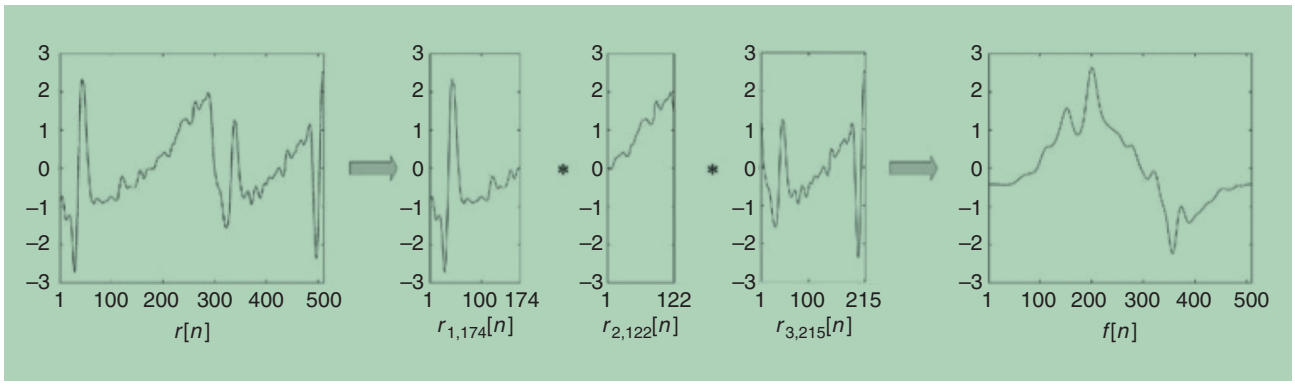

[FIG5] A sequence transformation using the BioConvolving approach [26].

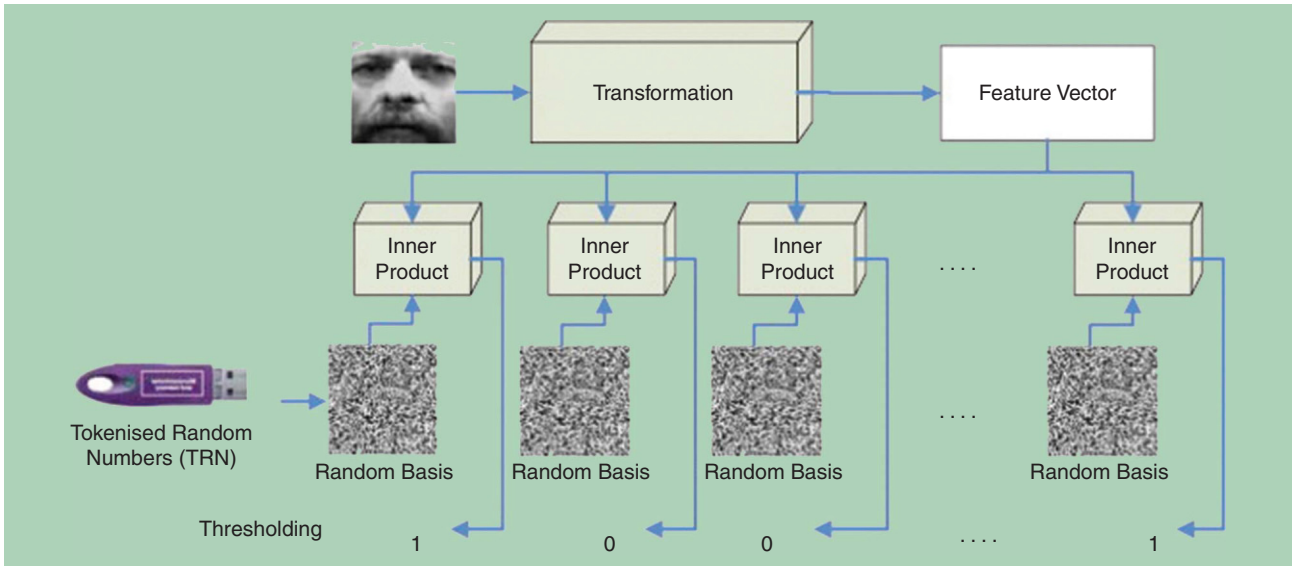

[FIG6] An overview of BioHashing [33].

BIOHASHING METHODS

BioHashing methods are essentially an extension of random projection. In BioHashing [9], [32]–[38], a feature extraction method such as wavelet transform is first used to extract the biometric feature $\mathrm{x} \in \mathbb{R}^N$ from the input biometric data. Using a user-specific tokenized random number (TRN), n orthogonal pseudo-random vectors, $b_i \in \mathbb{R}^N, i = 1, \cdots, n$ are generated, where $n \leq N$. The dot product of the feature vector and all the random vectors is then calculated. Finally, a binary discretization is applied to compute the n bit BioHash template as

$$\mathrm{c} = \mathrm{Sig}\left(\sum_i \mathrm{x}\mathrm{b}_i - \tau\right), \tag{6}$$

where $\mathrm{Sig}(\cdot)$ is defined as a signum function and τ is an empirically determined threshold. Equation (6) only applies to a user who holds the user-specific random vectors $\mathbf{b}_i \in \mathbb{R}^N, i = 1, \cdots, n$, and thus the formulation can be extended to introduce an ensemble of random subspaces, where each subspace represents different individual k. The resulting BioHash is given as

$$\mathrm{c}^k = \mathrm{Sig}\left(\sum_i \mathrm{x}^k \mathrm{b}_i^k - \tau\right), \quad k = 1, \cdots, g, \tag{7}$$

where g is the total number of users in the system. Finally, the BioHash code is compared by the Hamming distance for the similarity matching. Figure 6 shows the progression of BioHashing [33]. The BioHashing framework is demonstrated to be a one-way transform, hence providing a high degree of security to the biometric and external factors. A detailed statistical analysis of the BioHashing framework in terms of random multispace quantization operations can be found in [9].

RANDOM PERMUTATIONS

Another common approach for generating cancelable biometric templates is based on a random permutation of features. In [39], two such methods were proposed for generating cancelable iris templates. The first method, GRAY-COMBO, transforms the Gabor features by circularly shifting and adding rows at random. BIN-COMBO, the second method, applies similar transformations on the iris codes by random shifting and XOR-ing. As pointed out by the authors, these methods gradually reduce the amount of information available for recognition. Since these methods employ linear transformations on the Gabor feature vectors, they are also sensitive to outliers in the form of eyelids, eyelashes, and specular reflections. Reference [18] proposes to overcome this limitation by dividing the feature into different regions and permuting them randomly in a dictionary. Without prior knowledge of the locations of sectored features in a dictionary, it is impossible to perform recognition. A similar approach was also proposed in [16], where each block of the target texture is mapped to a block from the source texture. In this method, a remapping of blocks instead of a permutation is performed, as it is not reversible. Source blocks that are not part of the mapping are not contained in the transformed texture. As a result, it is impossible to reconstruct the original iris texture. Another permutation-based cancelable method for fingerprint biometric was presented in [40]. This method permutes a binary vector obtained from fingerprint features and stores them in the database. During authentication, the binary vector obtained from the fingerprints of the user are permuted using the key provided by the user and matched with the database.

In these methods, key security is essential for protecting the privacy of individuals. One of the advantages of these methods is that since permutations are merely rearranging the feature vector, authentication accuracy is not affected by these operations.

SALTING METHODS

One of the simplest ways of generating cancelable biometric templates is by mixing in a totally artificial pattern. The mixing patterns can be pure random noise, a random pattern, or a synthetic pattern. Two such salting methods were proposed in [39] for iris recognition: GRAY-SALT and BIN-SALT. These methods add random patterns or synthetic iris patterns to the Gabor features and

iris codes, respectively. Unlike GRAY-COMBO and BIN-COMBO permutation-based methods, they do not suffer from the problem of outlier amplification and reduction of useful area. However, it is difficult to decide the relative strength of the noise patterns to be added. Adding very strong patterns will reduce the discriminative capacity of the original iris patterns and hence lead to lower recognition results. Adding weaker patterns can lower the noninvertibility property, making it easier to extract useful information about the original iris biometric from the transformed patterns. Also, if the added patterns are compromised, the original iris patterns could be extracted from the transformed patterns by a simple subtraction operation.

HYBRID METHODS

Several biometric template protection approaches make use of both cryptosystems and cancelable biometrics [41], [42]. One such hybrid system was proposed in [41] for face biometrics. They introduced *biotoken,* which is "the revocable identity token produced by applying a revocable transform to biometric data, such that identity matching is done in the encoded/revocable form" [41], [42]. Specifically, this approach combines the ideas of transformation of data, robust learning measures, and encryption of biometric data. The method essentially separates the data into two parts, the fractional part, which is retained for local distance computation, and the integer part, which is encrypted. It was shown that for face biometrics, this method significantly improved the performance of the methods based on principal component analysis (PCA) and linear discriminant analysis (LDA) algorithms. This work was later extended for fingerprints in [42].

SUMMARY OF CANCELABLE BIOMETRIC TEMPLATE PROTECTION SCHEMES

The cancelable biometric template protection schemes reviewed in this article can be broadly divided into two main categories as shown in Figure 7—methods that require a special matcher and methods that can work with the existing matchers. These schemes can be further classified into two categories: 1) registration-free methods and 2) methods that require good registration of biometric samples. Finally, these methods can be further divided into two types of schemes—1) schemes that work with the original biometric samples (denoted as signal) and 2) schemes that work with the features extracted from the biometric signals (denoted as feature).

Among the methods that require a special matcher and good registration of biometric samples, Biotokens [41], [42] is a signal-based method and BioConvolving [26], salting, PalmHashing [37], and PalmPhasor [37] methods are feature-based methods. On the other hand, the correlation-based MACE filter approach [23] is a signal-based, registration-free method that requires a special matcher. COMBO [39], block-remapping [16], image warping [16], noninvertible transforms [4], and dynamic random projection (RP) [43] methods fall under signal-based methods that require registration and can work with the existing matchers, whereas permutations [18], RP [17], BioHashing [35], [9] and PalmHashing [34] methods are feature based that can work with the existing matchers and require good registration. Registration-free methods that can work with existing matchers include minimum distance graph [44] and curtailed circular convolution [45] methods (signal-based methods) and a registration-free approach proposed in [15] (which is a feature-based method).

Furthermore, Table 1 summarizes key cancelable biometric template protection approaches in terms of their performances on various biometric data sets. Note that the performances of different methods are reported in terms of false rejection rate (FRR), equal error rate (EER), rank-1 recognition rate (RR), genuine accept rate (GAR), and false acceptance rate (FAR).

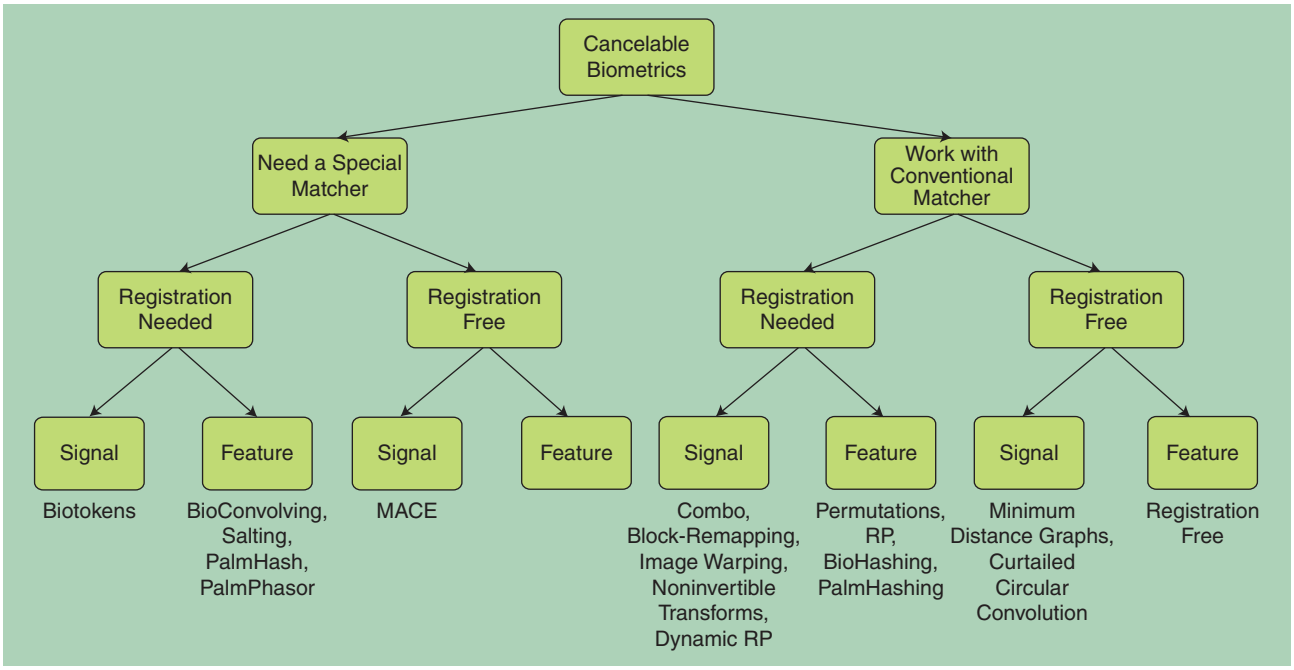

[FIG7] **The categorization of cancelable biometric template protection schemes.**

[TABLE 1] KEY CANCELABLE BIOMETRIC TEMPLATE PROTECTION SCHEMES.

METHOD	BIOMETRIC	DATA SET (SUBJECTS)	PERFORMANCE	REMARKS
NONINVERTIBLE TRANSFORMS [4]	FINGERPRINT	IBM-99 (188)	FRR: ~35, 15, 15	—
IMAGE WARPING [16]	IRIS	CASIA IRIS V3 (396)	EER: 1.6 – 6	—
RANDOM PROJECTIONS [17]	IRIS	MMU1 DATA SET (100)	RR: 97.7	—
BIOMETRIC FILTERS [23]	FACE	CMU PIE (65)	RR: 100	—
BIOCONVOLVING [26]	ONLINE SIGNATURE	MCYT (330)	EER: 6.33 – 15.40	—
BIOHASHING [35]	FINGERPRINT	FVC 2002 (100)	EER: ~ 0	FAR: 0
PALMHASHING [34]	PALMPRINT	PALMPRINT DATA SET (50)	EER: 0–0.222	FAR: 0
BIOHASHING [9]	FACE	FERET (1196)	EER: 0.002 – 7.51	—
GRAY-COMBO [39]	IRIS	MMU1 DATA SET (100)	GAR: ~ 0.995	FAR : 10^{-4}
BIN-COMBO [39]	IRIS	MMU1 DATA SET (100)	GAR: ~ 0.965	FAR : 10^{-4}
BLOCK REMAPPING [16]	IRIS	CASIA IRIS V3 (396)	EER: 0.2 – 1.6	—
GRAY-SALT [39]	IRIS	MMU1 DATA SET (100)	GAR: ~1	FAR : 10^{-4}
BIN-SALT [39]	IRIS	MMU1 DATA SET (100)	GAR: ~ 0.995	FAR : 10^{-4}
ATOM PERMUTATIONS [18]	IRIS	ND-IRIS-0405 (356)	RR: 99.17	—
BIOTOKENS [42]	FINGERPRINT	FCV 2000-4 (100)	EER: 0.012–0.086	HYBRID METHOD
BIOTOKENS [41]	FACE	FERET (1196)	EER: 0.9997	HYBRID METHOD
DYNAMIC RANDOM PROJECTIONS [43]	FINGERPRINT	FVC2002DB2-A (800)	EER: ~0.05	—
PALMHASH CODE [37]	PALMPRINT	POLYU DATA SET (7752)	EER: 0.38	2-D PALMHASH CODE
PALMPHASOR CODE [37]	PALMPRINT	POLYU DATA SET (7752)	EER: 0.32	2-D PALMPHASOR CODE
MINIMUM DISTANCE GRAPH [44]	FINGERPRINT	FVC2002-DB1A,B (100)	EER: 0.0227	—
CURTAILED CIRCULAR CONVOLUTION [45]	FINGERPRINT	FVC2002-DB1,2,3 (100)	EER: 0.02, 0.03, 0.0612	—

ATTACKS AGAINST CANCELABLE BIOMETRIC TEMPLATES

A generic biometric system consists of a sensor, a feature extraction module, a biometric template database, a matcher module, and an application device that is driven by the matcher's response. Researchers have identified different points of attacks in a biometric system as shown in Figure 8. The attacks can come in various forms such as: trojan horse attack, front-end attack, phishing and farming attacks, back-end attack, and communication channel attack. The unauthorized access to raw biometric templates is among the most serious threats to users' privacy and security. Some of the attacks can be averted using cancelable biometric systems while others are extremely difficult to detect. See [46] for more details on different types of attacks.

The cancelable biometric systems can be attacked by exposing the parameter (key) of the transformation being applied to biometric templates. In the case when the transformation is invertible, the original biometric can be reconstructed. In this case, security is in the secrecy of the key. If the transformation is not invertible, then an attacker can try to approximately recover the original biometric templates. For instance, it was shown in [47]–[49] that face images can be restored from encrypted templates. Attacks against the cancelable system using noninvertible transforms [4] are proposed in [50]. It was argued that when multiple transformed templates are generated from the same original template, they can be cracked by a method known as *attack via record multiplicity*. In particular, given a transformed template, an attacker can find the inverse solutions by inverting the transformation. Due to the many-to-one property of transform functions, there may be several solutions out of which is the original solution. The attacker can come up with a way to to pick out the right solution. A similar dictionary attack method is also proposed in [51] to recover the original templates from the cancelable templates. Also, convolution-based cancelable biometric systems' [23], [26] security depends on how well the blind deconvolution algorithms are able to recover the original biometric templates.

Several vulnerabilities in BioHashing-based systems have also been investigated [52]–[55]. One of the major limitations of BioHashing methods is their low performance when attackers are in possession of a secret key. To deal with this problem, [52] proposed an improved BioHashing method that is more robust than the original BioHashing method [32], [33]. In [53], it was shown that even without having a genuine users' private random vectors, a preimage of a BioCode can be easily calculated from a lost BioCode. As a result, an attacker can gain an illegal access to a system. It was observed that simple data dimension reduction and discretization as is done in most BioHashing methods may be vulnerable to preimage attacks. Similarly, a new way to approximate the original biometric feature from the transformed template in a cancelable biometric scheme was recently proposed in [54]. Their method is based on a genetic algorithm which essentially determines the optimal value of a criterion by simulating the evolution of a population and survival of best fitted individuals [54]. It was shown that a genetic algorithm can allow an intruder to recover a biometric template, similar to the original template, under some realistic assumptions.

In a related work, [43] analyzed the security concerns over random projection-based cancelable systems [49], [56] and proposed a dynamic random projection method to alleviate these concerns by

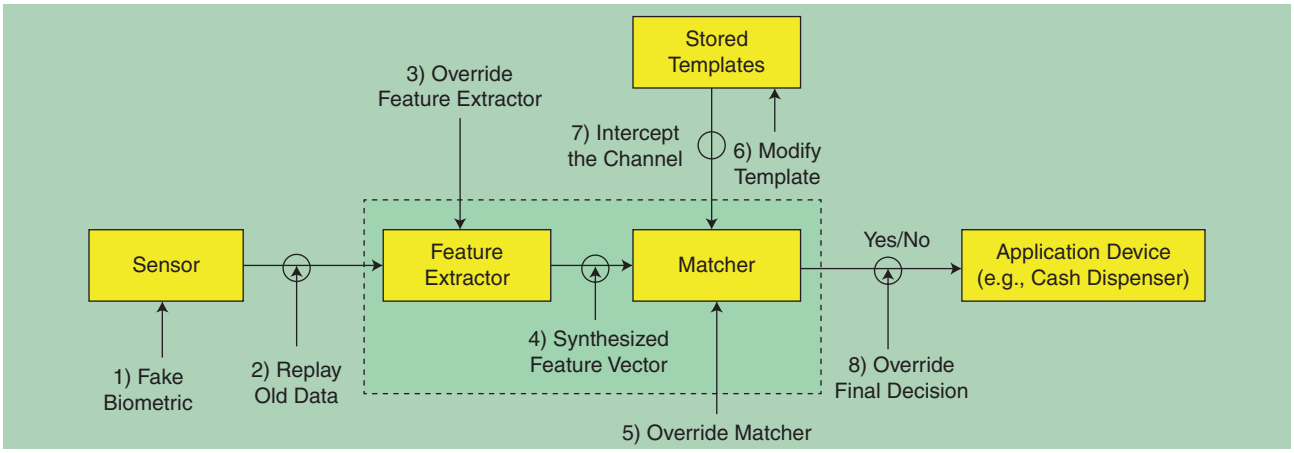

[FIG 8] **Possible attack points in a generic biometric system [8], [46].**

forming a nonlinear projection process that relates the random matrix's assembly to the biometric feature vector itself. The dynamic random projection method greatly increases the computational complexity to apply inversion attacks in the token stolen cases. Furthermore, it was shown that this method does not degrade the performance compared with the fixed matrix-based random projection [43].

In recent years, several biometrics protection schemes have been proposed in the literature that attempt to protect the privacy of biometric templates without using a key [57]–[60]. For instance, a visual cryptography method is introduced in [58] that decomposes a biometric image into two noise-like images, called *sheets*, that are stored in two different databases. During authentication, the two sheets are overlaid to create a temporary image for matching. One of the limitations of this method is that it requires two separate databases to work together, which may not be practical in some applications. Another method for protecting fingerprint biometrics combines two fingerprints from two different fingers to generate a new template [57]. For authentication, two query fingerprints are required and a two-stage matching process is proposed for matching the two query fingerprints against a combined template. One of the advantages of this method is that by using the combined template, the complete minutiae feature of a single fingerprint will not be compromised when the database is stolen [57].

To deploy a biometric template protection system, one needs to investigate the security strength of the template transformation technique and define metrics that facilitate security evaluation. Toward this end, six different evaluation metrics were defined in [49]. Furthermore, the security of BioHashing and cancelable fingerprint templates were analyzed based on these metrics. It was reported that both these schemes are vulnerable to intrusion and linkage attacks because it is relatively easy to obtain either an approximation of the original biometric template in the case of BioHashing or a preimage of the transformed template in the case of cancelable fingerprints.

In a related work, [61] presents several evaluation criteria, metrics and testing methodologies for assessing biometric template protection algorithms. In particular, criteria such as accuracy of the recognition algorithm, throughput, storage requirements, performance degradation of a biometric template protection scheme, diversity, and error rate of failing to generate a protected template are discussed in detail. These definitions will help researchers in designing robust biometric template protection schemes.

CONCLUSIONS AND FUTURE DIRECTIONS

We often use standard information security tools such as encryption or secure hashing methods to protect the biometric content. There are two issues with this approach. First, as the biometrics data (image, template) constantly change with every sample acquisition, the encrypted biometrics has to be decrypted for matching. If it is decrypted, that opens an opportunity for the hacker to attack at the output point of the decryption. If a secure hash function is used, the matching of the secure hashes is totally useless as biometrics signals never reproduce exactly. While the hash will be best in terms of privacy, the biometrics matching will not ever produce the positive authentication result. Cancelable biometrics is inspired by this approach but handles biometric variability. The transformation management in cancelable biometrics is equivalent to key management in information security. For example, a part of the transform can be retained by the user, another part can stay with the authentication system. Until the two come together, the biometrics authentication can't take place. But the keys in encryption or hash functions are derived totally differently than the cancelable biometrics transform. Second, because of the special construction, the matching of the cancelable biometrics signal or template is carried out in the transformed domain. In fact, the original biometrics signal is not required to be retained as both enrollment and authentication is carried out using the transformed biometrics.

This article presented a review of recent developments in such template protection schemes which included noninvertible transform-based methods, BioHashing, and hybrid methods. There are challenges that must be overcome before successfully designing a cancelable biometric system. Below we list a few.

- For the transform to be repeatable, the biometric signal must be positioned in the same coordinate system each time. This requires that an absolute registration be done for each biometric signal acquired prior to the transformation.

Registration-free cancelable biometric systems have also been proposed in the literature [15], [44], [45], [62]. However, some of these methods do not perform well in practice. For instance, a registration-free construction of cancelable fingerprint biometric templates [15] exhibits lower verification performances than the one proposed in [4], which requires registration. More robust registration-free noninvertible transform and BioHashing methods are needed.

- The recently introduced theory of compressive sampling allows one to reconstruct the original signal from a few random measurements provided that certain conditions are met. Many cancelable biometric template protection systems make use of random projections [17], [18], [32], [33]. It remains an interesting problem to study the vulnerability of such cancelable systems using compressive sampling.
- Over the past few years, we have witnessed an exponential growth in the use of mobile devices such as smartphones and tablets. Most mobile devices use passwords, PINs, or secret patterns for authenticating users. As long as the device remains active, there is no mechanism to verify that the user originally authenticated is still the user in control of the device. As a result, unauthorized individuals may improperly gain access to personal information of the user if the password is compromised. Active authentication systems deal with this issue by continuously monitoring the user identity after the initial access has been granted. Examples include systems based on screen touch gestures [63], gait recognition [64], and device movement patterns (as measured by the accelerometer) [65]. The development of cancelable active authentication systems is a nascent area of research.
- Blind deconvolution is an extremely ill-posed problem in which one attempts to recover the original signal from convolved outputs without the explicit knowledge of the convolution kernel. Recent advances in the signal processing community have shown that one can approximate the convolution kernel directly from the observations. These methods exploit some underlying structure of signals such as sparsity. It remains to be seen whether convolution-based cancelable systems are robust to these blind deconvolution methods.
- Most cancelable biometric template protection schemes have been evaluated on small and midsize data sets consisting of hundreds and thousands of samples. However, to really see the significance and impact of various biometric template protection schemes, they need to be evaluated on large-scale data sets containing millions of samples.
- As the research community advances biometric template protection schemes, third-party evaluation for security attacks and evaluation of the revocable methods are needed. Some efforts are being made [66], but more standardization efforts are needed to establish guidelines and procedures for testing and evaluating various cancelable biometric systems.

AUTHORS

Vishal M. Patel (pvishalm@umd.edu) received B.S. degrees in electrical engineering and applied mathematics (with honors) and the M.S. degree in applied mathematics from North Carolina State University, Raleigh, in 2004 and 2005, respectively. He received the Ph.D. degree from the University of Maryland in 2010. He is a member of the research faculty at the University of Maryland Institute for Advanced Computer Studies. His research interests are in signal processing, computer vision, and machine learning with applications to biometrics and object recognition. He was a recipient of the Oak Ridge Associated Universities postdoctoral fellowship in 2010. He is a Member of the IEEE, IEEE-Eta Kappa Nu, Pi Mu Epsilon, and Phi Beta Kappa.

Nalini K. Ratha (ratha@us.ibm.com) received the B.Tech. degree in electrical engineering and the M.Tech. degree in computer science and engineering from the Indian Institute of Technology, Kanpur. He received his Ph.D. degree from Michigan State University, East Lansing, in 1996. He is a research staff member at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, where he leads the biometrics research efforts in building efficient biometrics systems. In addition to publishing more than 80 papers in peer-reviewed journals and conferences, he is the co-inventor on 12 patents and has coedited two books on biometrics recognition. He is a Fellow of the IEEE, the International Association for Pattern Recognition, and a senior member of the Association for Computing Machinery. His current research interests include biometrics, computer vision, pattern recognition, and special-purpose architecture for computer vision systems.

Rama Chellappa (rama@umiacs.umd.edu) received the Ph.D. degree from Purdue University in 1981. He is a professor and chair of electrical and computer engineering and an affiliate professor of computer science at the University of Maryland, College Park. He is also affiliated with the Center for Automation Research, the Institute for Advanced Computer Studies (permanent member), and a Minta Martin Professor of Engineering. He is a Fellow of the IEEE, the International Association for Pattern Recognition, the Optical Society of America, the American Association for the Advancement of Science, the Association for Computing Machinery, and the Association for the Advancement of Artificial Intelligence. His current research interests are clustering, three-dimensional modeling from video, image- and video-based recognition of objects, dictionary-based inference, and domain adaptation.

REFERENCES

[1] N. K. Ratha, "Privacy protection in high security biometrics applications," in *Ethics and Policy of Biometrics*, (Lecture Notes in Computer Science, vol. 6005), A. Kumar and D. Zhang, Eds. Berlin, Germany: Springer, 2010, pp. 62–69.

[2] G. I. Davida, Y. Frankel, and B. J. Matt, "On enabling secure applications through off-line biometric identification," in *IEEE Symp. Security and Privacy*, May 1998, pp. 148–157.

[3] N. K. Ratha, J. H. Connel, and R. Bolle, "Enhancing security and privacy in biometrics-based authentication systems," *IBM Syst. J.*, vol. 40, no. 3, pp. 614–634, 2001.

[4] N. Ratha, S. Chikkerur, J. Connell, and R. Bolle, "Generating cancelable fingerprint templates," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 4, pp. 561–572, Apr. 2007.

[5] R. M. Bolle, J. H. Connel, and N. K. Ratha, "Biometrics perils and patches," *Pattern Recogn.*, vol. 35, no. 12, pp. 2727–2738, 2002.

[6] A. Juels and M. Sudan, "A fuzzy vault scheme," *Designs, Codes Cryptogr.*, vol. 38, no. 2, pp. 237–257, Feb. 2006.

[7] Y. Sutcu, Q. Li, and N. Memon, "Protecting biometric templates with sketch: Theory and practice," *IEEE Trans. Inform. Forensics Security*, vol. 2, no. 3, pp. 503–512, Sept. 2007.

[8] A. K. Jain, K. Nandakumar, and A. Nagar, "Biometric template security," *EURASIP J. Adv. Signal Process.*, vol. 2008, pp. 113:1–113:17, Jan. 2008.

[9] A. Teoh, A. Goh, and D. Ngo, "Random multispace quantization as an analytic mechanism for biohashing of biometric and random identity inputs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 12, pp. 1892–1901, Dec. 2006.

[10] U. Uludag, S. Pankanti, S. Prabhakar, and A. Jain, "Biometric cryptosystems: issues and challenges," *Proc. IEEE*, vol. 92, no. 6, pp. 948–960, June 2004.

[11] A. Juels and M. Wattenberg, "A fuzzy commitment scheme," in *Proc. ACM Conf. Computer and Communications Security*, New York, NY, 1999, pp. 28–36.

[12] S. Draper, A. Khisti, E. Martinian, A. Vetro, and J. S. Yedidia, "Using distributed source coding to secure fingerprint biometrics," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, Apr. 2007, vol. 2, pp. II–129–II–132.

[13] Y. Dodis, L. Reyzin, and A. Smith, "Fuzzy extractors: How to generate strong keys from biometrics and other noisy data," in *Advances in Cryptology—Eurocrypt 2004*. Berlin, Germany: Springer, 2004, pp. 523–540.

[14] C. Rathgeb and A. Uhl, "A survey on biometric cryptosystems and cancelable biometrics," *EURASIP J. Inform. Security*, vol. 2011, no. 3, pp. 1–25, 2011.

[15] S. Chikkerur, N. Ratha, J. Connell, and R. Bolle, "Generating registration-free cancelable fingerprint templates," in *Proc. IEEE Int. Conf. Biometrics: Theory, Applications and Systems*, Sept. 2008, pp. 1–6.

[16] J. Hmmerle-Uhl, E. Pschernig, and A. Uhl, "Cancelable iris biometrics using block re-mapping and image warping," in *Information Security* (Lecture Notes in Computer Science, vol. 5735), P. Samarati, M. Yung, F. Martinelli, and C. Ardagna, Eds. Berlin, Germany: Springer, 2009, pp. 135–142.

[17] J. K. Pillai, V. M. Patel, R. Chellappa, and N. K. Ratha, "Sectored random projections for cancelable iris biometrics," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, Mar. 2010, pp. 1838–1841.

[18] J. K. Pillai, V. M. Patel, R. Chellappa, and N. K. Ratha, "Secure and robust iris recognition using random projections and sparse representations," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 9, pp. 1877–1893, Sept. 2011.

[19] W. Johnson and J. Lindenstrauss, "Extensions of Lipschitz maps into a Hilbert space," in *Proc. Contemporary Mathematics*, 1984, pp. 189–206.

[20] S. Dasgupta and A. Gupta, "An elementary proof of a theorem of Johnson and Lindenstrauss," *Random Struct. Algorithms*, vol. 22, no. 1, pp. 60–65, Jan. 2003.

[21] D. Achlioptas, "Database-friendly random projections: Johnson-Lindenstrauss with binary coins," *J. Comput. Syst. Sci.*, vol. 66, no. 4, pp. 671–687, June 2003.

[22] B. V. K. V. Kumar, A. Mahalanobis, and R. D. Juday, *Correlation Pattern Recognition*. Cambridge, U.K.: Cambridge Univ. Press, 2005.

[23] M. Savvides, B. Kumar, and P. Khosla, "Cancelable biometric filters for face recognition," in *Proc. Int. Conf. Pattern Recognition*, Aug. 2004, vol. 3, pp. 922–925.

[24] K. Takahashi and S. Hirata, "Cancelable biometrics with provable security and its application to fingerprint verification," *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.*, vol. 94-A, no. 1, pp. 233–244, 2011.

[25] S. Hirata and K. Takahashi, " Cancelable biometrics with perfect secrecy for correlation-based matching," in *Advances in Biometrics* (Lecture Notes in Computer Science, vol. 5558), M. Tistarelli and M. Nixon, Eds. Berlin, Germany: Springer, 2009, pp. 868–878.

[26] E. Maiorana, P. Campisi, J. Fierrez, J. Ortega-Garcia, and A. Neri, "Cancelable templates for sequence-based biometrics with application to on-line signature recognition," *IEEE Trans. Syst., Man Cybern. A*, vol. 40, no. 3, pp. 525–538, May 2010.

[27] C. Rathgeb, F. Breitinger, C. Busch, and H. Baier, "On the application of bloom filters to iris biometrics," *IET J. Biometrics*, vol. 3, no. 4, pp. 207–218, 2014.

[28] C. Rathgeb and C. Busch, "Cancelable multi-biometrics: Mixing iris-codes based on adaptive bloom filters," *Comput. Security*, vol. 42, pp. 1–12, 2014.

[29] C. Rathgeb, F. Breitinger, and C. Busch, "Alignment-free cancelable iris biometric templates based on adaptive bloom filters," in *Proc. IAPR Int. Conf. Biometrics*, June 2013, pp. 1–8.

[30] W. Xu, Q. He, Y. Li, and T. Li, "Cancelable voiceprint templates based on knowledge signatures," in *Proc. Int. Symp. Electronic Commerce and Security*, Aug. 2008, pp. 412–415.

[31] J. Camenisch and M. Stadler, "Efficient group signature schemes for large groups," in *Proc. Int. Cryptology Conf. Advances in Cryptology*. London: Springer-Verlag, 1997, pp. 410–424.

[32] A. B. J. Teoh and C. T. Yuang, "Cancelable biometrics realization with multispace random projections," *IEEE Trans. Syst., Man, Cybern. B*, vol. 37, no. 5, pp. 1096–1106, 2007.

[33] A. B. Teoh, Y. W. Kuan, and S. Lee, "Cancellable biometrics and annotations on biohash," *Pattern Recogn.*, vol. 41, no. 6, pp. 2034–2044, 2008.

[34] T. Connie, A. Teoh, M. Goh, and D. Ngo, "Palmhashing: a novel approach for cancelable biometrics," *Inform. Process. Lett.*, vol. 93, no. 1, pp. 1–5, 2005.

[35] A. Teoh, D. N. C. Ling, and A. Goh, "Biohashing: Two factor authentication featuring fingerprint data and tokenised random number," *Pattern Recogn.*, vol. 37, no. 11, pp. 2245–2255, 2004.

[36] A. Kong, K.-H. Cheung, D. Zhang, M. Kamel, and J. You, "An analysis of biohashing and its variants," *Pattern Recogn.*, vol. 39, no. 7, pp. 1359–1368, 2006.

[37] L. Leng and J. Zhang, "Palmhash code vs. palmphasor code," *Neurocomputing*, vol. 108, pp. 1–12, 2013.

[38] L. Leng, A. B. J. Teoh, M. Li, and M. K. Khan, "Analysis of correlation of 2dPalmHash code and orientation range suitable for transposition," *Neurocomputing*, vol. 131, pp. 377–387, 2014.

[39] J. Zuo, N. Ratha, and J. Connell, "Cancelable iris biometric," in *Proc. Int. Conf. Pattern Recognition*, 2008, pp. 1–4.

[40] F. Farooq, R. Bolle, T.-Y. Jea, and N. Ratha, "Anonymous and revocable fingerprint recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, June 2007, pp. 1–7.

[41] T. Boult, "Robust distance measures for face-recognition supporting revocable biometric tokens," in *Proc. Int. Conf. Automatic Face and Gesture Recognition*, Apr. 2006, pp. 560–566.

[42] T. Boult, W. Scheirer, and R. Woodworth, "Revocable fingerprint biotokens: Accuracy and security analysis," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, June 2007, pp. 1–8.

[43] B. Yang, D. Hartung, K. Simoens, and C. Busch, "Dynamic random projection for biometric template protection," in *Proc. IEEE Int. Conf. Biometrics: Theory Applications and Systems*, Sept. 2010, pp. 1–7.

[44] P. Das, K. Karthik, and B. C. Garai, "A robust alignment-free fingerprint hashing algorithm based on minimum distance graphs," *Pattern Recogn.*, vol. 45, no. 9, pp. 3373–3388, 2012.

[45] S. Wang and J. Hu, "Design of alignment-free cancelable fingerprint templates via curtailed circular convolution," *Pattern Recogn.*, vol. 47, no. 3, pp. 1321–1329, 2014.

[46] N. Ratha, J. Connell, and R. Bolle, "An analysis of minutiae matching strength," in *Audio- and Video-Based Biometric Person Authentication* (Lecture Notes in Computer Science, vol. 2091), J. Bigun and F. Smeraldi, Eds. Berlin, Germany: Springer, 2001, pp. 223–228.

[47] A. Luong, M. Gerbush, B. Waters, and K. Grauman, "Reconstructing a fragmented face from a cryptographic identification protocol," in *Proc. IEEE Workshop on Applications of Computer Vision*, Jan. 2013, pp. 238–245.

[48] A. Adler, "Sample images can be independently restored from face recognition templates," in *Proc. IEEE Canadian Conf. Electrical and Computer Engineering*, May 2003, vol. 2, pp. 1163–1166.

[49] A. Nagar, K. Nandakumar, and A. K. Jain, "Biometric template transformation: A security analysis," in *Proc. SPIE, Electronic Imaging, Media Forensics and Security XII*, 2010, vol. 7541, pp. 75 410O–75 410O–15.

[50] F. Quan, S. Fei, C. Anni, and Z. Feifei, "Cracking cancelable fingerprint template of ratha," in *Proc. Int. Symp. Computer Science and Computational Technology*, Dec. 2008, vol. 2, pp. 572–575.

[51] S. Shin, M.-K. Lee, D. Moon, and K. Moon, "Dictionary attack on functional transform-based cancelable fingerprint templates," *ETRI J.*, vol. 31, no. 5, pp. 628–630, 2009.

[52] A. Lumini and L. Nanni, "An improved biohashing for human authentication," *Pattern Recogn.*, vol. 40, no. 3, pp. 1057–1065, 2007.

[53] Y. Lee, Y. Chung, and K. Moon, "Inverse operation and preimage attack on biohashing," in *Proc. IEEE Workshop on Computational Intelligence in Biometrics: Theory, Algorithms, and Applications*, Mar. 2009, pp. 92–97.

[54] P. Lacharme, E. Cherrier, and C. Rosenberger, "Reconstruction attack on biohashing," in *Proc. Int. Conf. Security and Cryptography*, 2013, pp. 1–8.

[55] R. Belguechi, E. Cherrier, and C. Rosenberger, "Texture based fingerprint biohashing: Attacks and robustness," in *Proc. IAPR Int. Conf. Biometrics*, Mar. 2012, pp. 196–201.

[56] X. Zhou and T. Kalker, "On the security of biohashing," in *Proc. SPIE, Electronic Imaging, Media Forensics and Security II*, 2010, vol. 7541, pp. 75 410Q–75 410Q–8.

[57] S. Li and A. Kot, "Fingerprint combination for privacy protection," *IEEE Trans. Inform. Forensics Security*, vol. 8, no. 2, pp. 350–360, Feb. 2013.

[58] A. Ross and A. Othman, "Visual cryptography for biometric privacy," *IEEE Trans. Inform. Forensics Security*, vol. 6, no. 1, pp. 70–81, Mar. 2011.

[59] A. Othman and A. Ross, "On mixing fingerprints," *IEEE Trans. Inform. Forensics Security*, vol. 8, no. 1, pp. 260–267, Jan. 2013.

[60] A. Othman and A. Ross, "Privacy of facial soft biometrics: Suppressing gender but retaining identity," in *Proc. European Conf. Computer Vision Workshops*, 2015, vol. 8926, pp. 682–696.

[61] K. Simoens, B. Yang, X. Zhou, F. Beato, C. Busch, E. Newton, and B. Preneel, "Criteria towards metrics for benchmarking template protection algorithms," in *Proc. IAPR Int. Conf. Biometrics*, Mar. 2012, pp. 498–505.

[62] N. Zhang, X. Yang, Y. Zang, X. Jia, and J. Tian, "Generating registration-free cancelable fingerprint templates based on minutia cylinder-code representation," in *Proc. IEEE Int. Conf. Biometrics: Theory, Applications and Systems*, Sept. 2013, pp. 1–6.

[63] H. Zhang, V. M. Patel, M. E. Fathy, and R. Chellappa, "Touch gesture-based active user authentication using dictionaries," in *Proc. IEEE Winter Conf. Applications of Computer Vision*, 2015, pp. 207–214.

[64] M. Derawi, C. Nickel, P. Bours, and C. Busch, "Unobtrusive user-authentication on mobile phones using biometric gait recognition," in *Proc. Int. Conf. Intelligent Information Hiding and Multimedia Signal Processing*, Oct. 2010, pp. 306–311.

[65] A. Primo, V. V. Phoha, R. Kumar, and A. Serwadda, "Context-aware active authentication using smartphone accelerometer measurements," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition Workshops*, 2014, pp. 98–105.

[66] S. Rane, "Standardization of biometric template protection," *IEEE MultiMedia*, vol. 21, no. 4, pp. 94–99, Oct. 2014.

[SP]

Mauro Barni, Giulia Droandi, and Riccardo Lazzeretti

Privacy Protection in Biometric-Based Recognition Systems

A marriage between cryptography and signal processing

©ISTOCKPHOTO.COM/GREYFEBRUARY

Biometrics Security and Privacy Protection

Systems employing biometric traits for people authentication and identification are witnessing growing popularity due to the unique and indissoluble link between any individual and his/her biometric characters. For this reason, biometric templates are increasingly used for border monitoring, access control, membership verification, and so on. When employed to replace passwords, biometrics have the added advantage that they do not need to be memorized and are relatively hard to steal. Nonetheless, unlike conventional security mechanisms such as passwords, biometric data are inherent parts of a person's body and cannot be replaced if they are compromised. Even worse, compromised biometric data can be used to have access to sensitive information and to impersonate the victim

Digital Object Identifier 10.1109/MSP.2015.2438131
Date of publication: 13 August 2015

1053-5888/15©2015IEEE

for malicious purposes. For the same reason, biometric leakage in a given system can seriously jeopardize the security of other systems based on the same biometrics. A further problem associated with the use of biometric traits is that, due to their uniqueness, the privacy of their owner is put at risk. Geographical position, movements, habits, and even personal beliefs can be tracked by observing when and where the biometric traits of an individual are used to identify him/her.

Processing biometric signals while they are encrypted provides a secure and elegant way to overcome the aforementioned problems [1], especially those related to privacy protection. Thanks to the opportunities offered by secure multiparty computation (SMPC) techniques [2], it is, in fact, possible to carry out the match between any two biometric templates by working only on encrypted data. Furthermore, it is also possible to design the underlying matching protocol in such a way that the final result of the match is known only to the intended party without leaking any information about the biometric templates or the identity of the biometric owner. The wide range of techniques allowing to process encrypted signals are usually known as signal processing in the encrypted domain (SPED).

As an example, let us consider a scenario in which a server is interested to know whether the owner of a biometric template is part of a list of enrolled individuals, e.g., the users who can access a certain service, or the criminals contained in a police record. The server has a database of plain biometric templates and the user submitting the query is interested to access the service without revealing his/her identity. Alternatively, the user submitting the query may be interested to know whether a biometric signal matches with one of the templates stored in the server—without that the server accesses the result of the query. According to the SPED paradigm, the aforementioned goals are achieved by letting the server comparing the templates in the database with the one provided by the user directly in the encrypted domain. While apparently impossible, a functionality like the aforementioned can be implemented by resorting to SMPC. It is known that virtually any computable function or algorithm can be evaluated by means of an SMPC protocol [3]. In the simplest cases, like those considered in this article, the protocol involves only two parties. In this case, we talk about secure two-party computation (STPC). In a general STPC setting, one party, say the client $\mathcal{C}$, owns a signal that must be processed in some way by the other party, hereafter referred to as the server $\mathcal{S}$. $\mathcal{S}$ must process $\mathcal{C}$'s signal without getting any information about it, in some cases not even the result of the computation. At the same time, $\mathcal{S}$ is interested to protect the information used to process the signal.

Two of the main approaches to SPED are homomorphic encryption (HE) [4] and garbled circuits (GCs) [5]. HE provides a way to evaluate linear operations on encrypted data, however when nonlinear operations are involved, it is necessary to resort to ad hoc, interactive, and usually complex protocols. On the other hand, GCs allow the evaluation of any function that can be represented with an acyclic boolean circuit. In some cases, however, the boolean circuit required to describe the functionality is so complex that it makes the use of GCs problematic. Given the complementary pros and cons of HE, oblivious transfer (OT), and GCs, the use of hybrid protocols has also been proposed to take advantage of the benefits offered by the two approaches [6]. Recently, fully HE (FHE) schemes [7] have been devised, allowing the evaluation of any function without any interaction between the involved parties. Unfortunately, FHE is still highly inefficient, principally due to the huge size of the public key.

Despite many recent advances and the introduction of more efficient cryptographic primitives, the complexity of SPED protocols is often high to prevent their use in practical applications. To reduce the complexity down to a manageable level, it is necessary that the underlying biometric processing algorithms and the STPC protocol are designed jointly by taking into account both the cryptographic and the signal processing facets of the problem. On the contrary, the most common approach used so far has been that of taking a classical biometric matching algorithm and transforming it into a protocol to be run in the encrypted domain. It is arguable that much better results can be obtained by developing a class of algorithms that are explicitly thought to ease a SPED implementation, e.g., by considering in advance which are the most complex operations to be carried out in a secure way and trying to avoid them.

In general, it is necessary that the biometric templates are represented through a vector of features of constant length and that a simple distance measure (e.g,. the Hamming or Euclidean distance) can be used to measure the degree of similarity between two vectors. If the previously mentioned conditions are satisfied, a biometric authentication or verification protocol can be developed easily by composing few blocks: distance computation, minimum selection, and comparison against a threshold [8], [9]. The search for efficiency is not limited to the choice of a suitable matching algorithm: representation issues must be considered as well. In the end, the complexity of SPED primitives depends on both the number of features the matching algorithm relies on and the number of bits used to represent them. By using fewer features and/or fewer bits, the complexity of the protocol decreases at the expense of matching accuracy. It is then necessary to find a proper configuration to couple efficiency and accuracy. Signal processing expertise can be exploited in several other ways: for example, it has been proven in [10] that using a common mask for iris recognition instead of a varying one dramatically simplifies the implementation of an iris-recognition system in the encrypted domain, with a very reduced impact on the performance of the system.

This article aims to illustrate the basics of STPC, including the way it can be applied to the protection of biometric templates, and to explain how the signal processing and cryptographic points of view can be considered together to obtain efficient, secure, and accurate SPED protocols. We also review some works in which such an approach has been used successfully for different biometric modalities, including fingerprint matching, iris recognition, and face recognition.

OVERVIEW OF BASIC SPED TOOLS

In this section, we provide a concise introduction to the basic primitives on which SPED technology relies. The tools presented here and the protocols described in the next sections are provably

secure in a semihonest setting [1], i.e., when the involved parties execute the protocol without deviating from it, but at the same time try to obtain as much information as possible about the other party's data. The choice of a semihonest model is due to the fact that while protocols providing security against a malicious party would be preferable, their implementation has a very high complexity. Moreover, at least in principle, protocols guaranteeing security in the semihonest model can always be modified to make them secure under more stringent threat models, even if such an increased security comes at the price of a higher complexity. Next we provide a qualitative description of various tools, focusing on their strengths and limitations.

HOMOMORPHIC ENCRYPTION

A cryptographic scheme (cryptosystem) is homomorphic [11] if an operation over encrypted data exists that correspond to another operation over the plain message. In other words, by indicating with $[\![x]\!]$ the encryption of a plain value x, we have $[\![x]\!] \boxtimes [\![y]\!] = [\![x \boxplus y]\!]$, for some operations $\boxtimes$ and $\boxplus$. Most HE schemes rely on asymmetric cryptography, and the homomorphic property holds under encryption with the public key of one of the parties involved in the protocol. Unless otherwise stated, in the following we assume that the private key is known only to the client $\mathcal{C}$, while the server $\mathcal{S}$ has access only to the public key.

The most common homomorphic cryptosystems (see, for instance, [12] and [13]) are additively homomorphic, i.e., $\boxtimes = \times$ and $\boxplus = +$. An additively homomorphic cryptosystem allows a party that does not know the decryption key to obtain the encryption of the sum between two values available to him only in encrypted form. In the same way, he can compute the encryption of the product between a known integer value c and a value available to him under encryption as $[\![cx]\!] = [\![x]\!]^c$. More complex operations can be implemented by resorting to an interactive protocol between $\mathcal{S}$ and $\mathcal{C}$.

Despite its elegance, the use of HE to compute with encrypted data comes at quite a high computational cost. In Paillier's cryptosystem, for instance, even plain values represented with few bits are encrypted in 2,048-bit-long ciphertexts (the plaintext after the encryption) so that sums and products between plain values are mapped respectively to products and exponentiations on very long ciphertexts. Nonlinear operations, such as products between encrypted values or comparisons, are even more complex and require interaction between the parties. For this reason, the communication complexity of an HE protocol depends on the number of transmitted ciphertexts, as well as on the number of communication rounds, while computation complexity is usually dominated by the number of exponentiations on encrypted values (the most expensive operation) required by the protocol. Multiplicative homomorphic cryptosystems exist as well [4], [14], allowing the evaluation of products between encrypted values ($\boxtimes = \times, \boxplus = \times$), but they have a lower practical utility with respect to additive HE.

Fully HE (FHE) schemes allow both the evaluation of additions and products in the encrypted domain. C. Gentry [7] developed the first secure somewhat HE (SHE) and FHE schemes, working on binary data. SHE allows the evaluation of a limited number of additions and multiplications, while FHE extends SHE to bypass such a restriction at the price of a huge increment of memory and computational complexity, thus making all FHE schemes proposed so far highly impractical.

By using Gentry's original SHE scheme and subsequent improvements, it is possible to evaluate binary circuits composed by up to a maximum number of XOR and AND gates directly on $\mathcal{S}$'s side without any interaction with $\mathcal{C}$, thus making protocols based on SHE very appealing for clients equipped with low-power devices. Efficient SHE solutions can be designed to evaluate circuits having a given (small) number of AND gates and then transformed into more expensive FHE solutions, if necessary. Luckily in most biometric recognition algorithms, the number of required operations is known in advance, making the use of protocols based on SHE possible.

A further simplification has been introduced in [15] where a SHE scheme operating on integer values has been proposed, thus allowing to encrypt each input directly, instead of decomposing it into bits and then using bitwise encryption. On the other hand, SHE (or FHE) schemes working on integers permit only the evaluation of polynomial functions (up to a certain degree for SHE).

OBLIVIOUS TRANSFER

An OT [16] is an STPC protocol that enables one party, say the server $\mathcal{S}$, to forward one out of n messages $(x_1, x_2, \ldots, x_n)$ to the client $\mathcal{C}$. $\mathcal{C}$ chooses the index i of the element that he would like to get. At the end of the protocol, the server gets no information on the index i, and the client does not get any information on the other x_js. The possibility to move great part of the computation to an offline phase, during which several OTs are evaluated on randomly chosen values, permits to greatly simplify the complexity of OT. The random values are replaced by the actual values during a much more efficient online phase [17]. Neglecting the offline complexity, and thanks to precomputation, the online communication of multiples one-out-of-two OTs is reduced to about 2ℓ bits for each OT, where ℓ is the message bit length, transmitted in parallel in two rounds. With regard to computational complexity, only simple XOR operations are required on both sides.

GARBLED CIRCUITS

The possibility of securely evaluating any binary circuits was proposed for the first time by Yao in his seminal paper [5]. Yao's protocol, the GC, involves both the parties in the computation and distributes the computation between $\mathcal{S}$ and $\mathcal{C}$. $\mathcal{S}$ encrypts (garbles) each gate of the circuit and maps each input bit into a random string. Then $\mathcal{S}$ sends the GC to $\mathcal{C}$ together with the secrets corresponding to $\mathcal{S}$'s inputs. The secrets associated to $\mathcal{C}$'s inputs are transmitted to $\mathcal{C}$s by means of OT. In the last phase of the protocol, $\mathcal{C}$ decrypts the gates by using the input secrets and obtains the final output of the circuit.

For a long time, GC were thought to be highly impractical. However, they have recently gained renewed popularity, thanks to several efficiency improvements (most of which summarized in [18]). The protocol associates a secret of 80 bits to each bit involved in the computation, making single-core operations lighter than in HE (we recall that a Paillier ciphertext is 2,048 bits long). Unluckily, even if most of the computation is performed on $\mathcal{S}$'s side, $\mathcal{C}$

must also take an active part in the protocol. The computational complexity depends linearly on the number of nonXOR gates composing the circuit (which in turn depends on the input bit lengths), in fact, XOR gates can be evaluated with negligible computational and communicational complexity. It is important to underline that a GC protocol requires only two rounds, regardless of the circuit size and the number of input bits (an additional round is necessary if the final result must be sent to $\mathcal{S}$). We also point out that circuit garbling does not depend on the actual inputs and, in some particular scenarios where the functionality to evaluate is known in advance, circuit encryption and transmission can be precomputed.

Given that the complexity depends on the number of gates composing the circuit, GCs are suited for operations such as sums and comparisons, for which the number of gates depends linearly on the input bit length. On the contrary, GCs are less efficient when the number of gates grows more than linearly with the input bit length. This is the case, for instance, of products and divisions for which the circuit size depends quadratically on the bit length of the inputs.

HYBRID PROTOCOLS

Sometimes complex protocols can be divided into subprotocols, and different tools can be used for their implementation to take the best from each approach. Such an idea has been applied to develop hybrid protocols working with HE and GCs in [6], but can also be extended to different tools. Hybrid protocols require the adoption of proper interfacing protocols to link subparts implemented by relying on different technologies. For instance, it may happen that an intermediate value x output by an HE protocol must be used as input in a GC subroutine, or vice versa. In this case, the different parts of the protocol must be connected in such a way that the security of the whole system is guaranteed. At the same time, the representation of the variable x must be adapted to the subprotocol requirements.

BIOMETRIC RECOGNITION PROTOCOLS

Biometric recognition protocols can be divided in two main categories: in the first scenario, usually referred to as *authentication*, the user is interested in demonstrating that he is who he claims to be, while in the second one, called *identification*, the goal of the protocol is to determine the identity of the user submitting the biometric template. To better protect the users' privacy, in some cases, SPED-based identification protocols simply verify whether the user is enrolled in the database or not. The server $\mathcal{S}$ owns a database of enrolled biometric feature vectors $(\{Y_i\}, i = 1, \ldots, n)$ and the client C owns a biometric vector X. In all cases, $\mathcal{S}$ and C are interested in protecting the privacy of their data.

In the authentication problem [Figure 1(a)], C submits a new instance of his biometrics. The fresh biometric template is processed to extract a feature vector X that is sent to $\mathcal{S}$ together with an identifier, used by $\mathcal{S}$ to select the corresponding enrolled template Y_{id} in the database. The distance $d(X, Y_{\text{id}})$ between the query X and the template Y_{id} is evaluated and the result is compared against an acceptance threshold.

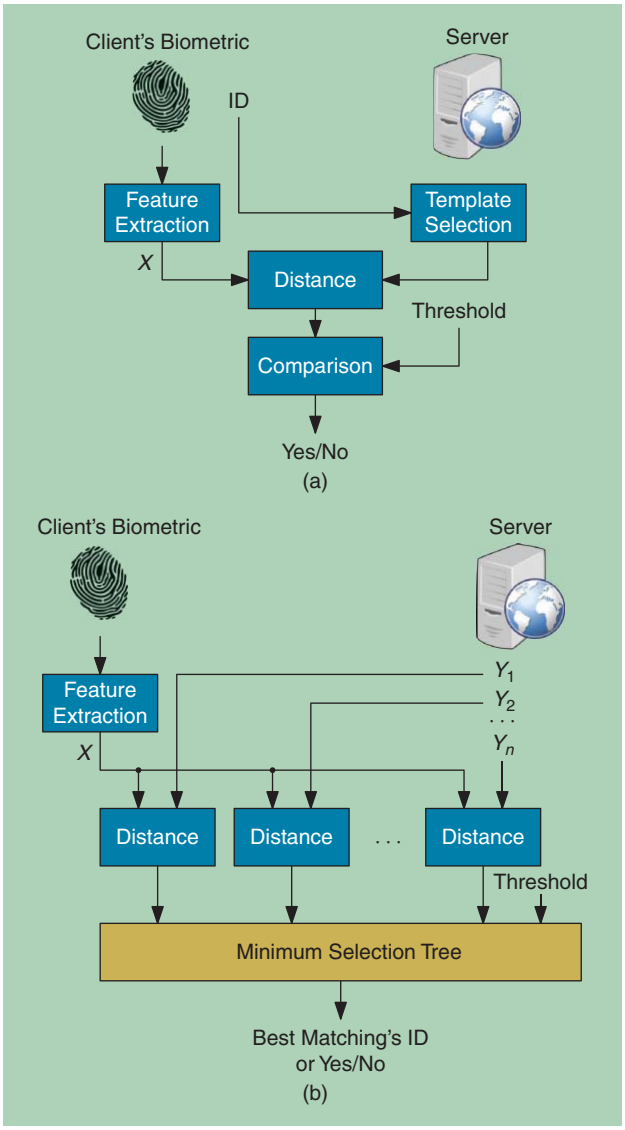

[FIG1] **Biometric recognition protocols: (a) authentication and (b) identification.**

In the identification scenario [Figure 1(b)], the client extracts the feature vector X from the fresh biometric template and submits it to the server without revealing his identity. The server must verify whether an index i exists such that $d(X, Y_i) < \varepsilon$. To do so, C and $\mathcal{S}$ first evaluate $d_i = d(X, Y_i)$ for all $i = 1 \ldots n$, then they find the minimum among all d_i and the threshold through a minimum selection tree returning "yes" if the minimum distance is below the threshold, and "no" otherwise. It is also possible to modify the minimum tree so that the output is a user's identification index instead of a yes/no answer.

As can be seen, a general recognition protocol is composed of a few basic blocks: feature extraction, distance computation, comparison, and minimum selection. Feature extraction involves only data provided by one party, and for this reason it is usually implemented in the plain domain. On the other hand, distance computation, comparison, and minimum selection involve private data owned by C and $\mathcal{S}$, and so they must be implemented by resorting to SPED.

[TABLE 1] COMPUTATIONAL AND COMMUNICATION COMPLEXITY OF PRIVACY-PRESERVING FACE RECOGNITION [19].

DATABASE SIZE	COMPUTATIONAL COMPLEXITY (SECONDS)			COMMUNICATION COMPLEXITY (KILOBYTES)	
n	FULL QUERY	WITH PRECOMPUTATION	PUBLIC EIGENFACES	FULL QUERY	PUBLIC EIGENFACES
10	24	8.5	1.6	2,725	149
200	34.2	14.5	11.4	5,497	2,921
320	40	18	18.2	7,249	4,674

There are many possibilities to implement these blocks in a privacy-preserving way. The choice depends on many factors, such as device configuration, network bandwidth, and latency, computational capabilities of $\mathcal{S}$ and $\mathcal{C}$. In this section, we provide a brief description of how the various blocks can be implemented, leaving a more detailed description to the next sections.

The Hamming and the squared Euclidean distances are the most commonly used distances because they can be easily implemented in a SPED setting. The Hamming distance is used whenever the biometric template corresponds to a binary vector, while the squared Euclidean distance is used on integer biometric vectors (the squared version is used to avoid the expensive computation of the square root). Both distances can been implemented by using GC, HE, or OT. In [8, Ch. 7] the authors show that, due to its binary nature, the Hamming distance can be efficiently implemented by using GC, while an HE implementation is preferable for the squared Euclidean distance [19], since HE allows an efficient computation of products. An efficient OT implementation of both Hamming and squared Euclidean Distance has been proposed in [20]. It is also possible to implement such distances through SHE [21], while, given the limited number of operations required in both cases, resorting to FHE is not necessary.

A comparison is needed to verify whether a certain distance is lower than the acceptance threshold (squared threshold if the squared Euclidean distance is used). Its implementation [8, Ch. 7] requires that the involved quantities are represented in binary form, thus making GC-based implementations more attractive. Implementations based on HE [19] have also been proposed, but they require several interactions between the parties.

Starting from a comparison protocol, it is possible to evaluate the minimum among two encrypted values by using the output of the comparison to select between two numbers x and y in a multiplexer. Given the necessity of a comparison operator, a GC implementation is usually preferable. The protocol for the selection of the minimum between two numbers can be easily extended to the computation of the minimum among n values using a reverse tree implementation [8, Ch. 7] where each node computes the minimum between the results of the previous left and right subtrees. The minimum selection tree can be modified to output the minimum value or the corresponding identifier.

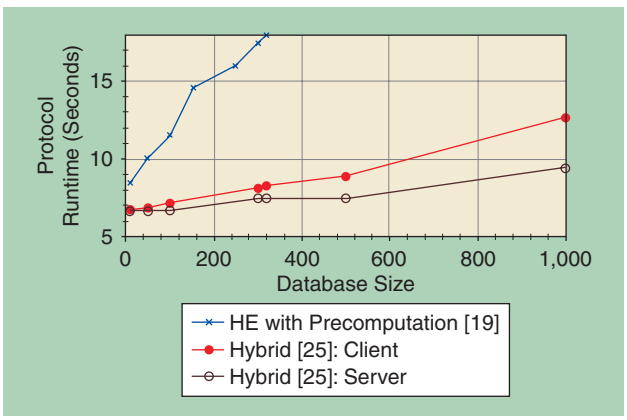

[FIG2] The runtime comparison of HE [19] and hybrid [25] implementations of the eigenface protocol.

OPTIMIZATION OF SPED PROTOCOLS THROUGH CRYPTOGRAPHIC PRIMITIVE SELECTION

In this section, we provide an overview of how the use of different cryptographic primitives can be exploited to improve the performance of biometric recognition protocols. For the sake of simplicity, we do not discuss the improvements in the implementation of the basic cryptographic primitives and we leave the description of signal processing optimizations to the next section.

One of the first papers addressing privacy-preserving biometric authentication is [22]. The protocol does not focus on a specific biometric modality, but rather on a general biometric representation consisting of a binary string. It then presents a secure implementation of the Hamming distance computation based on private information retrieval.

An implementation of privacy-preserving biometric identification protocols operating in the semihonest setting, implemented according to the overall scheme presented in the previous section, has been proposed by Erkin et al. in [19]. The recognition protocol is based on eigenfaces [23], it achieves 96% correct classification averaged over different lightning conditions, 85% when different face orientations are considered, and 64% when face size varies as well. In contrast to most SPED biometric-recognition protocols, the feature extraction step is carried out in the encrypted domain by relying on the homomorphic properties of the Paillier cryptosystem [12]. Squared Euclidean distance computation is also implemented by relying on the Paillier system, while the comparison protocol is implemented according to the scheme proposed by Damgard et al. in [13]. The protocol complexity was evaluated by running it on a computer with a 2.4-GHz dual-core processor, and using the "ORL Database of Faces" [24] obtaining a runtime of about 40 seconds for a single match. The runtime could be reduced to 18 seconds by resorting to precomputation. As shown in Table 1, the authors have demonstrated that it is possible to further reduce the computational and communication complexity by assuming that the parameters of the eigenface extraction protocol are public (such an assumption has been adopt by virtually all subsequent works on the same topic).

Erkin et al. protocol has been improved by Sadeghi et al. [25], who proposed a full-GC and a hybrid protocol for eigenface biometric recognition, where HE is used to compute the distance and GC for the comparison. As shown in Figure 2, the resulting protocol is 30% faster than [19], when implemented on a PC having a 2.6-GHz processor.

In [26], the authors propose a new technique for template extraction called *SCiFI*. The protocol evaluates distances between faces by using Paillier HE and then implements the comparison by using a one-out-of-d OT, where d is the maximum value that the distance can assume. The experiments were performed on two computers with a 2.6-GHz processor and a 2.8-GHz dual-core processor, respectively. The online time complexity is about 0.30 seconds for a single match.

Moving from face recognition to iris-based systems, Luo et al. [27] implemented an HE-based privacy-preserving iris identification protocol based on IrisCode [28] and tested it on the CASIA Iris database [29], containing 100 IrisCodes of 9,600 bits each. The resulting protocol needs 27.1 minutes on average for a single query on a computer equipped with a 2.4-GHz processor. Such a large complexity is justified by the very large bit length of IrisCodes (9,600-bit), which are bitwise encrypted by means of the Paillier cryptosystem. A different approach is presented in [30], where the authors use a hybrid (HE and GC) protocol for biometric identification and optimize it by precomputing most of the operations. Further improvements are obtained by optimizing the multiplication protocols and by using the DGK scheme [13] for comparison computation. A *C* implementation of the protocol has been tested on a 2.13-GHz dual-core processor obtaining results about 25% faster with respect to the same protocol implemented by using HE. Online computation times are summarized in Table 2. In particular, the comparison between two encrypted 2,048-bit IrisCodes requires only 0.15 seconds.

In [31] and [10], the authors present an iris identification protocol based on two different full-GC implementations (more details are given in the next section). In [31], the authors run a Java implementation of the protocol on a client with a 2.66-GHz quad-core processor connected through a local area network with a server equipped with a 2-GHz processor. They tested the protocol on databases of different sizes n obtaining a total bandwidth of $475n + 0.08n^2$ kilobytes and a runtime of about 2.4 seconds for each match.

The protocol described in [10] has been implemented in Java and run on a machine mounting a 3.00-GHz processor over IrisCodes of the CASIA Iris database [29] represented with 9600 and 2,048 bits. Thanks to offline computation of the circuit garbling phase and circuit transmission, the matching between two IrisCodes represented with 2,048 bits needs 0.56 seconds and the transmission of 571 kilobytes, while the matching between two IrisCodes represented with 9,600 bits needs 2.5 seconds and the transmission of 2,655 kilobytes.

We conclude this section by considering fingerprint matching. Given the necessity of working with finite-length feature vectors, most schemes proposed so far rely on the fingercode representation of fingerprints [32]. This is the case of the system proposed by Barni et al. [33], [34] implementing a Paillier-based identification protocol. The execution of the protocol on a database with 64 identities takes about 16 seconds on a PC equipped with a 2.4-GHz dual-core processor. Fingerprint identification is also addressed in [30], where protocols similar to those used for iris recognition are used. With respect to [34], the implementation based on fingercode is 35 times faster (client online runtime is 0.35 seconds while server's one is 0.45 seconds). The protocol has been also adapted to operate on minutiae [35] (results in [32] reports an FAR lower than 1%), but runtimes increase significantly. Table 2 shows the performance of the protocol when 32 minutiae are used to represent the fingerprint. Yet another hybrid implementation is described in [36] for fingercode-based identification. Table 3 shows the online computation time obtained with a Java implementation running on two machines equipped with a 2.0-GHz processor.

A somewhat different approach, relying on a different use of the available cryptographic primitives, has been proposed by Bringer et al. [20]. The new approach, called *GSHADE*, is based on a hybrid use of OT and GMW [37]. GMW is an SMPC primitive similar to Yao's GCs. It implements the to-be-computed functionality as a binary circuit; however, it performs the secure evaluation by relying on shares rather than encrypted gates. GSHADE has been tested by running a C++ implementation on two computers with 3.2-GHz precessor. By considering a database of 320 IrisCodes of 2,048 bits each, the communication complexity of GSHADE is around three times larger than that of the hybrid protocol described in [30]. However, the GSHADE protocol is 35 times faster than the system presented in [30]. Similar results have been obtained with fingercodes (runtime improves by a factor 500 with respect to [36]) and eigenfaces (with a runtime improvement of a factor ranging from 66 to 100 with respect to [19]).

With the increased popularity of FHE and SHE schemes, a few completely noninteractive solutions for privacy-preserving biometric recognition have been proposed. In [21], the first noninteractive biometric authentication protocol, based on an integer extension of the SHE scheme described in [38], is presented. All the computation is moved on the server's side, leaving only the encryption of the inputs and the decryption of the result to the

[TABLE 2] THE ONLINE PERFORMANCES OF IRISCODE-, FINGERCODE-, AND MINUTIAE-BASED FINGERPRINT IDENTIFICATION [30]. SOME OF THE OVERHEADS DEPEND ON THE SERVER'S DATABASE SIZE, IN WHICH CASE THE COMPUTATION ARE INDICATED PER RECORD ("/REC").

	SERVER RUNTIME	CLIENT RUNTIME	BANDWIDTH
IRISCODE	89 MILLISECONDS + 149.25 MILLISECONDS/REC	0 MILLISECONDS + 22.61 MILLISECONDS/REC	0.5 KILOBYTES + 19.9 KILOBYTES/REC
FINGERCODE	0.22 MILLISECONDS + 1.42 MILLISECONDS/REC	4.7 MILLISECONDS + 1.08 MILLISECONDS/REC	2.12 KILOBYTES + 0.86 KILOBYTES/REC
MINUTIAE	6 MILLISECONDS + 339 MILLISECONDS/REC	25 MILLISECONDS +1,876 MILLISECONDS/REC	16 KILOBYTES + 294 KILOBYTES/REC

[TABLE 3] ONLINE PERFORMANCES OF THE FINGERCODE IDENTIFICATION PRESENTED IN [36].

DATABASE SIZE	RUNNING TIME (SECONDS)	BANDWIDTH (KILOBYTES)
128	2.22	966.84
256	4.33	1,927.71
512	9.12	3,849.48
1,024	18.11	7,692.98

client. With regard to complexity, a $C++$ implementation of the protocol has been run on a machine mounting a 3.30-GHz processor. With respect to an equivalent implementation based on the Pailler cryptosystem, the computational complexity is considerably reduced (59 seconds for Troncoso et al. implementation versus the 420 seconds of an equivalent Paillier implementation), with the additional advantage of avoiding the interaction between the parties. On the other hand, due the larger expansion factor of a lattice-based cryptosystem like [38], the communication complexity is larger than the Paillier-based version: 393 megabytes in [21] and 16.4 megabytes for the Paillier-based version. Another authentication protocol based on SHE has been proposed in [39]. Thanks to a packed representation of the biometric templates, the protocol is able to compute the Hamming distance with only three products. Tests performed on a 3.07-GHz processor show that only 18.10 milliseconds are necessary for distance computation, which is not only faster than the SHE-based implementation of [21], but also faster than the Hamming distance computational time of SCiFI (310 milliseconds) [26] and [30] (150 milliseconds). In both [21] and [39], only the distance is computed by means of SHE operating on integers. Such schemes permit only the computation of polynomial functions of the inputs, and they cannot be used for comparisons. For this reason, in [21] and [39] the final comparison is carried out in plain domain by the client.

For completeness, we highlight that beyond papers strictly focusing on biometric recognition, other interesting privacy-preserving applications that can be also applied to biometric protocols have been developed. For example, [40] presents a new scheme for a privacy-preserving evaluation of a sample set similarity (EsPRESSo) that can be used for iris matching, while in [41] the authors address privacy-aware media classification, and also face recognition, on public databases.

SIGNAL PROCESSING OPTIMIZATION

Even if the development of more and more efficient cryptographic primitives and their adaptation to the specific needs of biometric-recognition protocols, has led to considerable complexity reduction, further ways to reduce the complexity of SPED protocols are needed to match the requirements set by practical applications. A less explored, but promising, strategy is the optimization of the signal processing aspects of the algorithms to be implemented in a SPED fashion. Generally speaking, signal processing optimization can be carried out at three different levels: 1) the algorithmic level, 2) the feature choice and distance selection, and 3) the feature representation level. (While this classification is quite general, in some cases the various levels cannot be clearly identified and optimizations operating at different levels may depend on each other in a complex way.) In the first case, the matching algorithm is designed in such a way to avoid the operations that most complicate a SPED implementation. As an example, when considering an HE-based implementation, algorithm designers should try to minimize the use of nonlinear operations. With regard to feature and distance selection, it is desirable that the computation of distances between feature vectors can be easily implemented by means of the available STPC primitives. In identification scenarios, the number of distances to be computed grows linearly with the size of the database [31], calling for a careful design of this part of the protocol. The last optimization level concerns the size of the feature vector and the number of bits used to represent the feature values. Both aspects have a great impact on protocol efficiency. Investigating the relationship between the size of the feature vector and the number of bits used to represent it on one side and the accuracy of the matching process on the other side may lead to a significant simplification of the resulting protocol. Of course, all of the above considerations are not independent from the STPC primitives on which the protocol relies. Hence the preferable tool for each algorithm configuration must be selected among all the available SPED tools. As shown in the previous section, this is often a hard choice depending on many factors such as the bandwidth and the latency of the network, the characteristics of the devices available at the client and server side, etc.

In the following, the various optimization levels are described in more detail. For each level, we provide one or more practical examples of its use in a biometric-matching protocol.

[TABLE 4] FALSE REJECTION RATES (FRRs) OF [31] ACCORDING TO THE NUMBER OF BIOMETRIC TEMPLATES SELECTED IN THE FILTERING PHASE AMONG THE 2,710 ELEMENTS IN THE DATABASE.

$k = 1$	$k = 10$	$k = 20$	NO FILTER
19.5%	8.2%	6.1%	3.1%

ALGORITHM-LEVEL OPTIMIZATION

Given a matching algorithm, some optimizations can be applied to improve its performance, trying to avoid the operations that are most expensive when implemented in a SPED setting.

In identification protocols, the complexity mainly depends on the number of biometric templates contained in the database, since this directly affects the number of matches that must be computed. In the iris-recognition protocol presented in [31], the matching between two IrisCodes is based on a normalized Hamming distance involving two iris masks (one for each iris template) that are used to remove the noninformative parts of the iris code, usually those impaired by reflexes, eyelashes, and shades. Given the binary nature of the IrisCode, a GC solution is very efficient with regard to Hamming distance computation, but the use of the two masks involves two nonfree AND gates for each bit, approximately tripling the complexity of the modified Hamming distance circuit. The idea put forward in [31] is to reduce the database size through a filtering phase during which only the most promising templates are selected. The nonmasked

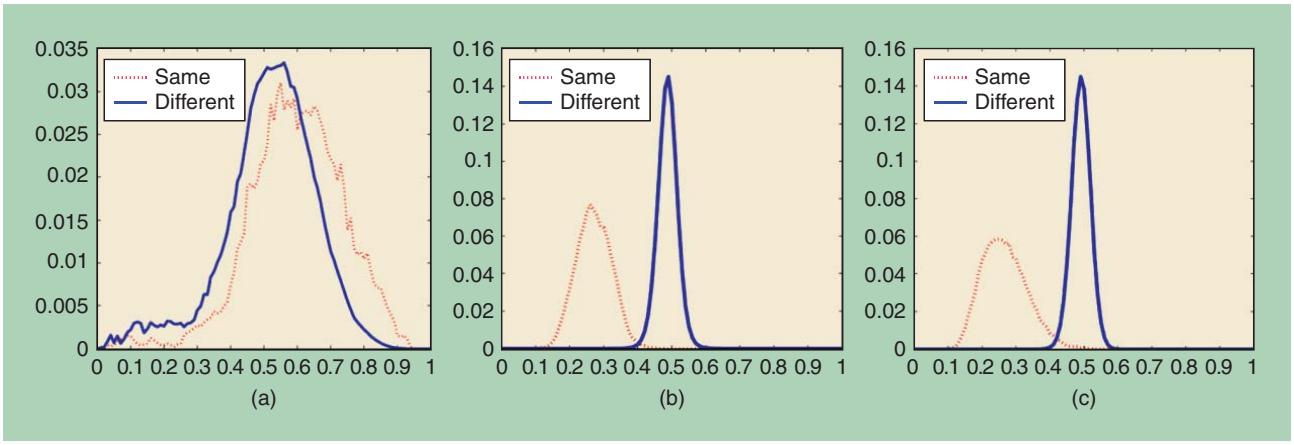

[FIG3] Distributions in IrisCode identification in [10] over IrisCodes in the CASIA Iris database [29]: (a) mask overlap sizes, (b) real masks, and (c) a common mask.

Hamming distance is evaluated on a subset of 128 bits, whose position is chosen between the usually unmasked bits, selected in the query and all the n templates in the database. Then the randomized indexes of the k templates with the smallest distances are passed to the client. After the filtering phase, $\mathcal{C}$ and $\mathcal{S}$ run an identification protocol where the masks are used to refine the distance computation and the input secrets of the k templates and masks are retrieved by $\mathcal{C}$ through an OT protocol. Thanks to the above solution, the complexity of the protocol is significantly reduced: with $k \approx n/10$ a total bandwidth of $475n + 0.08n^2$ kilobytes is reported, which is considerably lower than the $3.6n$ megabytes needed for an exhaustive comparison. On the negative side, the protocol does not guarantee that the correct biometrics are selected for the second phase, hence decreasing the accuracy of the identification. Table 4 shows the FRRs with different values of k and without filtering.

A different algorithmic optimization for iris-based identification has been proposed in [10]. It relies on the use of a common mask, estimated from all the masks associated to the IrisCodes in the database. Given a data set, the distribution of the mask overlap regions is computed. Figure 3(a) shows that masks from the same individuals have larger overlap than those from different individuals, concluding that among all masks, those of each individual have larger intercorrelation. On the other hand, as shown in Figure 3(a), masks also belonging to different individuals are quite similar. By relying on this observation, the authors proposed to simplify the circuit implementing the masked distance by using a common mask for all the IrisCodes. The common mask is set to "1" at all bit positions, where the percentage of the prealigned masks equal to "1" at those positions exceeds an empirically determined threshold λ. The common masks do not reveal information about the single templates in the database and can be publicly disclosed. Figure 3(b) and (c) shows the distribution of the distance when using individual masks and a common mask, respectively. By using a common mask, built by setting $\lambda = 0.8$, the overlap between the two distributions increases. Anyway, the best result with individual masks are obtained by using a similarity threshold ε between the iris templates equal to 0.41, providing an FAR equal to 0.53% and an FRR equal to 0.54%. By using a common mask, the best FAR and FRR are 1.44% and 1.47%, respectively, obtained with $\varepsilon = 0.43$, resulting in an accuracy loss lower than 1%. The protocol has been tested on two different data sets, one containing IrisCodes represented with 2,048 bits and the other containing IrisCodes represented with 9,600 bits. By using a common mask, a speedup factor of up to 8.7 can be achieved in the first data set and a speedup factor of up to 4.7 in the second one. In both cases the bandwidth is reduced by a factor ~4.3. As reported in the original paper [10], the online time for an iris match is 65 milliseconds and requires the transmission of 133.7 kilobytes.

Another example of algorithmic optimization has been proposed in the SHE-based face recognition protocol described in [21]. The authors use a Gabor filter (a linear filter used for edge detection) to build the feature vector. To minimize the amount of data to be processed, they discard the phase information and use a novel statistical characterization to model the magnitude of Gabor coefficients. Moreover, coefficient representation does not rely on quantization as usual but is obtained by dividing the probability density function into 2^ℓ numbered sections. A coefficient is represented through the index of the segment to which it belongs. The authors compared the performance of such an indexing procedure with classical quantization-based schemes while varying the coefficient bit length. Experiments were run on several databases. Results obtained on the XM2VTS data set [42] show that 4 bits are sufficient to produce a much better fit, equaling the original performance of [42] (~96%) when using a support vector machine (SVM) implemented as a weighted distance, while the accuracy decreases by ~3% if no SVM is used. On the other hand, the server runtime increases from 59 to 120 seconds when an SVM is used.

FEATURE AND DISTANCE CHOICE

The choice of the features used to represent the biometric templates has a major impact on the complexity of SPED biometric matching protocols, due to the strict correlation between the type of features used to represent the biometric signals and the distance function used to evaluate the match. Let us consider,

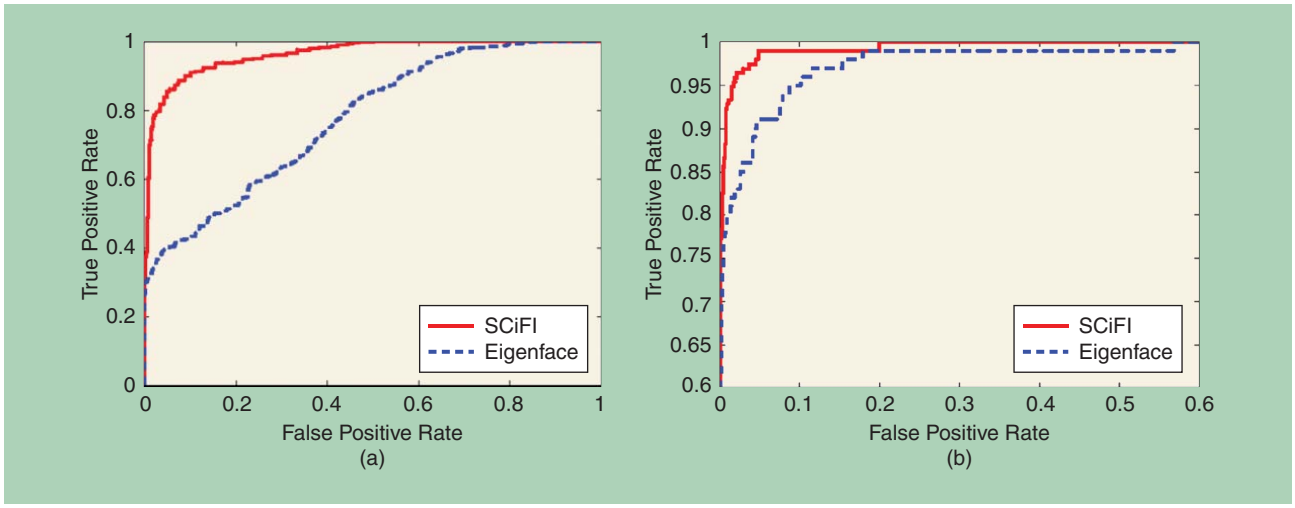

[FIG4] The robustness of SCiFI protocol [26] compared to eigenface [19]. Tests performed on the ORL "Database of Faces" from AT&T Laboratories Cambridge [24]. (a) A large illumination variation. (b) Near-frontal changes in pose, mild facial expressions, and mild illumination changes.

for example, fingerprint matching. The most popular and efficient matching algorithms are based on minutiae. However, in [33] and [34] the authors chose the fingercode representation. Even if the experiments show that filter-based matchers such as the fingercode tend to perform slightly worse than state-of-the-art minutiae-based matchers, the fingercode matching function has a much lower computational complexity and is more suitable for being implemented in an STPC setting. On the contrary, a privacy-preserving protocol operating on minutiae would be difficult to implement, mainly due to the variable length of the feature vector and the lack of a simple distance measure between minutiae features. The intuition of [33] and [34] was later validated in [30], wherein a hybrid implementation of both fingercode and minutia based identification protocols is described. As shown in Table 2, the runtime of the protocol based on minutiae is 100 times higher than that of the fingercode protocol.

Another example of protocol simplification through feature selection is the SCiFI protocol for face recognition [26]. The representation used by SCiFI is based on the idea of composing a face as a collection of fragments taken from a dictionary of facial features. The resulting feature vector consists of two parts: the first part with the indexes of the dictionary fragments that better represent the face, the second one with the position of each part with respect to the face center. The feature vector is then represented as a fixed length binary vector and matching is carried out by relying on the Hamming distance. Authors compared SCiFI with eigenface-based recognition [19] by evaluating its robustness to various factors such as large illumination variation and near-frontal changes in pose, mild facial expressions, and mild illumination changes. The results shown in Figure 4, where the recognition rate is plotted as a function of the false positive rate, demonstrate that is possible to improve the accuracy of the face recognition protocol, while, thanks to extensive precomputation, the online execution time required for the match of a query and a face in the database is reduced to about 0.31 seconds.

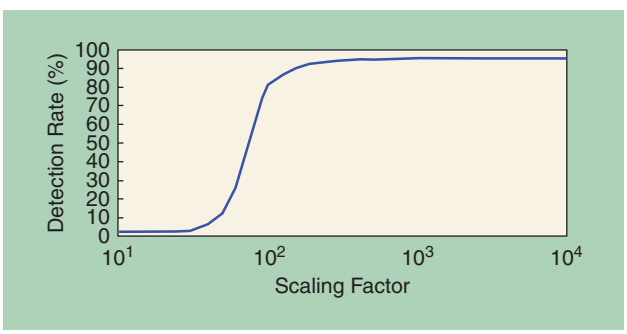

[FIG5] The correct detection rate versus representation accuracy in the face-recognition system described in [19].

FEATURE VECTOR SIZE AND REPRESENTATION ACCURACY

A further simplification can be obtained by decreasing the number of features used to represent the biometric template and the number of bits used to represent each feature. One example of such an approach is the HE face-recognition protocol proposed by Erkin et al. [19]. The signal processing analysis is limited to the definition of the scaling factor used to quantize the parameters of the protocol (which in turns determines the number of bits used to represent the parameters and hence the accuracy of the representation) and the number k of features used to represent a face. The authors aimed to obtain the same classification accuracy provided by a standard plain implementation—a correct recognition rate equal to 96%. As shown in Figure 5, such a goal is reached with a scaling factor ~1,000. Moreover, experiments proved that no improvement is observed by using $k > 12$. By relying on such an analysis, the authors show that matching a face image against a database of 320 biometrics takes roughly 40 seconds and requires the transmission of 7,249 kilobytes (see Table 1).

A more accurate signal processing analysis has been performed in the fingerprint recognition protocol described in [33]. Considering that a protocol computing the squared Euclidean

[TABLE 5] THE CONFIGURATION FOR FEATURE SIZE REDUCTION IN FINGERCODE PROTOCOL [33].

CONFIGURATION	FEATURES
A	640
B	384
C	192
D	96
E	48
F	32
G	16
H	8

[TABLE 6] THE PERFORMANCE OF PRIVACY-PRESERVING FINGERCODE PROTOCOL [33].

CONFIG.	FEATURE BIT LENGTH	EER	BANDWIDTH (BITS) 408 ENTRIES	RUNTIME (SECONDS) 100 ENTRIES
C	2	0.0715	6,902,008	44.43
	4	0.0673	8,135,800	53.66
D	2	0.0758	6,568,792	37.43
	4	0.0732	7,802,584	45.58

distances on 640 features would have a very high complexity, the authors checked if a lower number of features can be used without degrading significantly the matching accuracy and selected the minimum number of bits necessary to represent each feature. To this purpose, the matching algorithm was tested by using eight different fingercode configurations (Table 5) and varying the feature bit length between 1 and 8. Figure 6 shows the behavior of the equal error rate (EER) on the test set. As highlighted in the figure, it is evident that the accuracy of the system does not improve significantly when more than 96 features, each represented with 2 bits, are used. At the same time, the EER increases when only 1 bit is used for the representation, thus impeding the use of a more efficient protocol based on the Hamming distance. By the light of the above considerations, the authors chose to focus on configurations C and D, with 2 or 4 bits for feature representation. The results obtained in [33] are reported in Table 6. Moving from 192 features to 96 features and halving the number of bits, we observe a significant simplification of the protocol, with only a minor decrease of matching accuracy.

To improve the efficiency of a protocol, it is also possible to work on the representation of intermediate values. For example in the HE and GC hybrid protocols described in [36], the authors modify the protocol to use a more compact representation of intermediate distances. They assume that the acceptance threshold and its bit length κ are publicly known. After computing a distance by means of an HE protocol, they start the GC section by checking if the distance is greater than 2^{κ}. In this case, the distance value is replaced with the threshold. In such a way the minimum selection circuit can operate on shorter values hence reducing the total number of gates (results are given in Table 3).

CONCLUSIONS

As shown throughout this article, processing biometric signals in the encrypted domain provides an elegant and provably secure

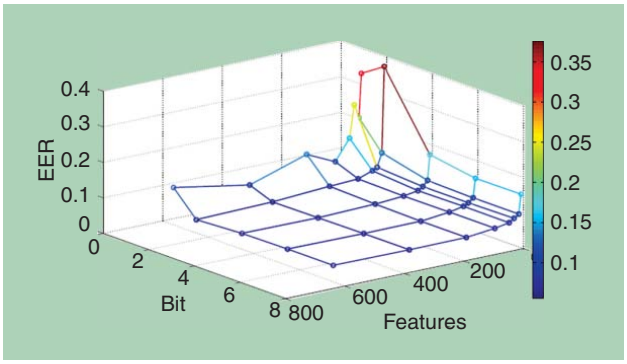

[FIG6] The EER of the different configurations of fingercode [33] on the fingerprint database [43].

mechanism to protect both the biometric data and the privacy of the individuals subject to biometric controls. Thanks to the use of STPC cryptographic primitives, biometric matching algorithms can be implemented in such a way that the parties involved in the matching do not get access to either the data owned by the other party or the result of the match. From a decade of research in the field, it is now well evident that the question is not whether a certain computation can be carried out in the encrypted domain, but whether such a computation can be carried out efficiently.

While the quest for efficiency has driven the agenda of researchers in the last years, research has been mainly focused on the development of more efficient STPC primitives and their use to implement conventional biometric matching algorithms in a SPED framework. We believe, though, that significant advantages can also be obtained by working at the signal processing level or, even better, by jointly considering the cryptographic and signal processing facets of the problem. It was the goal of this article to introduce the readers to the main concepts behind SPED biometric matching and to show how a clever design of the underlying matching protocol may help to fill the gap between the complexity of SPED protocols and the efficiency required for the deployment of such protocols in real systems. We hope that the readers appreciate our effort and will contribute to the future advancement of this exciting field.

AUTHORS

Mauro Barni (barni@dii.unisi.it) received the M.S. degree in electronics engineering in 1991 and the Ph.D. degree in information and communication engineering in 1995, both from the University of Florence, Italy. He is currently with the Department of Information Engineering and Mathematics of the University of Siena, Italy. His research activity focuses on multimedia and information security, with particular reference to copyright protection, multimedia forensics, and signal processing in the encrypted domain. He has coauthored almost 300 papers published in international journals and conference proceedings in addition to being a coauthor of *Watermarking Systems Engineering: Enabling Digital Assets Security and Other Applications* (CRC Press). From 2010 to 2011, he was the chair

of the IEEE Information Forensics and Security Technical Committee of the IEEE Signal Processing Society (SPS). He was appointed a Distinguished Lecturer by the SPS for 2013–2014. He is the editor-in-chief of *IEEE Transactions on Information Forensics and Security* and is a Fellow of the IEEE and a senior member of EURASIP.

Giulia Droandi (droandi@student.unisi.it) received the master's degree (cum laude) in mathematics from the University of Siena, Italy, in 2011, with a thesis on enumerative combinatory. She has been a Ph.D. student in the Department of Information Engineering and mathematics of the same university since 2012. Her research focuses on fully and somewhat homomorphic encryption and its possible applications to biometrics.

Riccardo Lazzeretti (lazzeretti@diism.unisi.it) graduated with a degree in computer science engineering from the University of Siena, Italy, in 2007, where he continued his studies as a Ph.D. student under the supervision of Prof. Mauro Barni in the Information Engineering Department. From November 2009 to May 2010, he was with Philips Lab in Eindhoven, The Netherlands. In 2012, he received a research grant and continued his research in the Information Engineering and Mathematics Department of the University of Siena. His research activity is mainly focused on privacy-preserving applications based on secure two-party computation tools.

REFERENCES

[1] R. Lagendijk, Z. Erkin, and M. Barni, "Encrypted signal processing for privacy protection: Conveying the utility of homomorphic encryption and multiparty computation," *IEEE Signal Processing Mag.*, vol. 30, no. 1, pp. 82–105, 2013.

[2] O. Goldreich, "Secure multi-party computation," manuscript, 1998. [Online]. Available: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.11.2201&rep=rep1&type=pdf

[3] O. Goldreich, S. Micali, and A. Wigderson, "How to play any mental game," in *Proc. 19th Annu. ACM Symp. Theory of Computing*, 1987, pp. 218–229.

[4] R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Commun. ACM*, vol. 21, no. 2, pp. 120–126, 1978.

[5] A. C. Yao, "How to generate and exchange secrets," in *Proc. 27th Annu. IEEE Symp. Foundations of Computer Science*, 1986, pp. 162–167.

[6] V. Kolesnikov, A.-R. Sadeghi, and T. Schneider, "How to combine homomorphic encryption and garbled circuits," in *Proc. Signal Processing in the Encrypted Domain–Ist SPEED Workshop,* Lausanne, Switzerland, 2009, pp. 100–121.

[7] C. Gentry, "Fully homomorphic encryption using ideal lattices," in *Proc. 41st Annu. ACM Symp. Theory of Computing*, 2009, pp. 169–178.

[8] P. Campisi, *Security and Privacy in Biometrics*. New York: Springer, 2013.

[9] J. Bringer, H. Chabanne, and A. Patey, "Privacy-preserving biometric identification using secure multiparty computation: An overview and recent trends," *IEEE Signal Processing Mag.*, vol. 30, no. 2, pp. 42–52, 2013.

[10] Y. Luo, S. S. Cheung, T. Pignata, R. Lazzeretti, and M. Barni, "An efficient protocol for private iris-code matching by means of garbled circuits," in *Proc. 19th IEEE Int. Conf. Image Processing (ICIP)*, 2012, pp. 2653–2656.

[11] C. Fontaine and F. Galand, "A survey of homomorphic encryption for nonspecialists," *EURASIP J. Inform. Security*, vol. 2007, no. 15, pp. 1–10, Jan. 2007.

[12] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *Proc. Advances in Cryptology (EUROCRYPT99)*, 1999, pp. 223–238.

[13] I. Damgard, M. Geisler, and M. Kroigard, "Homomorphic encryption and secure comparison," *Int. J. Appl. Cryptogr.*, vol. 1, no. 1, pp. 22–31, 2008.

[14] T. ElGamal, "A public key cryptosystem and a signature scheme based on discrete logarithms," in *Proc. Advances in Cryptology*, 1985, pp. 10–18.

[15] P. S. Pisa, M. Abdalla, and O. C. M. B. Duarte, "Somewhat homomorphic encryption scheme for arithmetic operations on large integers," in *Proc. Global Information Infrastructure and Networking Symp. (GIIS)*, 2012, pp. 1–8.

[16] M. O. Rabin, "How to exchange secrets by oblivious transfer," Tech. Rep. TR-81, Aiken Computation Lab., Harvard Univ., 1981.

[17] D. Beaver, "Precomputing oblivious transfer," in *Proc. Advances in Cryptology (CRYPT095)*, 1995, pp. 97–109.

[18] R. Lazzeretti and M. Barni, "Private computing with garbled circuits," *IEEE Signal Processing Mag.*, vol. 30, no. 2, pp. 123–127, Mar. 2013.

[19] Z. Erkin, M. Franz, J. Guajardo, S. Katzenbeisser, I. Lagendijk, and T. Toft, "Privacy-preserving face recognition," in *Proc. Privacy Enhancing Technologies*, 2009, pp. 235–253.

[20] J. Bringer, H. Chabanne, M. Favre, A. Patey, T. Schneider, and M. Zohner, "GSHADE: Faster privacy-preserving distance computation and biometric identification," in *Proc. 2nd ACM Workshop on Information Hiding and Multimedia Security*, 2014, pp. 187–198.

[21] J. Troncoso-Pastoriza, D. Gonzalez-Jimenez, and F. Perez-Gonzalez, "Fully private noninteractive face verification," *IEEE Trans. Inform. Forensics Security*, vol. 8, no. 7, pp. 1101–1114, July 2013.

[22] J. Bringer, H. Chabanne, D. Pointcheval, and Q. Tang, "Extended private information retrieval and its application in biometrics authentications," in *Proc. Cryptology and Network Security*, Singapore, 2007, pp. 175–193.

[23] M. A. Turk and A. P. Pentland, "Face recognition using eigenfaces," in *Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition*, 1991, pp. 586–591.

[24] AT&T Laboratories Cambridge. The database of faces (formerly "the ORL database of faces"). [Online]. Available: http://www.cl.cam.ac.uk/research/dtg/attarchive/facedatabase.html

[25] A. Sadeghi, T. Schneider, and I. Wehrenberg, "Efficient privacy-preserving face recognition," in *Proc. Information, Security and Cryptology (ICISC 2009)*, 2010, pp. 229–244.

[26] M. Osadchy, B. Pinkas, A. Jarrous, and B. Moskovich, "SCiFi—a system for secure face identification," in *Proc. IEEE Symp. Security and Privacy (SP)*, 2010, pp. 239–254.

[27] Y. Luo, S. S. Cheung, and S. Ye, "Anonymous biometric access control based on homomorphic encryption," in *Proc. IEEE Int. Conf. Multimedia and Expo (ICME)*, 2009, pp. 1046–1049.

[28] J. Daugman, "How iris recognition works," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 14, no. 1, pp. 21–30, 2004.

[29] T. Tan and Z. Sun. (2005). Casia-irisv3. Tech. Rep. [Online]. Chinese Academy of Sciences Institute of Automation. Available: http://www. cbsr.ia.ac.cn/IrisDatabase.htm

[30] M. Blanton and P. Gasti, "Secure and efficient protocols for iris and fingerprint identification," in *Proc. Computer Security (ESORICS 2011)*, pp. 190–209.

[31] J. Bringer, M. Favre, H. Chabanne, and A. Patey, "Faster secure computation for biometric identification using filtering," in *Proc. 5th IAPR Int. Conf. Biometrics (ICB)*, 2012, pp. 257–264.

[32] A. K. Jain, L. Hong, S. Pankanti, and R. Bolle, "An identity-authentication system using fingerprints," *Proc. IEEE*, vol. 85, no. 9, pp. 1365–1388, 1997.

[33] M. Barni, T. Bianchi, D. Catalano, M. Di Raimondo, R. Donida Labati, P. Failla, D. Fiore, R. Lazzeretti et al., "A privacy-compliant fingerprint recognition system based on homomorphic encryption and fingercode templates," in *Proc. 4th IEEE Int. Conf. Biometrics: Theory Applications and Systems (BTAS)*, 2010, pp. 1–7.

[34] M. Barni, T. Bianchi, D. Catalano, M. Di Raimondo, R. Donida Labati, P. Failla, D. Fiore, R. Lazzeretti et al., "Privacy-preserving fingercode authentication," in *Proc. 12th ACM Workshop on Multimedia and Security*, 2010, pp. 231–240.

[35] D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar, *Handbook of Fingerprint Recognition*. New York: Springer, 2009.

[36] D. Evans, Y. Huang, J. Katz, and L. Malka, "Efficient privacy-preserving biometric identification," in *Proc. 17th Conf. Network and Distributed System Security Symp. (NDSS)*, 2011.

[37] S. Goldwasser, S. Micali, and A. Wigderson, "How to play any mental game, or a completeness theorem for protocols with an honest majority," in *Proc. 19th Annu. ACM Symp. Theory of Computing*, 1987, vol. 87, pp. 218–229.

[38] C. Gentry and S. Halevi, "Implementing Gentry's fully-homomorphic encryption scheme," in *Proc. Advances in Cryptology (EUROCRYPT)*, 2011, pp. 129–148.

[39] M. Yasuda, T. Shimoyama, J. Kogure, K. Yokoyama, and T. Koshiba, "Packed homomorphic encryption based on ideal lattices and its application to biometrics," in *Proc. Security Engineering and Intelligence Informatics*, 2013, pp. 55–74.

[40] C. Blundo, E. De Cristofaro, and P. Gasti, "EsPRESSo: efficient privacy-preserving evaluation of sample set similarity," in *Proc. Data Privacy Management and Autonomous Spontaneous Security*, 2013, pp. 89–103.

[41] G. Fanti, M. Finiasz, G. Friedland, and K. Ramchandran, "Toward efficient, privacy-aware media classification on public databases," in *Proc. Int. Conf. Multimedia Retrieval*, 2014, p. 49.

[42] K. Messer, J. Matas, J. Kittler, J. Luettin, and G. Maitre, "XM2VTSDB: The extended M2VTS database," in *Proc. 2nd Int. Conf. Audio and Video-Based Biometric Person Authentication*, 1999, vol. 964, pp. 965–966.

[43] Neurotechnology, dataset cross match verifier 300. [Online]. Available: http://www.neurotechnology.com

[SP]

[Meng-Hui Lim, Andrew Beng Jin Teoh, and Jaihie Kim]

Biometric Feature-Type Transformation

[Making templates compatible for template protection]

©ISTOCKPHOTO.COM/GREYFEBRUARY

Biometrics Security and Privacy Protection

Biometrics refers to physiological (i.e., face, fingerprint, hand geometry, etc.) and behavioral (i.e., speech, signature, keystroke, etc.) traits of a human identity. As these traits are unique to individuals, biometrics can be used to identify users reliably in many authentication applications, such as access control and e-commerce. Most biometric authentication systems offer great convenience without requiring the users to possess or remember any secret credentials. For applications that demand greater security, biometrics can be used in complement with passwords and security tokens to offer a multifactor authentication.

Digital Object Identifier 10.1109/MSP.2015.2423693
Date of publication: 13 August 2015

Biometric measurement extracted from a user is subject to variations due to inconsistent environmental conditions or restrictiveness of the biometric representation used. The authentication decision of a biometric system is typically made upon how similar a sample is with reference to a template that is enrolled to the system at an earlier time. However, biometric templates used for similarity evaluation cannot be stored unprotected in the system because the consequences of biometric compromise is highly devastating. First, the stolen templates can be used to impersonate the corresponding users in the other applications. Second, biometrics is irrevocable and irreplaceable if it is compromised. A possible solution is to protect the biometric templates with a conventional encryption scheme such as the Rivest–Shamir–Adleman

Cryptosystem (commonly known as RSA). However, the similarity between samples before encryption cannot be preserved in the encrypted domain [12]. At the same time, the template must not be decrypted for matching because this may otherwise lead to exposure of the template. As a result, template protection techniques are developed to preserve template secrecy while allowing a similarity assessment to be carried out concurrently.

INTRODUCTION

Biometric template protection has been a prominent concern when biometric data is used in recognition. To securely match a biometric sample against a template, the similarity has to be computed in the encrypted domain without exposing the template throughout the matching process. A popular approach to protect biometric templates is to adopt a biometric cryptosystem [12]. These schemes are designed to bind a cryptographic key to the biometric or to generate a cryptographic key from the biometrics. Biometric cryptosystems that have been developed so far include fuzzy commitment [14], fuzzy extractor [9], and fuzzy vault [13]. In particular, fuzzy commitment binds a binary key to a binary biometric representation, and this key can only be recovered if a similar binary biometric representation is presented. Fuzzy extractor [secure sketch (SS) with strong randomness extractors] that operates in the Hamming domain converts nonuniform binary biometric inputs into stable binary strings, which can be used as encryption keys for subsequent cryptographic applications. Fuzzy vault projects a point set onto a polynomial whose coefficients are encoded by a selected binary key and conceals these genuine points within a set of chaff points that do not lie on the polynomial. The binary key can only be recovered for positive verification if the polynomial can be reconstructed through identifying sufficient points in another point set.

One of the main features of the biometric cryptosystems is that these schemes can only operate on a certain form of input, such as point-set or binary features. Incompatibility issues arise when the type of intended biometric features does not match the acceptable input type of a biometric cryptosystem. As an example, fuzzy vault only accepts as input an unordered feature set such as the fingerprint minutia set. Any other biometric modalities, such as face, palm print and iris, which are typically represented by an ordered feature set needs to be converted into the unordered counterpart before fuzzy vault can be applied. Further, this incompatibility issue could also occur when other means of template protection such as secure two-party computation mechanisms (that accept only ordered integer or binary inputs) [4], and multibiometric feature-level fusion [24] (that requires the same type of features from different biometric sources) are applied. This issue can be resolved only when the features can be transformed to the required feature type in an appropriate manner.

Apart from this problem, features are also transformed to meet the increasing demand of simplified, discriminative biometric representation for efficient processing. In a distributed and remote biometric verification scenario, lightweight devices such as portable devices may also need to protect templates using biometric cryptosystems in their local systems. As these devices are only equipped with restricted storage and processing capabilities, these transformations are essential for extracting compact discriminative representation for fast matching, thus allowing practical applications of these devices.

This article introduces the problem of biometric feature-type transformation and surveys the recent advances in several transformation approaches from both the signal processing and security points of view and is intended for both the generally knowledgeable and nonspecialist. We place more emphasis on one of the approaches—ordered-to-ordered transformation—because of its major advances over the past decade and its greater popularity over the other approaches. An empirical study will be presented to compare several ordered-to-ordered feature-type transformation techniques. Most challenges in this domain and directions for future work will also be discussed.

BIOMETRIC FEATURE-TYPE TRANSFORMATION

Biometric feature-type transformation is designed to satisfy at least one of the following objectives: 1) converts an incompatible biometric feature set into another to match the acceptable input

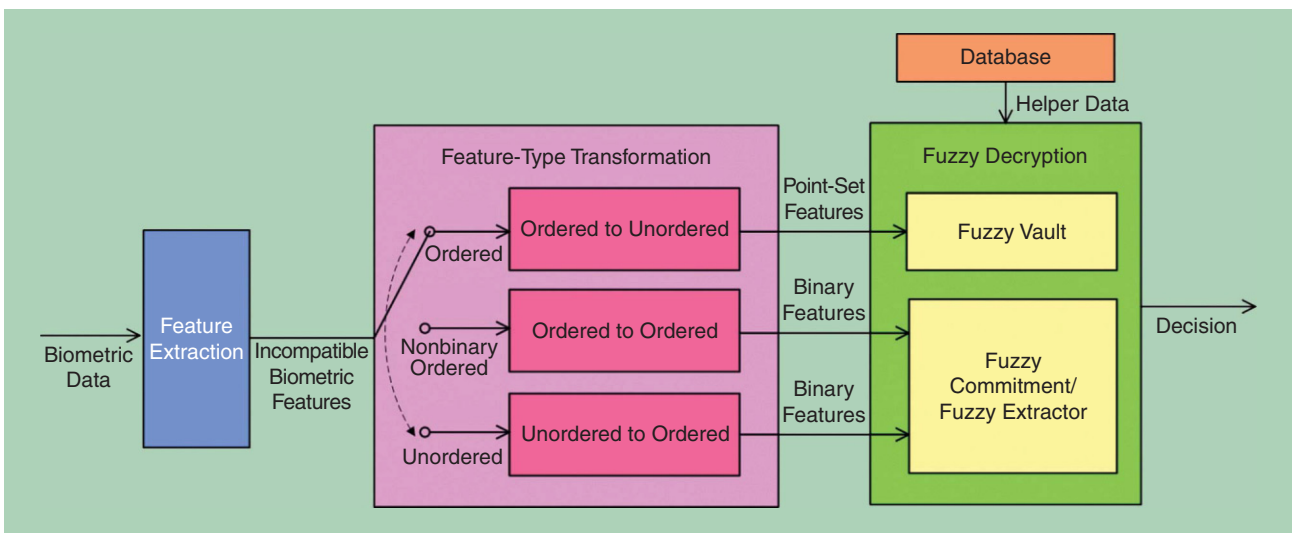

[FIG1] A generic biometric cryptosystem with feature-type transformation.

data type of a biometric cryptosystem and 2) represents features in a simplified yet discriminant form for lightweight processing. Unlike conventional biometric systems that rely on both feature extractor and a sophisticated matcher/classifier to resolve the intraclass and interclass variations of biometric data, a feature-type transformation algorithm is solely meant to address these problems. Hence, the matching can be done with plain distance measures such as Hamming distance or set difference.

Figure 1 shows a generic biometric cryptosystem with feature-type transformation. With the aid of auxiliary information or helper data of the biometric cryptosystem, this transformed feature set can be used by the fuzzy decryptor (decoder) of the biometric cryptosystem to determine if the query data is sufficiently similar to the template and whether a positive match has been obtained.

Biometric feature-type transformation has a threefold criterion: 1) preserve or improve the discriminability of the original features, 2) improve the randomness of the original features, and 3) avoid storage of biometric-correlated auxiliary data that may threaten biometric privacy if the system is compromised.

Generally, the discriminability of the original features is preserved or improved via further minimizing intrauser and maximizing interuser variations. To improve the randomness of the transformed features, uniformity of transformed element values is essential, and any interdependence among the original feature elements needs to be eliminated. In the case of ideal uniformity and zero-interdependence, adversarial guessing of the transformed features (input of the biometric cryptosystem) is the hardest, and this potentially optimizes the security of the system. Finally, the transformation should not depend on any auxiliary data that may leak information about the biometrics, since this data is usually stored in the system and could be compromised. A promising transformation scheme would warrant a negligible correlation between the auxiliary data and the user biometrics. While there is always an inevitable tradeoff among these requirements, the development of transformation techniques that could achieve a great balance among these requirements is among the most important problems in this field.

In general, biometric feature-type transformation can be classified into *ordered-to-ordered*, *unordered-to-ordered*, and *ordered-to-unordered* transformation. An ordered feature set is defined as a one-dimensional feature vector or a two-dimensional (2-D) feature matrix that contains ordered elements. Examples of the one-dimensional feature vector include subspace-projected feature vectors derived from principal component analysis (PCA) or linear discriminant analysis (LDA), normalized histogram feature vectors induced by local binary pattern (LBP), bag of features model, and time sequences, while examples of the 2-D feature matrix include a 2-D spectral transform feature matrix such as Fourier transform and wavelet transform, a 2-D filter such as Gabor filters, and a 2-D-subspace projected feature matrix. Both the one-dimensional vector and matrix are interchangeable as long as the order of their elements remains intact. Ordered feature sets are not limited to sets with fixed cardinality but also sets that are variable in length (e.g., online handwritten signature, speech signal time stamps).

An unordered feature set is defined as a feature set with unordered elements. Examples include fingerprint minutiae set, a variable quantity of timing statistics of a handwritten signature such as mean and standard deviation of x-y coordinates, and facial and palm print shift invariant feature transform (SIFT) keypoint descriptors.

ORDERED-TO-ORDERED TRANSFORMATION

Ordered-to-ordered transformation converts a type of ordered feature set into another type of ordered set mainly for fuzzy commitment and fuzzy extractor applications.

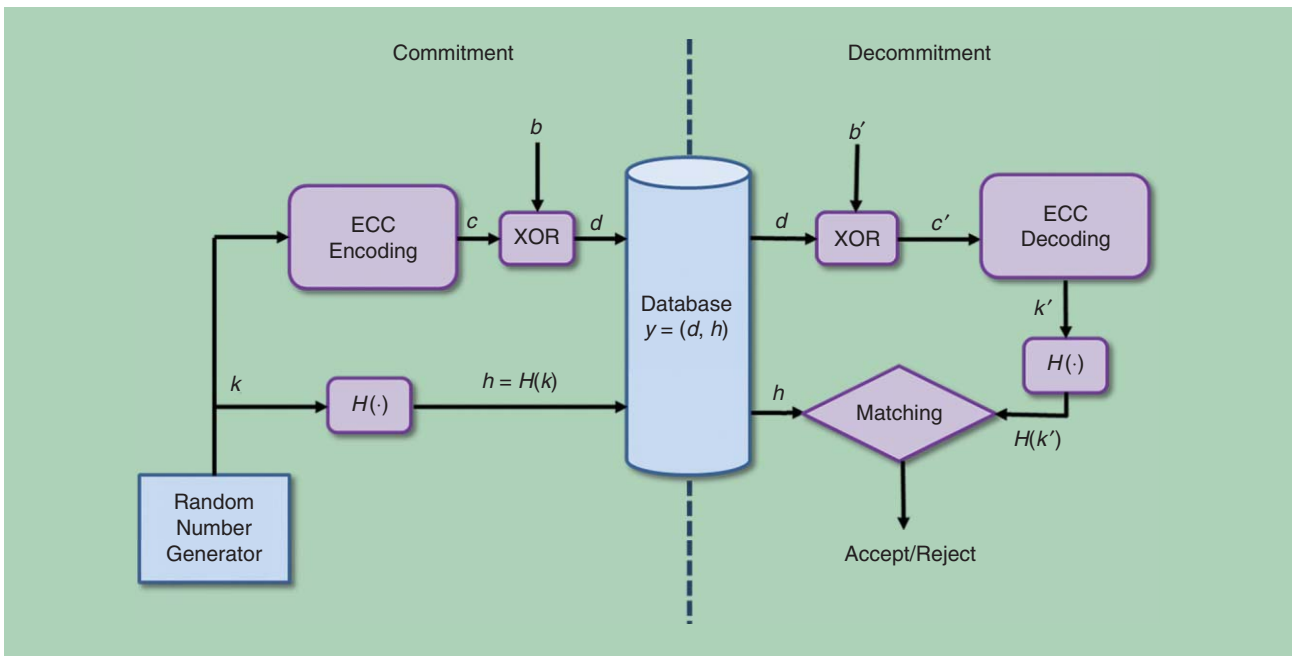

[FIG2] **A fuzzy commitment scheme.**

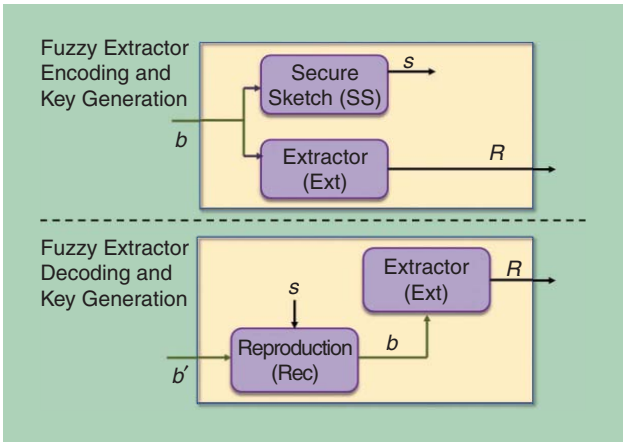

[FIG3] **A fuzzy extractor scheme.**

FUZZY COMMITMENT

Fuzzy commitment shown in Figure 2 is a cryptographic construction that binds a random codeword c among a set of equally distanced codewords to a biometric measurement b, where the codeword c is produced via encoding a secret binary key k with an error-correcting code (ECC). By computing an offset $d = c \oplus b$ and a one-way hashed value $h = H(k)$ of the key, the commitment $y = (d,h)$ can be produced. While it is not possible to reconstruct k using y alone, the key k can only be recovered in the presence of a measurement b' that is sufficiently close to b. In particular, to decommit y, one can obtain a "corrupted" codeword via $c' = d \oplus b'$, decode c' to obtain k', and verify if $h = H(k')$. A match can only be obtained when the Hamming distance between b and b' lies within the error-correcting capability of ECC.

FUZZY EXTRACTOR

Fuzzy extractor shown in Figure 3 consists of two major components: an SS and a strong extractor (Ext). The former handles the intrauser variations of biometric features, while the latter attempts to generate uniformly distributed data from a nonuniform input. In the encoding stage, an SS algorithm is applied to a biometric measurement b to obtain a syndrome (sketch) s, and a strong extractor with randomness r is applied to b to obtain the cryptographic key R. In decoding stage, b is first recovered using the sketch s and a sufficiently close measurement b'. Then, similar to the encoding stage, the same strong extractor can be applied to b to reproduce the same cryptographic key R.

Both fuzzy commitment and fuzzy extractor are based on the ECC, where the dissimilarity between a biometric sample and the template is treated as bit errors in a binary code and the success of decoding is dependent on whether the quantity of bit errors lies within the error correcting capability of the code. Hence, these cryptographic constructions require a continuous feature set to be converted into its binary counterpart before these cryptosystems can be applied.

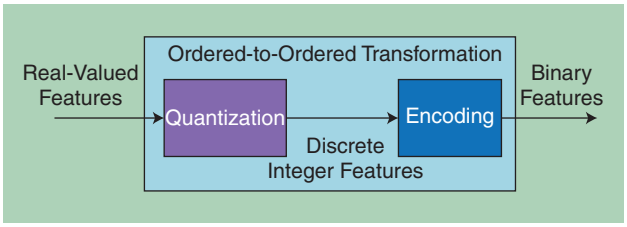

[FIG 4] **A continuous-to-binary ordered transformation.**

Figure 4 depicts the block diagram of a typical continuous-to-binary ordered transformation process. A continuous feature set is transformed into a discrete integer set through quantization. Each feature is quantized into integer-labeled intervals, where feature elements falling within an interval is mapped to the corresponding integer label. If an integer set is required, these individual integer labels can be concatenated to form an integer set. Otherwise, a binary encoding can be applied to map each integer to a short binary string. The final binary representation is formed through concatenating these individual binary strings.

QUANTIZATION ALGORITHMS

SIMPLE STATIC QUANTIZATION

Let us start by introducing a few simple yet representative quantization schemes to explain how a continuous feature set can be converted into a discrete integer set. The simplest quantization, known as *static quantization,* constructs equal partitions in each feature space, where these feature spaces are normally quantized independently.

Equal-width quantization is an instance of static quantization. This quantization directly segments a feature space V into S equal-width intervals. However, this quantization has several drawbacks: it is sensitive to feature outliers, as the quantization depends on the range of the feature values. If the corresponding population feature distribution is not uniform, the probability of certain integer values for such a feature would be higher, leading to a lower randomness in the output of such a component and an easier adversarial guessing attack. Privacy-wise, the storage of the auxiliary data would only leak the range of training features, but not information regarding user features. Hence, the auxiliary data is not correlated with any genuine user's biometrics.

Equal-probable quantization, on the other hand, segments a feature space V into S intervals encapsulating equal population probability mass $1/S$. As the probability of all the output values is equal, this quantization offers maximum randomness. Similar to equal-width quantization, the auxiliary data does not leak user-specific information.

To increase discriminability of the quantization output, labeling of training samples is utilized by entropy quantization [16], where this technique aims to induce a set of intervals through recursive space-splitting so that every interval encloses only feature components of a class (either genuine or imposter). As a result, the cutpoints υ are optimized, such that the entropy in each interval is minimal. Although this technique improves output discriminability, it is unable to guarantee that the output values are equally probable. Additionally, it offers lower privacy than equal-width and equal-probable quantization due to the fact that majority training features of a user in each interval

must have been enclosed within one of the intervals. An illustration of equal-width, equal-probable, and entropy quantization schemes is given in Figure 5.

EXTENSION OF SIMPLE QUANTIZATION ALGORITHMS

Static quantization standardizes the range of integer output over the feature components. However, when individual feature components do not offer the same discrimination power (e.g., Eigenface, Gabor features [15]), the discriminability of the transformed feature set can further be improved using dynamic quantization. Dynamic quantization weighs feature components distinctly according to the discriminability of the components, so that the discriminative components take a higher weight in the match score computation [18]. With this, for the same genuine pairwise similarity in two normalized spaces, the similarity of the higher-weighted space will have a larger influence in the final match score. If the weights are confined to binary [15], the dynamic quantization would then become a feature selection method. In dynamic quantization, the weight of a component is characterized by the corresponding range of integer outputs, i.e., the number of quantization intervals constructed in the corresponding feature space. Dynamic quantization seeks an optimum quantity of quantization intervals over the feature spaces to minimize the false acceptance rate (FAR) and/or false rejection rate (FRR).

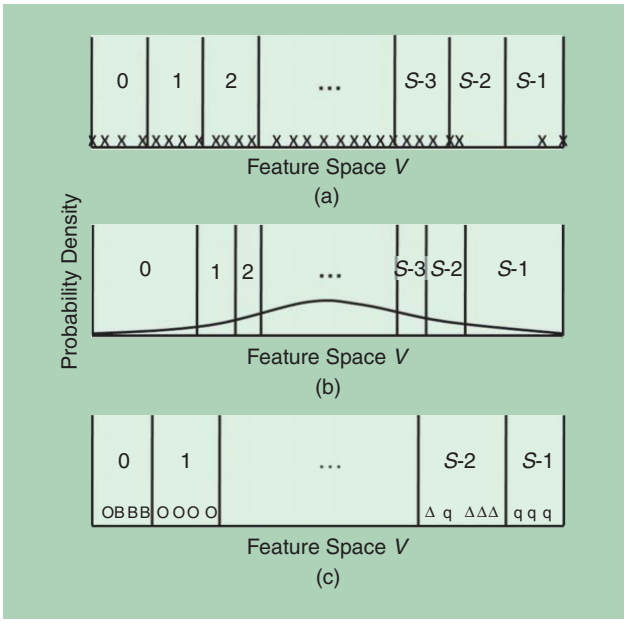

[FIG5] Illustrations of (a) equal width, (b) equal-probable quantization, and (c) entropy quantization, where "X" in (a) denotes unlabeled samples; probability density function in (b) denotes the population distribution; and "O," "B," "Δ," and "Q" in (c) denote labeled samples from four different classes.

Detection Rate Optimized Bit Allocation

Detection rate optimized bit allocation (DROBA) [7] is a representative instance, where this method minimizes FRR by maximizing the overall detection rate, which is the probability of a genuine user's feature component staying within the same partition. By assuming negligible dependence among the individual one-dimensional feature spaces, DROBA creates partitions on these feature spaces iteratively and weighs feature components through allocating binary bits to these spaces. The higher the number of bits is allocated to a feature space, the higher the weight is assigned to that feature component.

This algorithm begins by estimating the detection rate for two-interval quantization on all spaces and picks the space with the highest detection rate for a two-interval split (allocation of the first bit). In the next iteration, the detection rate for this selected space is updated with a 2^2-interval detection rate and then it is compared with two-interval detection rates of all other spaces for the selection of next best interval split (allocation of the second bit). This process is repeated until a terminating condition is met (e.g., bit-quantity requirement or detection rate threshold). Once the algorithm terminates, spaces that have not been selected for a split will be discarded.

An example of DROBA is given in Figure 6, where a 4-bit-allocation process is demonstrated. In this example, the 2-D user feature distribution has a wider spread along the horizontal axis (feature space 1), indicating a larger intrauser variation along that dimension. At the end of the algorithm, only one bit is allocated to the less discriminative feature space 1 and three bits are allocated to the more discriminative feature space 2.

When equal-probable quantization is adopted in the quantization process, DROBA offers maximum randomness to its quantization output. By labeling the constructed intervals with a set of integer values, DROBA extracts a larger range of integer values from discriminative spaces and a smaller range from less discriminative feature components for a higher matching performance. Compared to static quantization, the auxiliary data of DROBA reveals additional information regarding the importance of each feature space due to the additional storage of quantity of quantization intervals allocated on the feature spaces in the system databases [23].

Reliability-Dependent Bit Allocation

As the estimation of detection rate could be computationally intensive when a large number of spaces or intervals are concerned, an efficient dynamic quantization alternative could be reliability-dependent bit allocation (RDBA) [20]. Different from DROBA, which uses a top-down interval-splitting approach, RDBA uses a bottom-up interval-merging approach. RDBA uses reliability of binary-encoded training samples as the interval-merging measure. Hence, the main efficiency improvement of RDBA comes from its bit counting instead of estimation of maximum probability mass over the intervals. RDBA uses binary reflected gray code (BRGC) [11] to encode training measurements of each user's feature component based on a predefined maximum number of intervals that can be constructed on a space. Based on the encoded training measurements, a bit agreement probability is estimated for every bit and an interval merging strategy is used in RDBA based on the entire set of bit agreement probabilities (over all feature components). Similar to DROBA, the integer merging process is repeated until a terminating condition is met (e.g., bit-quantity requirement or bit-reliability threshold).

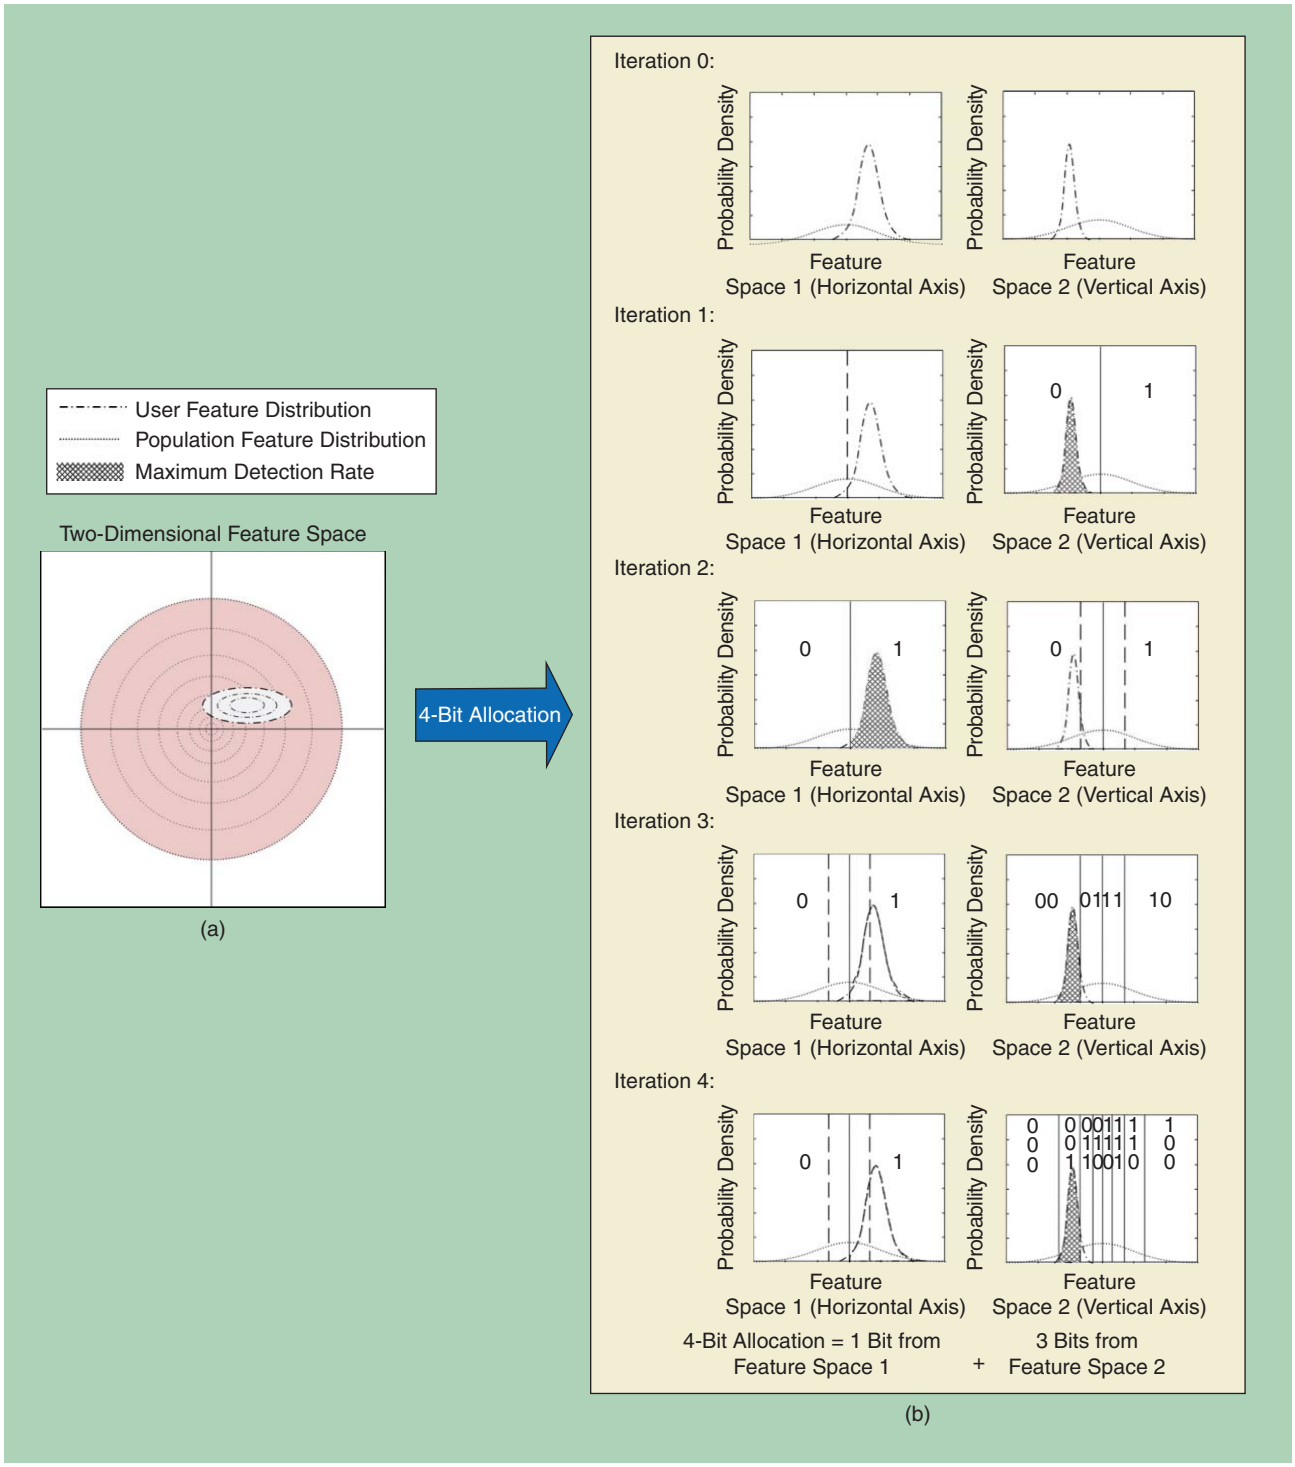

[FIG 6] **A simple bit-allocation scenario of DROBA.**

OTHER TYPES OF QUANTIZATION

The aforementioned static and dynamic quantization techniques that quantize the features individually did not take into account the dependency among the features. This approach falls into the category of univariate quantization. In fact, biometric features are rarely independent to each other. Hence, the univariate quantization schemes may overlook meaningful interaction among the components, which could be essential for classification. From a high-dimensional point of view, univariate quantization induces hypercubical/rectangular segments, which in many situations may fail to well-accommodate nonuniform intrauser samples distribution, causing them to perform suboptimally.

There are other quantization approaches that take into account the dependency among components when quantization is performed: semimultivariate quantization considers multiple subsets of features and quantizes these subsets individually; and multivariate quantization considers the feature components as a whole. For instance, Chen et al. [8] introduced a semimultivariate polar quantization technique in which every two feature components in Cartesian coordinates are paired up for extracting the corresponding polar coordinates (phase and magnitude features) for quantization. Lim et al. [21] put forward a multivariate quantization technique that induces meaningful convex hyperpolygonal Voronoi segments on the high-dimensional feature space using medoid-based clustering. Although it has better classification results, the high-dimensional population and user feature probability distribution function is often difficult to estimate. Hence, designing an effective quantization in this aspect remains a great challenge.

ENCODING

SIMPLE ENCODING SCHEMES

Single-bit encoding can be applied when there are only two possible integer labels. When multiple bits are desired for more than two integer possibilities, there are several binary encoding variants one can opt for. The simplest encoding scheme is to represent integers with their direct binary counterparts. This encoding is known as *direct binary representation* (*DBR*). For instance, integers 3_{10}, 4_{10}, and 7_{10} are converted to binary codewords 011_2, 100_2, and 111_2, respectively. However, the distance between two elements in the discrete domain cannot be preserved proportionally by their counterparts in the Hamming domain. For instance, the distances between 3_{10} and $4_{10}(=1)$ and the distance between 3_{10} and $7_{10}(=4)$ do not differ proportionally with the Hamming distances between 011_2 and $100_2(=3)$ and between 011_2 and $111_2(=1)$, respectively. This motivates the employment of a more effective encoding scheme called *BRGC* [11]. Given a bit length n, BRGC has all n-bit binary codewords reordered, such that every consecutive pair of codewords differs by a single bit. If a genuine feature component falls within a neighboring interval and tagged with an integer value that is adjacent to that of a reference genuine component, the BRGC encoding scheme ensures that the Hamming distance caused by such intrauser variation is precisely limited to one bit. However, even if BRGC encoding confines distance between neighboring elements in the discrete domain to one bit, the distance mapping between two nonneighboring elements is not precise. Several instances of BRGC code are shown in Table 1, where the required code length for S integer components is $n = \lceil \log_2 S \rceil$ bits.

In general, when an n-bit code is used to transform $S = 2^n$ integer components, distance between two points in the discrete domain cannot be preserved exactly in the Hamming domain. This is because, when the code length is not sufficiently large, multiple codewords would share a common Hamming distance with respect to a reference codeword for $n > 1$. Therefore, there is a significant possibility where imposter feature elements that are well-separated from a genuine feature element in the discrete feature space are mapped to closer codewords in the Hamming space, increasing the possibility of misclassification.

[TABLE 1] BRGC ENCODING.

BRGC			
$n = 2/S = 4$	$n = 3/S = 8$	$n = 4/S = 16$	
$0_{10} \rightarrow 00_2$	$0_{10} \rightarrow 000_2$	$0_{10} \rightarrow 0000_2$	$8_{10} \rightarrow 1100_2$
$1_{10} \rightarrow 01_2$	$1_{10} \rightarrow 001_2$	$1_{10} \rightarrow 0001_2$	$9_{10} \rightarrow 1101_2$
$2_{10} \rightarrow 11_2$	$2_{10} \rightarrow 011_2$	$2_{10} \rightarrow 0011_2$	$10_{10} \rightarrow 1111_2$
$3_{10} \rightarrow 10_2$	$3_{10} \rightarrow 010_2$	$3_{10} \rightarrow 0010_2$	$11_{10} \rightarrow 1110_2$
	$4_{10} \rightarrow 110_2$	$4_{10} \rightarrow 0110_2$	$12_{10} \rightarrow 1010_2$
	$5_{10} \rightarrow 111_2$	$5_{10} \rightarrow 0111_2$	$13_{10} \rightarrow 1011_2$
	$6_{10} \rightarrow 101_2$	$6_{10} \rightarrow 0101_2$	$14_{10} \rightarrow 1001_2$
	$7_{10} \rightarrow 100_2$	$7_{10} \rightarrow 0100_2$	$15_{10} \rightarrow 1000_2$

AN EFFECTIVE ENCODING SCHEME

To address the intrinsic limitation of DBR and BRGC encoding, codes with larger length $n > \log_2 S$ are employed to encode S integers. Linearly separable subcode (LSSC) [22], also known as the general unary code, is an effective code instance with S varying linearly with n $(n = S + 1)$. With this encoding scheme, the indefinite discrete-to-binary mapping of DBR and BRGC can completely be overcome and the norm-1 distance can be exactly computed by measuring the Hamming distance on LSSC-coded representation of the feature vector. For example, in a 5-bit LSSC-coded representation, the Hamming distance between "01111_2" representing 4_{10} and "00001_2" representing 1_{10} is 3 bits, which equals the norm-1 distance of their discrete counterpart. In fact, this code is the only binary encoding scheme that could convert the distance between two discrete integers into its equivalent distance in the Hamming domain.

With LSSC coding, dynamic quantization can generally gain better precision in its component-weighting or bit-allocation process. The linear property between S and n in LSSC could enable unity interval splitting/merging with each increment in its allocated weight, instead of double interval splitting/merging when DBR or BRGC encoding is used. Hence, this extra flexibility allows dynamic quantization schemes to search for a more appropriate number of quantization intervals that better accommodates the genuine user pdf on the feature space.

UNORDERED-TO-ORDERED TRANSFORMATION

Unordered-to-ordered transformation refers to the conversion of an unordered, variable-sized feature set (mainly minutiae point set) into an ordered, fixed-length feature vector or binary string for the application of fuzzy commitment, SS, and fuzzy extractor in the Hamming domain [12]. The general premise of this transformation is to protect minutiae locations and orientation so that recovering these minutiae features from the transformed features is infeasible. This transformation complements the need of privacy (hardness of feature reversing to recover original biometric data) for most biometric cryptosystems. However, this transformation would lead to certain information loss and a drop in performance accuracy. Hence, performance and privacy preservation are the two main concerns in the design of unordered-to-ordered transformation

schemes. Note that unordered-to-ordered transformation can be assimilated with ordered-to-ordered transformation to yield more discriminative ordered fixed-length binary string.

In general, the unordered-to-ordered transformation methods for fingerprint minutia can be broadly divided into two categories: reference based and spectral transform. These methods can also be adapted to face and palm print biometrics when unordered keypoint feature descriptors such as SIFT are employed.

REFERENCE-BASED APPROACH

In the reference-based approach, a reference is first defined and the fingerprint minutiae are quantized into an ordered, fixed-length representation with respect to this reference. For instance, Sutcu et al. [29] demonstrated that random cuboids on a fingerprint can also be used as the reference. Based on the number of minutiae points in each of the m cuboids, a binary string, which is transformed from the m-dimensional integer vector, can be generated. With reference to random cuboids, Nagar et al. [25] suggested a more robust set of features by considering the average minutia coordinate in a random cuboid, the standard deviation of the minutiae coordinates, and the aggregate wall distance. However, all of these methods either require 1) prealignment of fingerprint images, which is usually unavailable when the enrolled template is not stored in the clear; or 2) registration points (e.g., high curvature points) for alignment, which may leak information about the minutiae points.

To eliminate the need of prealignment, Bringer and Despiegel [3] used m minutiae vicinities as the reference, where a vicinity is the neighborhood around a minutia within a radius. The correspondence between each of the m vicinity in the enrolled template and the corresponding vicinity within the vicinities in the query template is sought for similarity assessment. Based on the m matching scores, this method generates an ordered vector, which can then be transformed into a binary string. However, this representation requires a high storage capability because the resultant bit length of the binary string can be very large (e.g., 50,000 bits). Farooq et al. [10] generated a binary fingerprint representation based on the histograms of triangular features generated from minutiae triplets, e.g., side length, and relative difference in angle. This method requires a high computational cost due to the exhaustive feature computation over all possible minutiae triplets. A better alternative [1] draws a line through the core point at an angle, θ specified by a user-specific key. Then, the minutiae are projected onto several axes: one parallel to the x-axis and another parallel to the y-axis, or an additional axis parallel to θ. These axes are segmented and the segments are indexed using a user-specific key. Finally, a histogram is built using the quantity of projected minutiae in the segments.

SPECTRAL-TRANSFORM-BASED APPROACH

The spectral transform approach performs Fourier transform on a minutia set and remaps the Fourier spectral magnitudes onto polar-logarithmic coordinates [33]. By doing so, the spectral minutiae representation is invariant to rotation, shifting, and scaling variations. An analytical representation of minutiae is further proposed to minimize error, which can be directly evaluated on polar-logarithmic grids. As the number of grids is fixed, a fixed-length representation can be derived. However, the discriminability of the transformed feature is not as par to its minutia counterpart. Instead of using magnitude spectrum, Nandakumar adopted phase spectrum of the minutiae: the binarized phase spectrum (BiPS) [26]. By incorporating fuzzy commitment and reliable bit selection schemes, BiPS achieves state-of-the-art accuracy performance over other biometric cryptosystems. However, BiPS is not rotation-, shift- and scale-invariant, and it requires proper image alignment for focal point estimation.

ORDERED TO UNORDERED

Ordered-to-unordered transformation converts an ordered feature set into an unordered feature set, where each transformed feature is usually represented by an integer or a binary string. This conversion is required when fuzzy vault is used to protect an ordered biometric feature set.

Fuzzy vault is another prominent fuzzy template protection primitive whose security relies upon the hardness of polynomial reconstruction. Fuzzy vault is primarily used in unordered sets of fingerprint minutia [27]. This scheme encodes a given secret K into coefficients of a polynomial P of degree D, projects an unordered genuine feature point set onto P, and conceals these genuine points within a much larger set of chaff points that do not lie on P. The secret K can only be retrieved through polynomial reconstruction if at least $D+1$ points can be identified. The uniqueness of fuzzy vault for an ordered biometric feature set lies in its tolerance to unstable, unreliable, or missing elements of features extracted during different acquisitions of biometrics. Fuzzy vault permits matching (during unlocking genuine points) in both plain [27] and encrypted domains [30], thus offering greater flexibility than the fuzzy commitment primitive.

A chief requirement of ordered-to-unordered transformation for fuzzy vault application is the independency and uniformity among the transformed elements because fuzzy vault is designed for randomly occurring feature points that are uniformly distributed and independent to one another. Many ordered biometric feature representations (e.g., Gabor features, LDA coefficients) are inherently highly correlated and nonuniform. While an ordered feature set can be naively treated as unordered point set, these correlated and nonuniform points in the set could weaken the security of fuzzy vault because it leads to easier genuine-point guessing and record multiplicity attacks. Another requirement of the transformation is performance preservation. When the genuine point unlocking is performed, the loss of discriminability due to quantization and binarization of the original feature set should not be significant after transformation. In addition, making features uncorrelated and uniform during the ordered-to-unordered transformation may cause inevitable degradation in performance due to loss of ordering information in the transformed features.

Lee et al. [19] converts an iris sample that can be represented by multiple iris patches (a 2-D ordered feature matrix) to an unordered point set. For each ordered row feature vector, an individual independent component analysis (ICA) feature vector is extracted

and mapped to an integer. The integer output from all the ordered feature vectors are concatenated to produce the corresponding unordered feature set. In this example, the accuracy performance can be preserved. However, an ICA projection is applied to each ordered feature vector individually. This causes each ICA feature vector to be mapped to an unordered point. As a result, the original correlation between different ordered vectors could not be reduced in the transformed feature set.

To decorrelate elements in the original feature set, a random projection can be applied to obtain uncorrelated points in the unordered feature set. Wang and Plataniotis [31] suggested a 2-D quantization of Euclidean-distance vectors between an ordered face feature vector and multiple pairs of random vectors. Each pair of quantization output is concatenated to yield a point on a high-dimensional Hamming space.

A COMPARATIVE STUDY

A comparative study has been conducted on the ordered-to-ordered (continuous-to-binary) feature transformation. In this study, PCA is used to extract features from the facial images of Facial Recognition Technology (FERET) data set and several ordered-to-ordered transformation techniques are applied to obtain different sets of binary features for performance evaluation. In Figure 7(a), the equal error rate performances of static quantization scheme, i.e., equal-probable and entropy quantization are quite close. Despite using different encoding schemes, equal-probable quantization achieves a better overall performance. The (supervised) entropy quantization, however, suffers a greater accuracy loss due to the overfitting problem caused by limited training samples. In addition, due to imprecise distance preservation by BRGC, static quantization, when incorporated with LSSC encoding, performs better. However, the bit redundancy of LSSC encoding-based static quantization shown in Figure 7(b) is almost twice as high.

On the other hand, dynamic quantization schemes such as Flexi-DROBA (DROBA with unity interval increment due to LSSC employment) and RDBA outperform the static quantization schemes in Figure 7(a). As unity interval increment is not feasible for RDBA, Flexi-DROBA performs the best. However, these dynamic quantization schemes inflict a significant amount of bits redundancy. Of the two dynamic quantization schemes, Flexi-DROBA+LSSC achieves the lowest bit redundancy due to more uniform interval allocation to the feature spaces.

ALIGNMENT AND PREPROCESSING

Feature-type transformation is employed to adapt biometrics of various forms of representation to the required type of input feature of a template protection scheme. However, for many biometrics, the query and enrollment biometrics need to be preprocessed and aligned using a complicated procedure before matching. This is very much neglected in the current implementation of a secure biometric system, causing an inevitable accuracy drop in matching. For instance, circular shifts of an IrisCode caused by the nonzero relative orientation of two irises are typically aligned during matching through various bit-shift comparisons. Some noisy bits extracted from the part of the iris occluded by the eyelid are also omitted from contributing to the final matching score through a "masking" process at the matching stage. However, biometric cryptosystems can only employ a simple Hamming distance metric that does not allow masking and bit-shift comparisons. Because these important processes at the matching stage cannot be performed when a biometric cryptosystem is applied, the recognition performance is likely to significantly degrade. To deal with these issues,

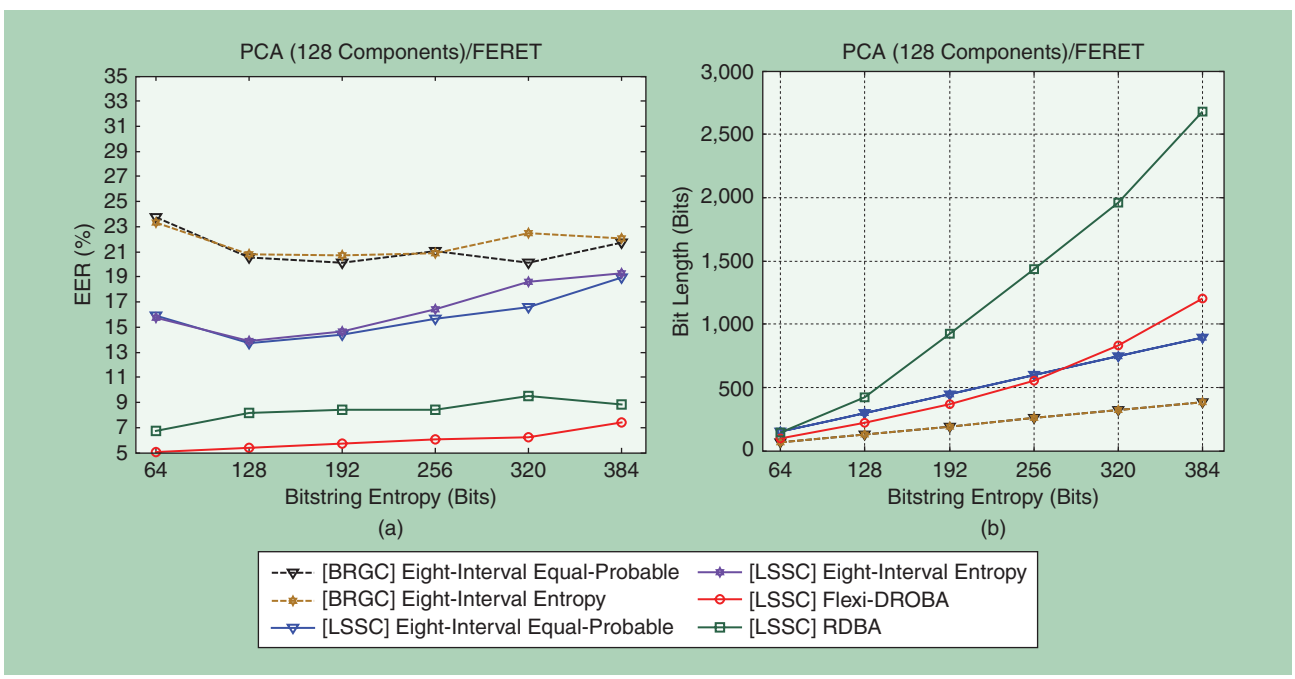

[FIG7] A comparison of ordered continuous-to-binary transformation schemes in (a) performance and (b) bit length of binary representation based on the FERET face data set.

many existing feature extractors have to be redesigned to avoid the need of extra processing at the matching stage when the template protection technique is incorporated. In this aspect, if one would want to adopt a nonalignment-free state-of-the-art feature extraction method in a secure biometric system, it is necessary to develop a general transformation method that could convert potentially misaligned features into alignment-free features without affecting the discriminability of the original features.

RELATED AND POTENTIAL APPLICATIONS

In template protection schemes that are based on secure two-party computation such as a homomorphic cryptosystem and garbled circuit [4], biometrics needs to be represented in integers or binary digits. However, unlike biometric cryptosystems, biometrics is the message to be protected in secure two-party computation-based template protection without simultaneously serving as an encryption/decryption key. In this case, the security and privacy requirements of the transformation can be neglected—leaving the performance preservation to be the sole requirement of the feature-type transformation.

Furthermore, biometric feature-type transformation comes into play when features of different modalities need to be fused at the feature level. For instance, unordered features of different biometric modalities are transformed to ordered features so that they are compatible for fusion [24].

Biometrics is also transformed into an ordered fixed-length binary string for speedy match and light storage, which is especially useful in large-scale biometric indexing for identification [32] and processing for portable and ubiquitous devices. These binary strings need to be of high stability and the bit length must be sufficiently lengthy to accommodate all biometric users.

CHALLENGES AND FUTURE DIRECTION

Over the past decade, significant progress has been made in biometric feature-type transformation for biometric cryptosystems. However, to date, there are still a number of challenges before these approaches can be reliably implemented in practice.

ORDERED-TO-ORDERED TRANSFORMATION

Most existing quantization techniques neglect dependency among feature components, leading to limited system security and recognition accuracy. Moreover, most existing dynamic quantization schemes such as DROBA and RDBA only focused on minimizing intrauser variation without concurrently maximizing interuser variation. An approach to extend this line of research is to look into static or dynamic quantization based on high-dimensional probability estimations for both population and/or user distributions (e.g., using a deep density model) and to consider population distribution in the bit-allocation strategy of the dynamic quantization techniques.

In addition, current feature-type transformation techniques are designed in isolation with the error-correcting capability (system decision threshold) of the biometric cryptosystem. To obtain optimal performance in practical biometric systems, transformation schemes should be designed with respect to an optimal quantity of tolerable bit errors, such that the quantity of genuine bit errors would lie within the error correcting capability of a biometric cryptosystem and the quantity of imposter bit errors would consistently exceed this error correction threshold.

Another line of research could be to explore efficient and effective image-hashing techniques (e.g., locality-sensitive hashing variant [2]) to extract compact discriminative representation for large-scale biometric indexing and lightweight processing for processors with limited storage and computational capabilities.

Furthermore, direct computation of binary descriptors such as binary robust independent elementary features (BRIEF) [5], oriented fast and rotated BRIEF descriptor [28], and ordinal spatial information of regional invariants [34] are also attracting great interest in the computer vision community. Unlike conventional techniques that follow the quantization-encoding pipeline such as DROBA and image-hashing techniques, direct computation of binary descriptors requires less computational time because it converts raw images to a binary string without undergoing the quantization process. However, how this technique can be applied effectively to the biometric data remains an open research problem.

UNORDERED-TO-ORDERED TRANSFORMATION

Many well-performing methods require prealignment or registration [25], [28], [33], which are either privacy-threatening or incompatible with the biometric cryptosystem framework (because template alignment is not possible with an encrypted enrolled template). Some methods even require high computational power and large template storage [3], [10]. As a result, lightweight, alignment-free, accuracy- and privacy preserving transformation methods still need to be sought. One promising direction is to adopt alignment-free minutia descriptors such as minutia cylinder code [6] so that each minutia in a fingerprint can be represented as a locally ordered fixed-length feature vector. However, these feature vectors extracted from a minutiae point set are globally unordered and may vary in number. In this respect, machine-learning and signal processing-based approaches such as kernel PCA, bag-of-features model, or dictionary learning can be applied to produce an ordered vector from the minutiae descriptors. If a binary string is desired, an ordered-to-ordered transformation can be stacked. These approaches can potentially be made privacy-preserving using privacy-preserving machine-learning and encrypted-signal processing techniques [17].

ORDERED-TO-UNORDERED TRANSFORMATION

This transformation is relatively less studied compared to the former two. While the three main requirements of this approach are accuracy preservation, independence, and uniformity among the transformed elements, current methods are not well evaluated according to these requirements. Thus, it would be interesting to study these methods in depth and explore new techniques for effective fulfillment of these requirements. One of the possible remedies could be to apply random projection prior to feature-type transformation. By using proper random projection matrix such as independent realization of ± 1 Bernouli random variables, independence and uniformity of feature elements can concurrently be achieved with certain performance preservation guaranteed by the Johnson–Lindenstrauss lemma.

ACKNOWLEDGMENTS

This research was supported by the Basic Science Research Program through the National Research Foundation of Korea funded by the Ministry of Science; Information, Communication, and Technology; and Future Planning (2013006574).

AUTHORS

Meng-Hui Lim (menghuilim@comp.hkbu.edu.hk) received his Ph.D. degree from Yonsei University, South Korea, in 2012. He joined the Department of Computer Science at Hong Kong Baptist University as a postdoctoral research fellow, where he has been a research assistant professor since 2013. His research interests include pattern recognition, cryptography, and biometric security. He is a Member of the IEEE.

Andrew Beng Jin Teoh (bjteoh@yonsei.ac.kr) is an associate professor in the Electrical and Electronic Department, College Engineering of Yonsei University, South Korea. His research, for which he has received government and industry funding, focuses on biometric security, specifically in biometric template protection and biocrypto key computation. He has published more than 220 international refereed journals, conference articles, and several book chapters in the areas mainly in biometric security and biometric systems. He is a senior member of the IEEE Signal Processing Society. He is the chair of *IEEE Biometric Council Newsletter*.

Jaihie Kim (jhkim@yonsei.ac.kr) received the Ph.D. degree in electrical engineering from Case Western Reserve University, Cleveland, Ohio, in 1984. Since 1984, he has been a professor in the School of Electrical and Electronic Engineering, Yonsei University, South Korea. Currently, he is the director of the Biometric Engineering Research Center in South Korea. His general research interests are biometrics, pattern recognition, and computer vision. Some of his recent research topics include mobile biometrics, biometric template protection, 2-D-to-3-D face conversion, and face age estimation/synthesis. He has been an author and associate editor of many international technical journals.

REFERENCES

[1] T. Ahmad and J. Hu, "Generating cancelable biometric templates using a projection line," in *Proc. Int. Conf. Control, Automation, Robotics and Vision*, 2010, pp. 7–12.

[2] A. Andoni, P. Indyk, H. L. Nguyen, and I. Razenshteyn, "Beyond locality-sensitive hashing," in *Proc. Annual ACM-SIAM Symp. Discrete Algorithms*, 2014, pp. 1018–1028.

[3] J. Bringer and V. Despiegel, "Binary feature vector fingerprint representation from minutiae vicinities," in *Proc. IEEE Int. Conf. Biometrics: Theory Applications and Systems*, 2010, pp. 1–6.

[4] J. Bringer, H. Chabanne, and A. Patey, "Privacy-preserving biometric identification using secure multiparty computation: An overview and recent trends," *IEEE Signal Processing Mag.*, vol. 30, no. 2, pp. 42–52, 2013.

[5] M. Calonder, V. Lepetit, M. Ozuysal, T. Trzcinski, C. Strecha, and P. Fua, "BRIEF: Computing a local binary descriptor very fast," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 7, pp. 1281–1298, 2012.

[6] R. Cappelli, M. Ferrara, and D. Maltoni, "Minutia Cylinder-Code: A new representation and matching technique for fingerprint recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 12, pp. 2128–2141, 2010.

[7] C. Chen, R. Veldhuis, T. Kevenaar, and A. Akkermans, "Biometric quantization through detection rate optimized bit allocation," *EURASIP J. Adv. Signal Process.*, vol. 2009, e784834, pp. 1–16, May 2009.

[8] C. Chen and R. Veldhuis, "Binary biometric representation through pairwise adaptive phase quantization," *EURASIP J. Inform. Security*, vol. 2011, Article ID 543106, pp. 1–16, Feb. 2011.

[9] Y. Dodis, L. Reyzin, and A. Smith, "Fuzzy extractors: How to generate strong keys from biometrics and other noisy data," in *Proc. Int. Conf. Theory and Applications of Cryptographic Techniques: Advances in Cryptology (EUROCRYPT)*, LNCS, 2004, vol. 3027, pp. 523–540.

[10] F. Farooq, R. Bolle, T. Jea, and N. Ratha, "Anonymous and revocable fingerprint recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2007, pp. 1–7.

[11] F. Gray, "Pulse code communications," U.S. patent 2,632,058, 1953.

[12] A. K. Jain, K. Nandakumar, and A. Nagar, "Biometric template security," *EURASIP J. Adv. Signal Process.*, Jan. 2008, pp. 1–17.

[13] A. Juels and M. Sudan, "A fuzzy vault scheme," in *Proc. IEEE Int. Symp. Information Theory*, 2002, pp. 408.

[14] A. Juels and M. Wattenberg, "A fuzzy commitment scheme," in *Proc. ACM Conf. Computer and Communications Security*, 1999, pp. 28–36.

[15] T. A. M. Kevenaar, G. J. Schrijen, M. van der Veen, A. H. M. Akkermans, and F. Zuo, "Face recognition with renewable and privacy preserving binary templates," in *Proc. IEEE Workshop on Automatic Identification Advanced Technologies*, 2005, pp. 21–26.

[16] A. Kumar and D. Zhang, "Hand geometry recognition using entropy-based discretization," *IEEE Trans. Inform. Forensics Security*, vol. 2, no. 2, pp. 181–187, 2007.

[17] R. L. Lagendijk, Z. Erkin, and M. Barni, "Encrypted signal processing for privacy protection: Conveying the utility of homomorphic encryption and multiparty computation," *IEEE Signal Processing Mag.*, vol. 30, no. 1, pp. 82–105, 2013.

[18] H. Lee, A. B. J. Teoh, H. G. Jung, and J. Kim, "A secure biometric discretization scheme for face template protection," *Future Gen. Comput. Syst.*, vol. 28, no. 1, pp. 218–231, 2012.

[19] Y. J. Lee, K.-R. Park, S. J. Lee, K. Bae, and J. Kim, "A new method for generating an invariant iris private key based on the fuzzy vault system," *IEEE Trans. Syst., Man, Cybern., Part B*, vol. 38, no. 5, pp. 1302–1313, 2008.

[20] M.-H. Lim, A. B. J. Teoh, and K.-A. Toh, "An efficient dynamic reliability-dependent bit allocation for biometric discretization," *Pattern Recognit.*, vol. 45, no. 5, pp. 1960–1971, 2012.

[21] M.-H. Lim and A. B. J. Teoh, "Non-user-specific multivariate biometric discretization with medoid-based segmentation," in *Proc. Chinese Conf. Biometric Recognition*, LNCS, 2011, vol. 7098, pp. 279–287.

[22] M.-H. Lim and A. B. J. Teoh, "A novel encoding scheme for effective biometric discretization: Linearly separable SubCode," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 2, pp. 300–313, 2013.

[23] M.-H. Lim, A. B. J. Teoh, and K.-A. Toh, "Dynamic detection-rate-based bit allocation with genuine interval concealment for binary biometric representation," *IEEE Trans. Cybern., Part B*, vol. 43, no. 3, pp. 843–857, 2013.

[24] A. Nagar, K. Nandakumar, and A. Jain, "Multibiometric cryptosystems based on feature level fusion," *IEEE Trans. Inform. Forensics Security*, vol. 7, no. 1, pp. 255–268, 2011.

[25] A. Nagar, S. Rane, and A. Vetro, "Privacy and security of features extracted from minutiae aggregates," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, 2010, pp. 524–531.

[26] K. Nandakumar, "A fingerprint cryptosystem based on minutiae phase spectrum," in *Proc. IEEE Workshop on Information Forensics and Security*, 2010, pp. 1–6.

[27] K. Nandakumar, A. K. Jain, and S. Pankanti, "Fingerprint-based fuzzy vault: Implementation and performance," *IEEE Trans. Inform. Forensics Security*, vol. 2, no. 4, pp. 744–757, 2007.

[28] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "ORB: An efficient alternative to SIFT or SURF," in *Proc. IEEE Int. Conf. Computer Vision*, 2011, pp. 2564–2571.

[29] Y. Sutcu, S. Rane, J. S. Yedidia, S. C. Draper, and A. Vetro, "Feature extraction for a Slepian-Wolf biometric system using LDPC codes," in *Proc. IEEE Int. Symp. Information Theory*, 2008, pp. 2297–2301.

[30] U. Uludag and J. Anil, "Securing fingerprint template: Fuzzy vault with helper data," in *Proc. Computer Vision and Pattern Recognition Workshop*, 2006, pp. 163–171.

[31] Y. Wang and K. N. Plataniotis, "Fuzzy vault for face based cryptographic key generation," in *Proc. Biometric Symp.*, 2007, pp. 1–6.

[32] Y. Wang, L. Wang, Y.-M. Cheung, and P. C. Yuen, "Fingerprint geometric hashing based on binary minutiae cylinder codes," in *Proc. IEEE Int. Conf. Pattern Recognition*, 2014, pp. 690–695.

[33] H. Xu, R. Veldhuis, T. Kevenaar, A. Akkermans, and A. Bazen, "Spectral minutiae: A fixed-length representation of a minutiae set," in *Proc. IEEE Computer Vision and Pattern Recognition Workshop on Biometrics*, 2008.

[34] X. Xu, L. Tian, J. Feng, and J. Zhou, "OSRI: A rotationally invariant binary descriptor," *IEEE Trans. Image Processing*, vol. 23, no. 7, pp. 2983–2995, 2014.

[SP]

Karthik Nandakumar and Anil K. Jain

Biometric Template Protection

Bridging the performance gap between theory and practice

Biometrics Security and Privacy Protection

Biometric recognition is an integral component of modern identity management and access control systems. Due to the strong and permanent link between individuals and their biometric traits, exposure of enrolled users' biometric information to adversaries can seriously compromise biometric system security and user privacy. Numerous techniques have been proposed for biometric template protection over the last 20 years. While these techniques are theoretically sound, they seldom guarantee the desired noninvertibility, revocability, and nonlinkability properties without significantly degrading the recognition performance. The objective of this work is to analyze the factors contributing to this performance divide and highlight promising research directions to bridge this gap. The design of invariant biometric representations remains a fundamental problem, despite recent attempts to address this issue through feature adaptation schemes. The difficulty in estimating the statistical distribution of biometric features not only hinders the development of better template protection algorithms but also diminishes the ability to quantify the noninvertibility and nonlinkability of existing algorithms. Finally, achieving nonlinkability without the use of external secrets

Digital Object Identifier 10.1109/MSP.2015.2427849
Date of publication: 13 August 2015

1053-5888/15©2015IEEE

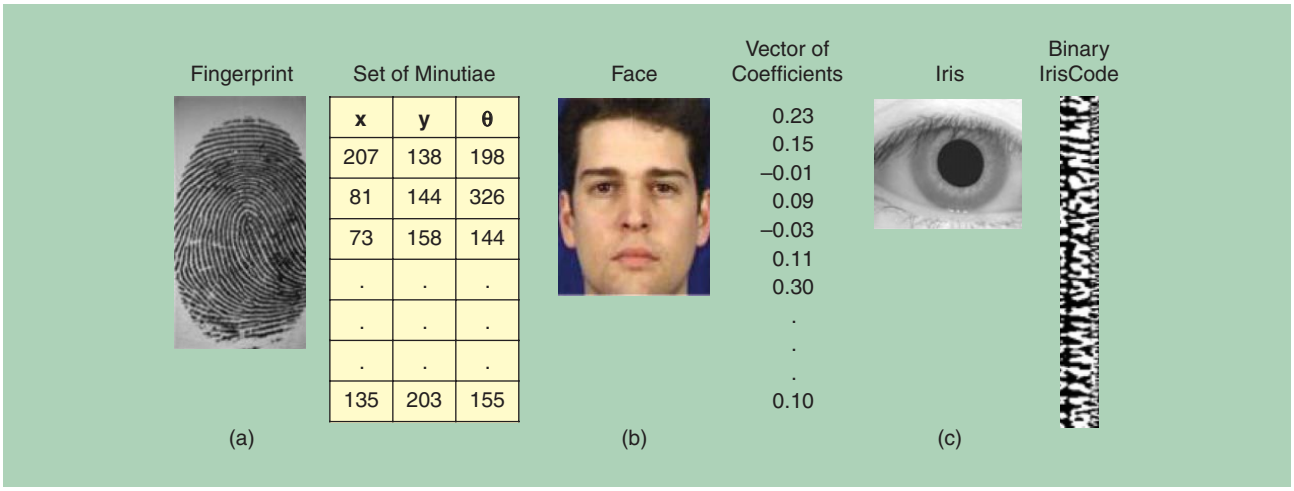

[FIG1] Examples of biometric templates extracted from (a) fingerprint, (b) face, and (c) iris images. A fingerprint image is typically represented as an unordered set of minutiae, which encodes the location (x, y) and orientation (θ) of friction ridge discontinuities. Face images are often represented as a linear combination of basis faces, with the vector of weight coefficients constituting the template. An iris image is usually represented as a fixed-length binary string called the *IrisCode*, which is obtained by binarizing the phase responses of Gabor filters applied to the given image.

(e.g., passwords) continues to be a challenging proposition. Further research on the above issues is required to cross the chasm between theory and practice in biometric template protection.

INTRODUCTION

Biometric recognition, or biometrics, refers to the automated recognition of individuals based on their biological and behavioral characteristics (e.g., face, fingerprint, iris, palm/finger vein, and voice) [1]. While biometrics is the only reliable solution in some applications (e.g., border control, forensics, covert surveillance, and identity deduplication), it competes with or complements traditional authentication mechanisms such as passwords and tokens in applications requiring verification of a claimed identity (e.g., access control, financial transactions, etc.). Though factors such as additional cost and vulnerability to spoof attacks hinder the proliferation of biometric systems in authentication applications, security and privacy concerns related to the storage of biometric templates have been major obstacles [2].

A template is a compact representation of the sensed biometric trait containing salient discriminatory information that is essential for recognizing the person (see Figure 1). Exposure of biometric templates of enrolled users to adversaries can affect the security of biometric systems by enabling presentation of spoofed samples [3] and replay attacks. This threat is compounded by the fact that biometric traits are irreplaceable in nature. Unlike passwords, it is not possible to discard the exposed template and re-enroll the user based on the same trait. Moreover, it is possible to stealthily cross-match templates from different databases and detect whether the same person is enrolled across different unrelated applications. This can severely compromise the privacy of individuals enrolled in biometric systems.

In most operational (deployed) biometric systems, the biometric template is secured by encrypting it using standard encryption techniques such as the Advanced Encryption Standard (AES) and RSA cryptosystem. This approach has two main drawbacks. First, the encrypted template will be secure only as long as the decryption key is unknown to the attacker. Thus, this approach merely shifts the problem from biometric template protection to cryptographic key management, which is equally challenging. Even if the decryption key is secure, the template needs to be decrypted during every authentication attempt because matching cannot be directly performed in the encrypted domain. Consequently, an adversary can glean the biometric template by simply launching an authentication attempt.

One way to address the limitations of the standard encryption approach is to store the encrypted template and decryption key in a secure environment within a smart card or a secure chip [e.g., A8 chip on Apple iPhone6 (http://support.apple.com/en-sg/HT5949), Privaris plusID (http://www.privaris.com/products/indeX.html)], which is in the possession of the user. When biometric matching is performed on the card (or chip), the template never leaves the secure environment. While this solution addresses the security and privacy concerns, it requires the user to carry an additional authentication token (smart card or a mobile device), thereby reducing user convenience and restricting the range of applications. Due to the above limitations of existing solutions, biometric template protection has emerged as one of the critical research areas in biometrics and computer security communities.

BIOMETRIC TEMPLATE PROTECTION REQUIREMENTS

The general framework of a biometric system with template protection is shown in Figure 2. Rather than storing the biometric template in its original form (x), a biometric template protection algorithm generates and stores a protected biometric reference (v) derived from the original template. Note that the term *protected biometric reference* not only includes the protected biometric

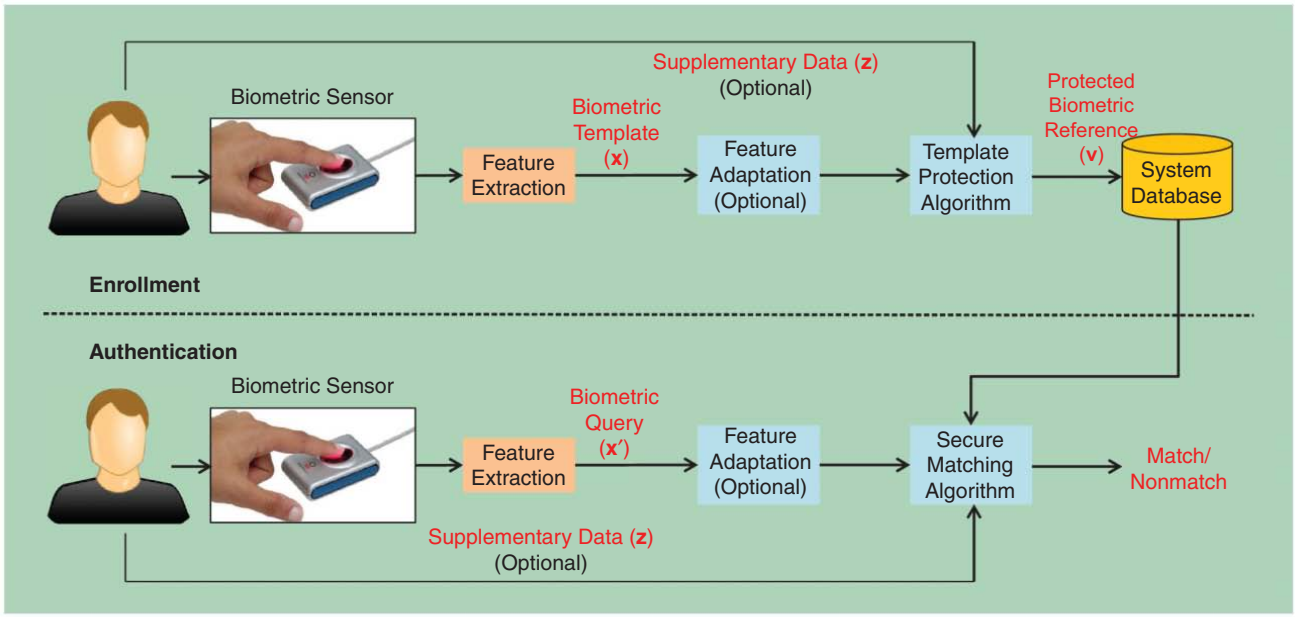

[FIG2] The general framework of a biometric system with template protection.

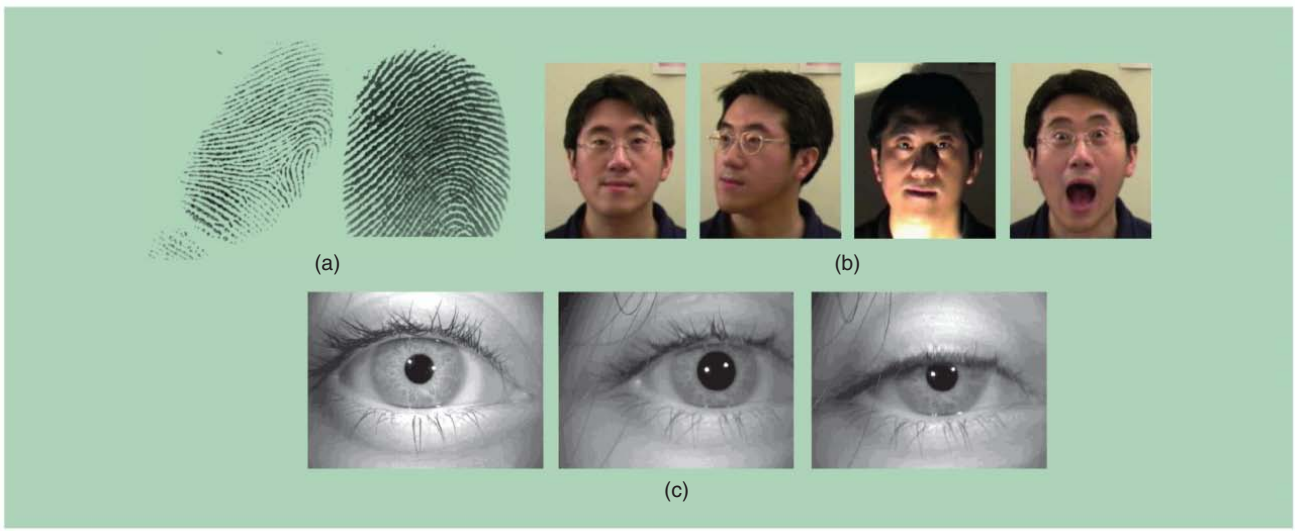

[FIG3] An illustration of intrasubject variations observed in biometric samples. (a) Images of the same finger may exhibit variations in translation, rotation, and nonlinear distortion. (b) Pose, illumination, and facial expression changes may change the appearance of face images obtained from the same person. (c) Iris images of the same eye may exhibit differences due to pupil dilation, partial closure of eyelids, and change in gaze angle.

information, but also other system parameters or values (e.g., cryptographic hashes) that need to be stored, as well as any biometric side information (e.g., information required for alignment, quality of the biometric features, etc.) that directly does not leak information about the user identity. On the other hand, supplementary data ($\mathbf{z}$) refers to entities that are not stored in the database, but are required during both enrollment and authentication. Examples of supplementary data include a password or secret key provided by the user in addition to his biometric trait. The use of supplementary data is optional, but if used, it provides an additional factor of authentication.

Feature adaptation is also an optional step in a template protection scheme. It is well known that biometric samples exhibit intrasubject variations due to various factors like sensor noise, differences in user interaction, environmental changes, and trait aging (see Figure 3). The objective of feature adaptation is to minimize intrasubject variations in the sensed biometric signal and/or represent the original features in a simplified form (e.g., a binary string) without diluting their distinctiveness. It must be emphasized that distinctiveness of a biometric representation is a function of both intrasubject variations and intersubject variations. A highly distinctive representation should

have small intrasubject variations (features extracted from multiple acquisitions of the same biometric trait of a person should be similar), but large intersubject variations (features extracted from the same biometric trait of different individuals should be different). When minimizing intrasubject variations, care must be taken to preserve intersubject variations. Otherwise, distinctiveness of the features may degrade, resulting in lower recognition performance.

In the context of template security, the protected biometric reference (ν) is typically considered as public information that is available to any adversary. Hence, ν should satisfy the following three properties:

- *Noninvertibility or irreversibility*: It should be computationally difficult to obtain the original biometric template from an individual's protected biometric reference. A problem can be considered to be computationally hard or difficult if it cannot be solved using a polynomial-time algorithm. The noninvertibility prevents the abuse of stored biometric data for launching spoof or replay attacks, thereby improving the security of the biometric system.
- *Revocability or renewability*: It should be computationally difficult to obtain the original biometric template from multiple instances of protected biometric reference derived from the same biometric trait of an individual. This makes it possible to revoke and reissue new instances of protected biometric reference when a biometric database is compromised. Moreover, this prevents an adversary from obtaining the original template by compromising multiple biometric databases where the same individual may be enrolled.
- *Nonlinkability or unlinkability*: It should be computationally difficult to ascertain whether two or more instances of protected biometric reference were derived from the same biometric trait of a user. The nonlinkability property prevents cross-matching across different applications, thereby preserving the privacy of the individual.

Apart from satisfying the aforementioned three properties, an ideal template protection algorithm must not degrade the recognition performance of the biometric system. In many applications of biometric recognition, especially those involving millions of enrolled identities (e.g., border crossing and national identity programs), recognition accuracy is of paramount importance. Moreover, issues such as throughput (number of biometric comparisons that can be performed in unit time) and template size must also be considered in real-world applications.

BIOMETRIC TEMPLATE PROTECTION APPROACHES

Numerous template protection techniques have been proposed in the literature with the objective of ensuring noninvertibility, revocability, and nonlinkability without compromising the recognition performance. The ISO/IEC Standard 24745 on Biometric Information Protection provides a general guidance for the protection of biometric information. According to this standard, a protected biometric reference is typically divided into two parts: pseudonymous identifier (PI) and auxiliary data (AD). Depending on how these two components are generated, biometric template protection schemes can be broadly categorized as 1) feature transformation approach and 2) biometric cryptosystems. A detailed review of biometric template protection approaches is beyond the scope of this article; we refer the readers to [4]–[6] for such in-depth analysis.

In the feature transformation approach [see Figure 4(a)], a noninvertible or one-way function is applied to the biometric template ($\mathbf{x}$). While the transformed template is stored in the database as PI, the transformation parameters are stored as AD. During authentication, the AD makes it possible to apply same transformation function to the biometric query ($\mathbf{x}'$) and construct PI′, which is compared to the stored PI. Thus, the biometric matching takes place directly in the transformed domain. Biohashing [7], cancelable biometrics [8], and robust hashing [9] are some of the well-known schemes that can be grouped under feature transformation. Some feature transformation schemes [7] are noninvertible only when the supplementary data (e.g., key or password) is assumed to be a secret. Techniques that can generate noninvertible templates without the need for any secrets (e.g. [8]) are sometimes referred to as *keyless biometric template protection schemes*. Such schemes can be useful in applications (e.g., law enforcement) where it may not be feasible or desirable to allow user-specific supplementary data.

In biometric cryptosystems, the AD is often referred to as a *secure sketch* [see Figure 4(b)], which is typically derived using error correction coding techniques. While the secure sketch in itself is insufficient to reconstruct the original template, it does contain adequate information to recover the original template in the presence of another biometric sample that closely matches with the enrollment sample [10]. The secure sketch is either obtained as the syndrome of an error correction code applied to the biometric template or by binding the biometric template with a error correction codeword that is indexed by a cryptographic key (e.g., fuzzy vault [11] and fuzzy commitment [12]). A cryptographic hash of the original template or the key used to index the error correction codeword is stored as PI. Matching in a biometric cryptosystem is performed indirectly by attempting to recover the original template ($\mathbf{x}$) using the secure sketch (AD) in conjunction with the query biometric features ($\mathbf{x}'$). The recovered template is used to regenerate a new pseudonymous identifier (PI′), which is compared to the stored PI to determine whether the template and query match. Secure sketch constructions have been proposed for various biometric modalities, including fingerprint [13], face [14], and iris [15], [16].

Both template protection approaches have their own strengths and limitations. The primary challenge in the feature transformation approach is finding an appropriate transformation function that provides noninvertibility, but at the same time is tolerant to intrasubject variations [17]. The strength of biometric cryptosystems is the availability of bounds on the information leaked by the secure sketch if we assume that the biometric data distribution is known [10], [18]. On the flip side, most biometric cryptosystems require the features to be represented in standardized data formats like binary strings and point

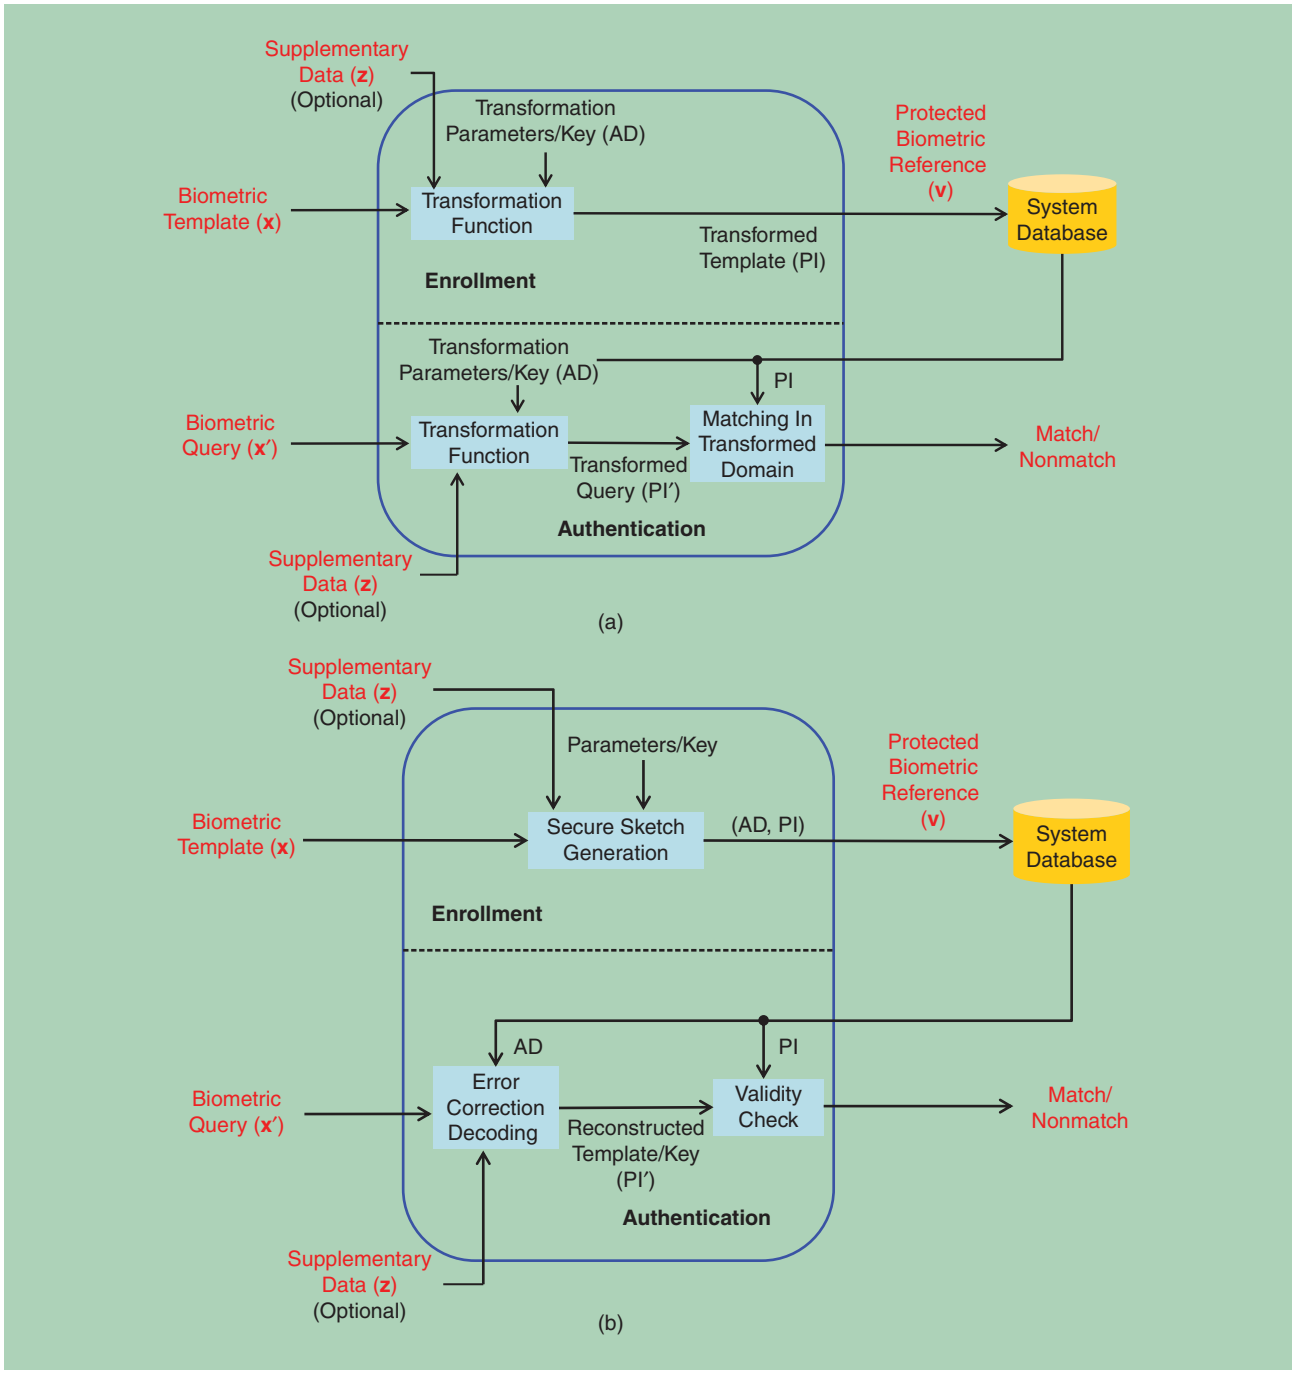

[FIG4] There are two broad approaches for biometric template protection: (a) feature transformation and (b) biometric cryptosystem. The protected biometric reference (denoted by v) generally consists of two distinct parts: pseudonymous identifier (PI) and auxiliary data (AD).

sets, which often leads to loss of discriminatory information and consequent degradation in recognition accuracy. Due to the properties of linear error correction codes that are commonly used in secure sketch constructions, it is difficult to achieve nonlinkability in biometric cryptosystems. In a linear error correcting code, any linear combination of codewords is also a codeword. Consequently, if two secure sketches are derived from the biometric data of the same user using different codewords, a suitable linear combination of these two sketches is highly likely to result in a decodable codeword. This paves the way for verifying whether the two secure sketches belong to the same user, thereby making them linkable.

One way to overcome the aforementioned limitations is to apply a feature transformation function to the biometric template before it is protected using a biometric cryptosystem. Since this involves both feature transformation and secure

sketch generation, such systems are known as hybrid biometric cryptosystems [19], [20]. Another promising approach is secure computation based on homomorphic encryption. While this approach offers the attractive proposition of performing biometric matching directly in the encrypted domain, it typically comes at the cost of a significant increase in the computational burden and communication overhead [21].

THE GAP BETWEEN THEORY AND PRACTICE

Most of the existing techniques do not satisfy the desired template protection requirements in practice. As an example, consider the results published by the ongoing Fingerprint Verification Competition (FVC-onGoing); see https://biolab.csr.unibo.it/fvcongoing/UI/Form/Home.aspx for more information. Six algorithms were able to achieve an equal error rate (EER) of less than 0.3% on the FVC-STD-1.0 benchmark data set when operating without any template protection. On the other hand, the lowest EER achieved by a fingerprint verification system with template protection on the same data set was 1.54%, which is more than five times higher. Reduction in accuracy was also observed during independent testing of template protection algorithms in [22].

Even if we assume that a small degradation in the recognition performance is acceptable in some applications, it is imperative to precisely quantify (in terms of bits) the noninvertibility and nonlinkability of the protected biometric reference. This is necessary to benchmark the utility of a biometric template protection scheme. In cryptography, "security strength" (the measure of the computational effort required to break a cryptosystem using the most efficient known attack) is one of the metrics used to compare different cryptosystems. It is well known that an AES system with a 128-bit key or a RSA cryptosystem with a 3,072-bit key can provide a security strength of approximately 128 bits [51]. However, there is no consensus within the biometrics community on analogous metrics that can be used to measure the noninvertibility, revocability, and nonlinkability properties of biometric template protection algorithms as well as the methods to compute these metrics [23]. Consequently, practical template protection schemes neither have proven noninvertibility/nonlinkability guarantees nor do they achieve satisfactory recognition performance. This explains why despite 20 years of research, operational biometric systems do not go beyond encrypting the template using standard encryption techniques and/or storing them in secure hardware.

The gap between theory and practice of template protection can be attributed to three main reasons:

1) The template protection schemes generally require the use of simple distance metrics such as Hamming distance or a measure of set difference to compute the similarity between biometric features [10]. Consequently, the burden of handling intrasubject variations observed in the biometric samples shifts completely to the feature extraction stage. Thus, the foremost challenge in biometric template protection is the design of feature extractors, which not only need to extract highly robust and distinctive features, but also represent them in a simplified form (e.g., a fixed-length binary string) that is suitable for applying the template protection construct.

2) Template protection techniques typically result in a trade-off between noninvertibility and recognition performance [17], [24] due to the following reason. Maximizing noninvertibility implies that the protected biometric reference should leak as little information about the original template as possible. However, high recognition performance can be achieved only when the protected biometric reference retains all the discriminatory information contained in the original template. This conundrum can only be solved by understanding the statistical distribution of biometric features and designing template protection schemes that are appropriate for the underlying feature distribution. For example, it is well known that bits in an IrisCode [25] or the minutiae locations in a fingerprint [26] are neither independent nor do they follow a uniformly random distribution. This inherent redundancy in the biometric features could be exploited to handle intrasubject variations without compromising on intersubject variations. In many biometric cryptosystems, the template is protected by adding noise to the true biometric information. In this case, knowledge of the feature distribution could be useful in selecting the appropriate noise distribution. Modeling the biometric feature distribution is also required for obtaining realistic estimates for the noninvertibility and nonlinkability of a protected biometric reference. If the biometric feature distribution is known, it may be possible to formulate biometric template protection as an optimization problem and systematically find solutions that maximize both recognition performance and noninvertibility. Thus, knowledge of the statistical distribution of biometric features is beneficial for biometric template protection. However, estimating the feature distributions is a challenging task.

3) Compared to the issue of noninvertibility, the problem of ensuring nonlinkability and revocability of protected biometric reference has not been adequately addressed in the literature. While many template protection constructs claim to provide nonlinkability and revocability, a deeper analysis indicates that this is often achievable only with the involvement of an additional authentication factor (supplementary data) such as a password or secret key [27].

The primary contribution of this article is to provide an in-depth analysis of the three previously mentioned challenges, discuss some of the solutions that have been proposed to overcome them, and identify unresolved issues that require further research.

DESIGNING INVARIANT FEATURE REPRESENTATIONS

A traditional biometric system accounts for intrasubject variations in two ways. First, the feature extraction algorithm attempts to extract an invariant representation from the noisy biometric samples. Second, the matching algorithm is designed to further suppress the effect of intrasubject variations and focus only on features that are distinctive across individuals. Consider the example of a fingerprint recognition system (see Figure 5). An accurate fingerprint matcher not only handles missing and spurious minutiae, but also other intrasubject

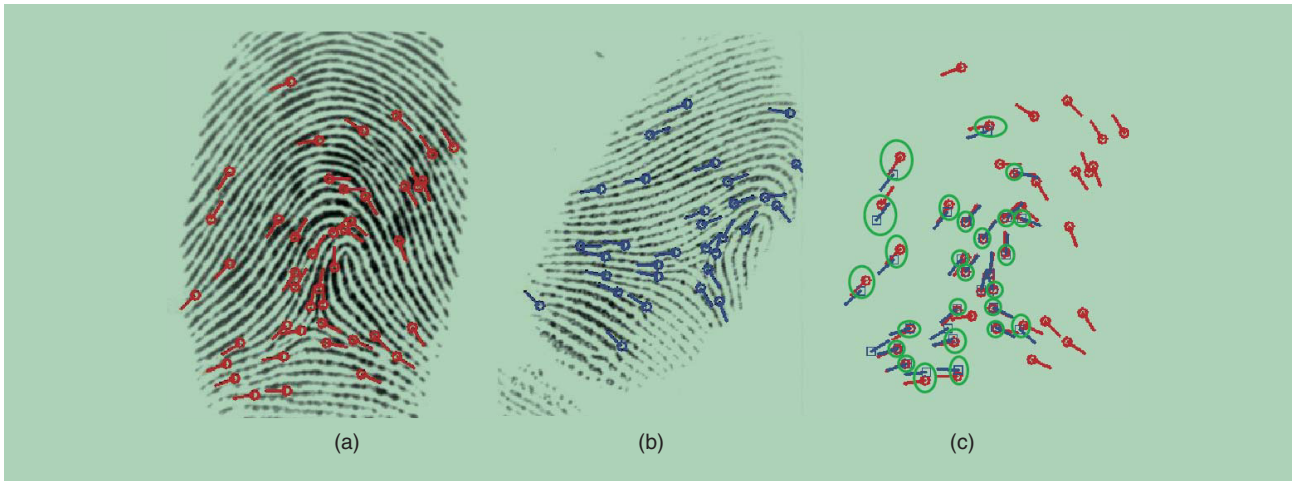

[FIG5] Complexity in fingerprint minutiae matching. (a) and (b) are two fingerprint images from the same finger with minutiae features marked on them. The two minutiae sets after global alignment are shown in (c). Apart from missing and spurious minutiae that can be captured well using the set difference metric, one can observe that the matching minutiae (marked by green ellipses) are not perfectly aligned due to nonlinear distortion. This explains why a simple set difference metric is unlikely to provide accurate recognition.

variations like rotation, translation, and nonlinear distortion [see Figure 5(c)]. When this matcher is replaced by a simple set difference metric (that accounts for only missing and spurious minutiae), it becomes imperative to represent the extracted minutiae in a form that is invariant to rotation, translation, and nonlinear distortion without affecting their distinctiveness. Failure to do so will naturally lead to significant degradation in the recognition performance.

Even in the case of iris recognition, it is not possible to achieve good recognition performance by directly computing the Hamming distance between two IrisCodes. Practical iris recognition systems compute normalized Hamming distance (that ignores bit locations erased by noise) over multiple cyclical shifts applied to one of the IrisCodes (to account for rotation variations). If this practical subtlety is ignored and a simple Hamming distance metric is enforced, the iris recognition accuracy is likely to decrease substantially.

Rather than developing new invariant feature extractors, which in itself is one of the fundamental problems in biometric recognition, researchers working on biometric template protection often implement a feature adaptation step on top of the original feature extractor. It must be emphasized that feature adaptation is not the same as feature transformation. In feature transformation, the goal is to obtain a noninvertible and revocable template. In contrast, adapted templates need not satisfy the noninvertibility and revocability properties. Instead, feature adaptation schemes are designed to satisfy one or more of the following three objectives: 1) minimize intrasubject variations without diluting their distinctiveness, 2) represent the original features in a simplified form, and 3) avoid the need for biometric side information (e.g., alignment parameters). While a feature transformation scheme may employ feature adaptation in the process of securing the template, the converse is not true.

The simplest and most common feature adaptation strategy is quantization and reliable component (feature) selection. The quantization of Gabor phase responses to generate a binary IrisCode and selection of reliable bits within an IrisCode [28] is a good illustration of this adaptation strategy. Another typical example is the quantization of fingerprint minutiae location and orientation features and selection of good quality minutiae [13] when designing a fingerprint cryptosystem. Though the process of quantization and feature selection reduces intrasubject variations, it is also likely to decrease intersubject variations. Thus, the challenge is to strike an optimum balance between reducing intrasubject variations and preserving intersubject variations. Moreover, if quantization and reliable component selection is user specific, the quantization parameters and selected components need to be stored as AD, which is likely to decrease the noninvertibility and nonlinkability of the protected biometric reference [29].

Other strategies for feature adaptation include biometric embedding and alignment-free representation. In biometric embedding, the goal is to obtain a new representation for the given biometric features so that simple distance metrics (e.g., Hamming distance or set difference) can be used to compare biometric samples in the modified representation space. Conversion of a real/complex vector or point set into a fixed-length binary string is an example of biometric embedding. On the other hand, the objective of an alignment-free representation is to generate templates that can be directly matched without the need for any alignment parameters. Such a need often arises when dealing with biometric traits like fingerprint and palmprint. Many practical feature adaptation schemes involve a combination of different adaptation strategies. For instance, quantization and feature selection are often applied in conjunction with biometric embedding or alignment-free representation to obtain the adapted features. Similarly, some alignment-free representations proposed in the literature also perform embedding in a new feature space.

BIOMETRIC EMBEDDING

Biometric embedding algorithms can be classified based on their input and output representations. Two types of embedding algorithms

that are commonly used for biometric feature adaptation are: 1) real vector into a binary string and 2) point set into a binary string.

REAL VECTOR TO BINARY STRING

Conversion of a real vector into a binary string involves two essential steps: 1) quantization-mapping continuous values into discrete values and 2) encoding the discrete values as bits. The critical parameters in quantization are the number of quantization levels and the quantization intervals. The detection rate optimized bit allocation (DROBA) scheme [30] proposes an adaptive bit allocation strategy, where the total number of bits in the binary string is fixed and the number of bits allocated to each feature dimension is varied based on the feature distinctiveness. Specifically, a higher number of bits (i.e., more levels of quantization) is allocated to a particular feature dimension if the mean feature value of that subject is very different from the population mean. Furthermore, this scheme advocates the use of equal-probability quantization intervals to maximize the entropy of the resulting binary string. While the DROBA approach optimizes the detection rate (genuine accept rate) at the minimum (low) FAR, it requires many training samples per subject to determine user-specific feature statistics. Furthermore, the need for storing user-specific quantization information increases information leakage when the resulting binary string is eventually secured using a template protection scheme [29].

While the DROBA scheme focuses on the quantization step, the linearly separable subcodes (LSSC) method attempts to develop a better encoding scheme for encoding the discrete values as bits. The gray coding scheme, which is traditionally used for binary encoding, maps the discrete values into bits such that adjacent quantization levels differ only by a single bit. The problem with the gray code approach is that it does not preserve the distances between the samples after encoding. Though the Hamming distances between genuine samples is likely to remain small (because feature values of two samples from the same subject can be expected to be similar), it is possible that two dissimilar feature values may also have a small Hamming distance. Consequently, the recognition performance based on the resulting binary string will degrade significantly. A unary coding scheme solves this problem, but it does not produce a compact representation. The LSSC method attempts to generalize the idea of unary coding. A partially linearly separable subcode was also proposed in [31] to obtain a better compromise between compactness and distance preservation.

POINT SET TO BINARY STRING

The most well-known example of point-set-based biometric representation is a collection of fingerprint minutia. Techniques for converting unordered point sets (especially fingerprint minutiae) into fixed-length binary strings include local point aggregates [32] and spectral minutiae [33]. In the local aggregates approach [32], the fingerprint region is divided into a fixed number of randomized local regions (could be overlapping) and aggregate features are computed based on the minutiae falling within each local region. The resulting feature vector is then converted into a binary string using the techniques described in the section "Real Vector to Binary String." The main limitation of this approach is that it requires the fingerprints to be aligned before feature adaptation.

The spectral minutiae representation is obtained by considering the minutiae set as a collection of two-dimensional Dirac-delta functions and obtaining its Fourier spectrum after low-pass filtering [33]. Only the magnitude spectrum is considered and it is sampled on a log polar grid to obtain a fixed-length vector. Theoretically, the magnitude spectrum is invariant to rotation and translation due to the shift, scale, and rotation properties of the Fourier transform. Hence, it is possible to perform matching between two spectral minutiae vectors without aligning them first. However, in practice, alignment based on singular points (core and delta) is required to achieve good recognition performance [33] because large rotation or translation may lead to partial overlap between different impressions of the same finger. Another variation of the spectral minutiae approach is the binarized phase spectrum representation [34], where the phase spectrum is considered instead of the magnitude spectrum (see Figure 6). However, this approach also requires prior fingerprint alignment.

ALIGNMENT-FREE REPRESENTATION

A possible solution to the problem of fingerprint alignment is the use of local minutiae structures, which consist of features that characterize the relative information between two or more minutiae (e.g., distance between two minutiae) [35]. Since such features are relative, they are invariant to global rotation and translation of the fingerprint and hence, no a priori alignment is needed before matching. An additional benefit is that such features are robust to nonlinear distortion. However, if the matching is based only on the local minutiae information and the global

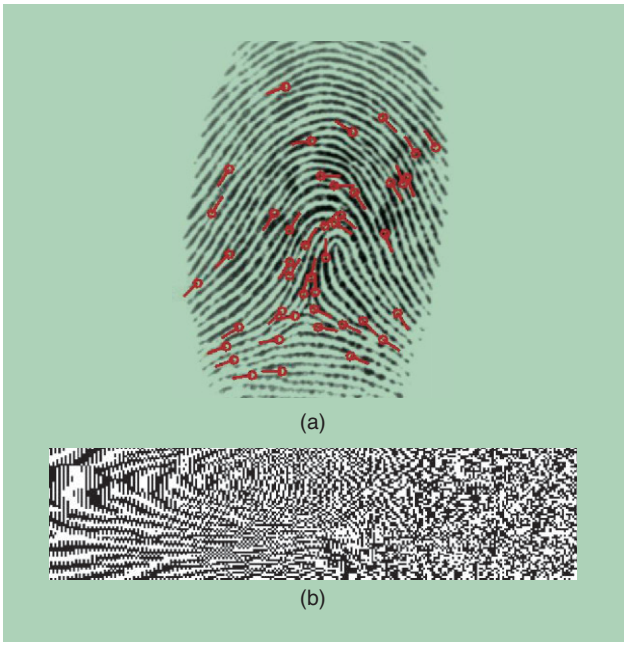

[FIG6] **An example of embedding a point set as a binary string. (a) Fingerprint with minutiae (point set) marked on it and (b) the corresponding binary string representation obtained using the binarized phase spectrum technique [34].**

spatial relationships between minutiae are ignored, some degradation in the recognition accuracy may occur.

The simplest local minutiae structure is based on minutia pairs, where the distance between the pair and the orientation of each minutia with respect to the line connecting them can be used as the invariant attributes [19]. The most commonly used local minutiae structure is the minutia triplet, where relative features (distances and angles) are computed from combinations of three minutiae. Rather than defining the local neighborhood based on a fixed number of minutiae, it is also possible to construct a local descriptor by considering all minutiae that fall within a fixed radius of a minutia point. An example of this latter approach is the minutia cylinder code (MCC) [35]. The MCC is obtained by dividing the cylindrical region (with its axis along the minutia orientation) around each minutia into a finite number of cells and encoding the likelihood of another minutia in the fingerprint with a specific angular difference from the reference minutia being present in the specific cell. It is also possible to binarize the MCC to get a fixed-length binary string describing each minutia point.

OPEN ISSUES IN FEATURE ADAPTATION

Though a significant amount of research effort has been devoted toward feature adaptation, three main issues remain unresolved. First, existing feature adaptation techniques invariably result in loss of some discriminatory information leading to lower recognition performance. A possible reason for this phenomenon is that most of these techniques focus only on minimizing intrasubject variations while ignoring the need to preserve intersubject variations. Hence, there is a strong need for distance-preserving feature adaptation strategies.

The second unresolved issue is the coupling between the feature adaptation strategy and the template protection technique. Recall that the main objective of feature adaptation is to generate an invariant representation that can be easily secured using an existing template protection scheme. Therefore, it is essential to carefully consider the requirements of the template protection scheme while designing the feature adaptation strategy. For instance, the error correction scheme used in a biometric cryptosystem may have the ability to correct a limited amount of errors. Since this error correction capability implicitly determines the system threshold, the feature adaptation scheme must be designed such that the number of errors between samples of the same user fall below this threshold, while the number of errors encountered during impostor comparisons is greater than the error correction capability. A feature adaptation scheme that is designed in isolation may not satisfy the above requirement. Alternatively, one can argue that it may be better to design a biometric template protection scheme that directly secures the template in its original representation rather than attempting to adapt the template to fit the template protection scheme. As an illustration, suppose that we wish to protect a biometric template represented as a real vector. This template can be protected either by converting it into a binary string and applying a fuzzy commitment scheme [12] to the binary template or by directly applying a secure sketch designed for the continuous domain [36]. It is not clear which of the these two strategies will lead to a better outcome.

Finally, the statistical properties of the adapted features is seldom given attention in the design of a feature adaptation scheme. For example, consider the case of a feature adaptation scheme generating a binary string as output. Apart from having low intrasubject variations and high distinctiveness, it would be ideal if the resulting binary string is uniformly random (i.e., has high entropy). Such a representation is likely to have better noninvertibility properties when it is eventually secured using a biometric cryptosystem (cf. the section "Solving the Recognition Performance Versus Security Conundrum"). Moreover, one of the implicit benefits of feature adaptation could be a new representation that makes it easier to characterize the statistical distribution of biometric features. However, the design of such feature adaptation strategies is still an open research problem.

SOLVING THE RECOGNITION PERFORMANCE VERSUS SECURITY CONUNDRUM

The main limitation of state-of-the-art template protection techniques is the trade-off between recognition performance and the level of security offered by them. The first step toward solving this problem is to clearly define the notion of security, establish metrics to quantify security properties such as noninvertibility and nonlinkability, and develop methodologies to compute such metrics. Once this is achieved, algorithms need to be developed to jointly maximize performance and security.

The lack of a well-accepted notion of security is a critical lacuna in the area of template protection. It is important to emphasize that a biometric template protection scheme is not designed to prevent other adversary attacks on a biometric system such as spoofing or zero-effort impostor attack. Therefore, the vulnerability of a biometric system to such attacks cannot be considered as the sole basis for evaluating a template protection scheme. For instance, a false acceptance rate (FAR) of 0.01% implies that 1 in 10,000 zero-effort impostor attempts is likely to succeed. At this FAR, it is possible to argue that the noninvertibility of a template protection scheme can be no more than $\log_2(10^4)$ bits because, on average, only 10,000 attempts would be required to find a biometric sample that closely matches with the stored template. However, such an argument is unfair since it is based on the assumption that an attacker has access to a large biometric database and is able to mount an off-line zero effort impostor attack. [Since most practical biometric systems restrict the number of failed authentication attempts, it is usually not possible to mount online zero-effort impostor (FAR) attacks.] Therefore, it may be better to consider vulnerability to zero-effort attacks as a distinct threat and report the FAR of the biometric system before and after the application of biometric template protection. Ideally, the FAR should be included as part of the recognition performance and not security analysis. Furthermore, the FAR of the biometric system after template protection should be reported based on the assumption that the attacker has full knowledge about the system, including access to any supplementary data (if used).

In the context of biometric template protection, the terms *security* and *privacy* have been used ambiguously in the literature. One of the reasons for this ambiguity is that many biometric cryptosystems are motivated by the desire to generate a cryptographic key from the biometric data or securely bind a key together with the biometric data. Template protection is only a by-product of this key generation/binding process. Therefore, in biometric cryptosystems, security is often defined in terms of the *secret key rate*, which measures the amount of randomness in the key bound to the template or extracted from the biometric data [18], [37]. While the term *privacy leakage* is commonly used in biometric cryptosystems as a proxy for measuring noninvertibility, one can find instances where the term *privacy* actually refers to nonlinkability. To further complicate matters, notions such as weak (or conditional) and strong (or unconditional) biometric privacy have been proposed [38]. Here, weak biometric privacy refers to noninvertibility given only the protected biometric reference $\mathbf{v}$, whereas strong biometric privacy refers to noninvertibility given $\mathbf{v}$ and the associated cryptographic key (one that is bound to the template or extracted from the biometric data). In the literature on feature transformation, the terms *security* and *privacy* typically refer to *noninvertibility* and *nonlinkability*, respectively. To avoid confusion, it has been suggested that specific properties such as noninvertibility (or irreversibility) and nonlinkability (or unlinkability) must be considered instead of employing generic terms like security and privacy [39].

METRICS FOR MEASURING NONINVERTIBILITY

Noninvertibility refers to the difficulty in obtaining (either exactly or within a small margin of error) the original biometric template from an individual's protected biometric reference. This is also referred to as full-leakage irreversibility in [39]. A number of metrics have been proposed in the literature to measure noninvertibility of a protected biometric reference.

A direct measure of noninvertibility is the conditional Shannon entropy of the original template $\mathbf{x}$ given the protected biometric reference $\mathbf{v}$, i.e., $H(\mathbf{x}\,|\,\mathbf{v})$. This quantity measures the average uncertainty in estimating $\mathbf{x}$ given the knowledge of $\mathbf{v}$. Note that $H(\mathbf{x}\,|\,\mathbf{v}) = H(\mathbf{x}) - I(\mathbf{x};\mathbf{v})$, where $H(\mathbf{x})$ is the entropy of the unprotected template $\mathbf{x}$ and $I(\mathbf{x};\mathbf{v})$ is the mutual information between $\mathbf{x}$ and $\mathbf{v}$. In the literature, a normalized quantity called the *privacy leakage rate* [37], which can be expressed as $H(\mathbf{x}\,|\,\mathbf{v})/H(\mathbf{x})$, has also been proposed to measure noninvertibility.

In the context of biometric cryptosystems, $I(\mathbf{x};\mathbf{v})$ is also referred to as *entropy loss*, which measures the amount of information leaked by the secure sketch about the biometric template. Entropy loss is a useful measure to compare multiple template protection schemes applied to the same biometric data. In this scenario, since $H(\mathbf{x})$ is constant, the scheme with a lower entropy loss should be preferred because it will lead to larger $H(\mathbf{x}\,|\,\mathbf{v})$. Furthermore, when the secure sketch is obtained by binding the template with a secret cryptographic key $(\mathbf{K})$, it is also important to consider $H(\mathbf{K}\,|\,\mathbf{v})$. Many biometric cryptosystems (e.g., fuzzy vault and fuzzy commitment) do not offer strong biometric privacy [40] in the sense that it is trivial to recover the original biometric template given $\mathbf{K}$ and $\mathbf{v}$. In such cases, the noninvertibility should be defined as the minimum of $H(\mathbf{K}\,|\,\mathbf{v})$ and $H(\mathbf{x}\,|\,\mathbf{v})$.

While the conditional Shannon entropy is a good measure of the average difficulty in inverting a protected biometric reference, researchers have also proposed the use of min-entropy [10] to account for the worst case scenario. For a discrete random variable A with probability mass function P, Shannon entropy is defined as $H(A) = \mathbb{E}_a(-\log_2(P(A=a)))$ and min-entropy is defined as $H_\infty(A) = (-\log_2(\max_a P(A=a)))$. Thus, min-entropy measures the uncertainty in predicting the most likely value of a discrete random variable. The conditional min-entropy is defined as $\tilde{H}_\infty(A\,|\,B) = -\log(\mathbb{E}_{b\to B}[2^{-H_\infty(A|B=b)}])$ and the corresponding entropy loss is computed as $H_\infty(A) - \tilde{H}_\infty(A\,|\,B)$.

In the case of feature transformation, it is difficult to theoretically measure the entropy loss introduced by the transformation scheme. Consequently, the noninvertibility of feature transformation schemes is typically measured empirically based on the computational complexity of the best-known template inversion attack. In particular, the coverage-effort curve [17] was proposed to analyze the noninvertibility of transformed templates. The coverage-effort (CE) curve measures the number of guesses (effort) required to recover a fraction (coverage) of the original biometric data. This measure is analogous to the normalized privacy leakage rate [37] defined earlier. The main pitfall of such empirical measures is that it is impossible to guarantee that the attacker cannot come up with a better template inversion strategy than what is known to the system designer.

Recall that one of the goals of biometric template protection is to prevent the attacker from launching spoof and replay attacks using the compromised template. To launch such attacks, it may not be necessary to exactly recover the original template from the protected biometric reference. Instead, it is sufficient for the attacker to obtain a close approximation (also known as a *preimage*), which can be replayed to the system to gain illegitimate access. Note that in a biometric cryptosystem, it is often straightforward to recover the original template if a close approximation of this template is available. Thus, the vulnerability of a biometric cryptosystem to preimage attacks is already factored into the noninvertibility analysis of such a system. Therefore, analysis of preimage attacks may be valid only for the feature transformation approach. Metrics to evaluate the difficulty in carrying out such attacks have been discussed in [17], [39], and [41]. However, for the sake of simplicity, we avoid a detailed discussion of these metrics in this article.

METHODS FOR COMPUTING NONINVERTIBILITY METRICS

Since the noninvertibility metrics for feature transformation schemes are generally computed empirically, this section will focus only on methods to compute the noninvertibility metrics for biometric cryptosystems. While the metrics for measuring noninvertibility discussed earlier are theoretically sound, they are not easy to compute for an arbitrary biometric template protection scheme. In most biometric cryptosystems, the inherent properties of the underlying error correction technique can be used to

establish upper bounds on the entropy loss [10], [18], [37], [40]. Typically, the entropy loss is an increasing function of the error correction capability of the system. In other words, if larger tolerance for intrasubject variations is desired, the entropy loss will be higher. Consequently, the resulting protected biometric references will leak more information about the original template. Since the above bounds are usually derived based on simplifying assumptions about the biometric feature distribution, their utility will depend on the extent to which the given biometric features conform to these assumptions. Even when a reliable estimate for the entropy loss is available, it is still difficult to directly compute $H(\mathrm{x}|\mathrm{v})$. This is because of the complexity in estimating the entropy of biometric features $(H(\mathrm{x}))$.

BIOMETRIC ENTROPY ESTIMATION

The primary difficulty in estimating the entropy of biometric features is the lack of statistical models to accurately characterize the intra- and intersubject variations. A few attempts have been made in the literature to characterize the distribution of minutiae points in a fingerprint [26], [42]. However, these models were proposed in the context of estimating fingerprint individuality. (More precisely, the goal in [26] and [42] is to estimate the probability of a false correspondence/match between minutiae templates from two arbitrary fingerprints belonging to different fingers.) Moreover, they rely on some simplifying assumptions to keep the problem tractable. Therefore, such models cannot be directly used to infer the entropy of a fingerprint minutiae template.

Entropy of a biometric template can be estimated by computing the relative entropy (also known as *Kullback-Leibler divergence*) between the feature distributions of a specific user and the feature distribution of the population as a whole [43]. This quantity measures the reduction in uncertainty about the identity of the user due to the knowledge of his/her biometric feature measurements. The average relative entropy among all the users enrolled in the system can be used as an estimate of the biometric feature entropy. However, the main drawback of the work in [43] is the use of a simple Gaussian model to characterize the feature distributions, which does not hold true for most biometric modalities.

An alternative to modeling the complex feature distributions is to compute the entropy based on match score distributions. A good example of estimating entropy based on match scores is the analysis of impostor score distribution using IrisCodes extracted from 632,500 different iris images [44]. Based on this approach, it has been estimated that a 2,048-bit IrisCode representation contains approximately 249 degrees of freedom. However, this result is based on a simple matching model that ignores the need to test multiple relative rotations of the IrisCode. Therefore, one cannot directly conclude that the entropy of an IrisCode template is 249 bits. Moreover, it is not straightforward to obtain a precise estimate of individuality of the IrisCode representation using the above result because it fails to take into account the genuine score distribution (consequently, intrasubject variations are not modeled). A simple extension of the above approach is to measure the relative entropy between genuine and impostor match score distributions [45]. But this approach may grossly underestimate the entropy of the biometric features and the resulting entropy estimates should be considered as a very loose lower bound.

OPEN ISSUES IN NONINVERTIBILITY ANALYSIS

Despite significant progress in analyzing the noninvertibility of template protection schemes, there is no consensus yet on the standard metrics to be used for measuring noninvertibility and well-defined methodologies to compute these metrics. Efforts to standardize these metrics are still in progress [23]. Once such metrics are standardized, the focus should shift toward the development of a suitable framework that allows joint optimization of recognition performance and noninvertibility for both feature transformation schemes and biometric cryptosystems.

One way to overcome the inherent tradeoff between noninvertibility and recognition performance is to develop techniques for multibiometric template protection. Multibiometric systems accumulate evidence from more than one biometric identifier (multiple traits like fingerprint and iris or multiple fingers/ irides) to recognize a person. It is well known that multibiometric systems lead to a significant improvement in the recognition performance. When multiple templates are secured together as a single construct, the inherent entropy of the template is also likely to higher, thereby leading to stronger noninvertibility. While a few solutions have been proposed recently for multibiometric cryptosystems [46], the fundamental challenge lies in overcoming the compatibility issues between different biometric templates and generating a combined multibiometric template from different modalities, which preserves the distinctiveness of individual templates. The advancements in the area of feature adaptation can also play a key role in overcoming the aforementioned challenge.

ACHIEVING REVOCABILITY AND NONLINKABILITY

While revocability and nonlinkability are also core requirements of a template protection scheme, the analysis of these two properties has received considerably less attention in the literature compared to noninvertibility. Recently, it has been demonstrated that many well-known biometric cryptosystems do not generate revocable or nonlinkable templates [24], [27], [47], [48]. Though feature transformation schemes are widely proclaimed as "cancelable biometrics" in acknowledgment of their strengths in achieving revocability and nonlinkability, the real capability of such schemes to guarantee these two properties is still questionable. If we assume that the attacker has full knowledge of the protected biometric reference and any supplementary data involved, the revocability and nonlinkability of feature transformation schemes appear to depend on the difficulty in obtaining a preimage of the transformed template. When the preimage is easy to compute given the transformation parameters and the transformed template, it may be possible to correlate the preimages obtained from multiple transformed templates to invert and/or link them [17]. Therefore, there is a critical need to develop one-way transformation functions that do not allow easy computation of a preimage.

One possible way to achieve revocability and nonlinkability is to use hybrid biometric cryptosystems [19], [20]. While a combination of secure sketch and feature transformation enhances the noninvertibility of the protected biometric reference, the feature transformation step ensures the revocability and nonlinkability properties. However, this may come at the cost of a degradation in the recognition performance.

Another practical solution for achieving revocability and nonlinkability is the use of two- or three-factor authentication protocols. In such protocols, either the supplementary data is assumed to be a secret [7] or the AD and PI are not stored together to prevent the possibility that both AD and PI are compromised simultaneously [49]. For example, the transformation parameters in a feature transformation scheme can be dynamically generated based on a password or personal identification number supplied by the user or derived based on a key stored on a smart card held securely by the user. Similarly, the AD in a biometric cryptosystem could be stored on a smart card, while the PI is stored in a central database. Apart from ensuring revocability and nonlinkability, an additional advantage of such protocols is improved robustness against zero-effort impostor (FAR) attacks because the attacker must be able to obtain more than one authentication factor (biometrics and password or biometrics and smart card) for successful authentication. However, if we assume that all the other factors except the biometric trait is available to the attacker, the advantages of such multifactor authentication protocols vanish, and their properties are no better than those of the underlying template protection scheme.

SUMMARY AND FUTURE RESEARCH DIRECTIONS

While biometric template protection has been an active research topic over the last 20 years, existing solutions are still far from gaining practical acceptance. The key reason for this failure is the unacceptable degradation in the recognition performance combined with unprovable security claims. In this article, we have identified three main issues that must be addressed to bridge this gap. Designing invariant biometric representations with high entropy will not only improve the recognition performance, but also enhance the noninvertibility of the protected template. This is because the information leaked by a protected template is often proportional to the tolerance allowed to account for intrasubject variations. Furthermore, standardized metrics are required for measuring the security properties of a template protection scheme, especially noninvertibility. Systematic formulation of such metrics and methodologies to compute them, followed by independent benchmarking of template protection algorithms based on these metrics, will greatly enhance the public confidence in biometric template protection technologies. Finally, practical solutions must be devised to ensure revocability and nonlinkability of protected templates.

Apart from the open research issues identified earlier in the context of feature adaptation (cf. the section "Open Issues in Feature Adaptation") and noninvertibility analysis (cf. the section "Open Issues in Noninvertibility Analysis"), a number of other questions remain unanswered.

- There is a greater need for template security in scenarios where the biometric data is stored in centralized repositories. Such databases are commonplace in large-scale identification systems [e.g., India's Aadhaar program, Office of Biometric Identity Management (formerly the US-VISIT) program]. However, almost all existing template protection techniques have been designed for the authentication use-case (one-to-one verification) as opposed to identification (one-to-many matching). It is not clear if such techniques can be scaled up to meet the requirements of an identification system, especially given the stringent constraints on accuracy and throughput in such applications.
- Another lacuna in template security is the absence of an entity similar to public key infrastructure, which can create, manage, and revoke biometric information [50]. A related issue is how to revoke and reissue a protected biometric reference without re-enrolling the user, which is often impractical.
- An attack on the template is just one of the possible adversarial attacks on a biometric system [4]. It is possible that efforts to secure the template may have a direct impact on other types of attacks [6]. Therefore, a system-level analysis of the effect of template protection algorithms is required.
- Finally, smartphones are turning out to be the preferred platform for the integration of biometric technologies. For example, the Touch-ID fingerprint recognition system in Apple's iPhone 6 enables phone unlocking capability as well as mobile payments via the Apple Pay service. In the near future, it may be possible to capture face, fingerprint, iris, and voice biometric modalities using a commodity smartphone. The ability to securely authenticate a smartphone user using multibiometrics can be expected to open up a number of new applications involving mobile commerce and transactions. In this context, it is necessary to review whether the current state of the art (storing the encrypted biometric templates on a secure chip) is adequate for the range of applications envisioned and develop novel template protection strategies as well as remote biometric authentication protocols suitable for this domain.

AUTHORS

Karthik Nandakumar (nkarthik@sg.ibm.com) is a research staff member with IBM Research, Singapore. Prior to joining IBM Research, he was a scientist with the Institute for Infocomm Research, A*STAR, Singapore, for more than six years. His research interests include pattern recognition, computer vision, and biometric recognition. He has coauthored two books and received a number of awards including the 2010 IEEE Signal Processing Society Young Author Best Paper Award.

Anil K. Jain (jain@cse.msu.edu) is a University Distinguished Professor in the Department of Computer Science and Engineering at Michigan State University. His research interests include pattern recognition, computer vision, and biometric recognition. He has received a Guggenheim fellowship, Humboldt Research Award, Fulbright fellowship, IEEE Computer Society Technical Achievement Award, IEEE W. Wallace McDowell Award, International Association for Pattern Recognition (IAPR)

King-Sun Fu Prize, and the IEEE International Conference on Data Mining Research Contribution Award for contributions to pattern recognition and biometrics. He is a Fellow of the IEEE, ACM, the American Association for the Advancement of Science, IAPR, and SPIE. He is the author of several books, and has served on The National Academies Panels on Information Technology, Whither Biometrics, and Improvised Explosive Devices.

REFERENCES

[1] A. K. Jain, A. Ross, and K. Nandakumar, *Introduction to Biometrics*. New York: Springer, 2011.

[2] P. Campisi, Ed., *Security and Privacy in Biometrics*. New York: Springer, 2013.

[3] R. Cappelli, A. Lumini, D. Maio, and D. Maltoni, "Fingerprint image reconstruction from standard templates," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 9, pp. 1489–1503, 2007.

[4] A. K. Jain, K. Nandakumar, and A. Nagar, "Biometric template security," *EURASIP J. Adv. Signal Process.*, vol. 2008, no. 1, p. 57941.

[5] C. Rathgeb and A. Uhl, "A survey on biometric cryptosystems and cancelable biometrics," *EURASIP J. Inform. Security*, vol. 2011, no. 1, pp. 1–25, 2011.

[6] S. Rane, Y. Wang, S. C. Draper, and P. Ishwar, "Secure biometrics: Concepts, authentication architectures, and challenges," *IEEE Signal Processing Mag.*, vol. 30, no. 5, pp. 51–64, Sept. 2013.

[7] A. B. J. Teoh, A. Goh, and D. C. L. Ngo, "Random multispace quantization as an analytic mechanism for biohashing of biometric and random identity inputs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 12, pp. 1892–1901, 2006.

[8] N. K. Ratha, S. Chikkerur, J. H. Connell, and R. M. Bolle, "Generating cancelable fingerprint templates," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 4, pp. 561–572, Apr. 2007.

[9] S. Tulyakov, F. Farooq, P. Mansukhani, and V. Govindaraju, "Symmetric hash functions for secure fingerprint biometric systems," *Pattern Recogn. Lett.*, vol. 28, no. 16, pp. 2427–2436, 2007.

[10] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith, "Fuzzy extractors: How to generate strong keys from biometrics and other noisy data," *SIAM J. Comput.*, vol. 38, no. 1, pp. 97–139, 2008.

[11] A. Juels and M. Sudan, "A fuzzy vault scheme," in *Proc. IEEE Int. Symp. Information Theory*, Lausanne, Switzerland, 2002, p. 408.

[12] A. Juels and M. Wattenberg, "A fuzzy commitment scheme," in *Proc. 6th ACM Conf. Computer and Communications Security*, Singapore, Nov. 1999, pp. 28–36.

[13] K. Nandakumar, A. K. Jain, and S. Pankanti, "Fingerprint-based fuzzy vault: Implementation and performance," *IEEE Trans. Inform. Forensics Security*, vol. 2, no. 4, pp. 744–757, Dec. 2007.

[14] Y. Sutcu, Q. Li, and N. Memon, "Protecting biometric templates with sketch: Theory and practice," *IEEE Trans. Inform. Forensics Security*, vol. 2, no. 3, pp. 503–512, Sept. 2007.

[15] F. Hao, R. Anderson, and J. Daugman, "Combining crypto with biometrics effectively," *IEEE Trans. Comput.*, vol. 55, no. 9, pp. 1081–1088, Sept. 2006.

[16] J. Bringer, H. Chabanne, G. Cohen, B. Kindarji, and G. Zemor, "Theoretical and practical boundaries of binary secure sketches," *IEEE Trans. Inform. Forensics Security*, vol. 3, no. 4, pp. 673–683, 2008.

[17] A. Nagar, K. Nandakumar, and A. K. Jain, "Biometric template transformation: a security analysis," in *Proc. SPIE, Electronic Imaging, Media Forensics and Security XII*, San Jose, vol. 7541, 2010, pp. 1–15.

[18] T. Ignatenko and F. M. J. Willems, "Biometric systems: Privacy and secrecy aspects," *IEEE Trans. Inform. Forensics Security*, vol. 4, no. 4, pp. 956–973, 2009.

[19] T. E. Boult, W. J. Scheirer, and R. Woodworth, "Fingerprint revocable biotokens: Accuracy and security analysis," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, Minneapolis, MN, June 2007, pp. 1–8.

[20] Y. C. Feng, P. C. Yuen, and A. K. Jain, "A hybrid approach for generating secure and discriminating face template," *IEEE Trans. Inform. Forensics Security*, vol. 5, no. 1, pp. 103–117, Mar. 2010.

[21] J. Bringer, H. Chabanne, and A. Patey, "Privacy-preserving biometric identification using secure multiparty computation: An overview and recent trends," *IEEE Signal Processing Mag.*, vol. 30, no. 2, pp. 42–52, Mar. 2013.

[22] D. Gafurov, B. Yang, P. Bours, and C. Busch, "Independent performance evaluation of pseudonymous identifier fingerprint verification algorithms," in *Proc. Int. Conf. Image Analysis and Recognition*, June 2013, pp. 63–71.

[23] S. Rane, "Standardization of biometric template protection," *IEEE MultiMedia*, vol. 21, no. 4, pp. 94–99, Oct. 2014.

[24] Y. Wang, S. Rane, S. C. Draper, and P. Ishwar, "A theoretical analysis of authentication, privacy, and reusability across secure biometric systems," *IEEE Trans. Inform. Forensics Security*, vol. 7, no. 6, pp. 1825–1840, Dec. 2012.

[25] A. W. K. Kong, "A statistical analysis of IrisCode and its security implications," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 3, pp. 513–528, Mar. 2015.

[26] C. Su and S. Srihari, "Evaluation of rarity of fingerprints in forensics," in *Proc. Advances in Neural Information Processing Systems*, 2010, pp. 1207–1215.

[27] M. Blanton and M. Aliasgari, "Analysis of reusability of secure sketches and fuzzy extractors," *IEEE Trans. Inform. Forensics Security*, vol. 8, no. 9, pp. 1433–1445, Sept. 2013.

[28] S. Yang and I. Verbauwhede, "Secure iris verification," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, vol. 2, 2007, pp. 133–136.

[29] E. J. C. Kelkboom, K. T. J. de Groot, C. Chen, J. Breebaart, and R. N. J. Veldhuis, "Pitfall of the detection rate optimized bit allocation within template protection and a remedy," in *Proc. IEEE 3rd Int. Conf. Biometrics: Theory, Applications, and Systems*, 2009, pp. 1–8.

[30] C. Chen, R. N. J. Veldhuis, T. A. M. Kevenaar, and A. H. M. Akkermans, "Biometric quantization through detection rate optimized bit allocation," *EURASIP J. Adv. Signal Process.*, vol. 2009, no. 1, p. 784834.

[31] M.-H. Lim and A. B. J. Teoh, "A novel encoding s cheme for effective biometric discretization: Linearly separable subcode," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 2, pp. 300–313, Feb. 2013.

[32] A. Nagar, S. Rane, and A. Vetro, "Privacy and security of features extracted from minutiae aggregates," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Dallas, TX, Mar. 2010, pp. 524–531.

[33] H. Xu, R. N. J. Veldhuis, A. M. Bazen, T. A. M. Kevenaar, T. A. H. M. Akkermans, and B. Gokberk, "Fingerprint verification using spectral minutiae representations," *IEEE Trans. Inform. Forensics Security*, vol. 4, no. 3, pp. 397–409, Sept. 2009.

[34] K. Nandakumar, "A fingerprint cryptosystem based on minutiae phase spectrum," in *Proc. 2nd IEEE Workshop on Information Forensics and Security*, Seattle, WA, Dec. 2010.

[35] R. Cappelli, M. Ferrara, and D. Maltoni, "Minutia cylinder-code: A new representation and matching technique for fingerprint recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 12, pp. 2128–2141, Dec. 2010.

[36] I. R. Buhan, J. M. Doumen, P. H. Hartel, and R. N. J. Veldhuis, "Fuzzy extractors for continuous distributions," in *Proc. ACM Symp. Information, Computer and Communication Security*, Singapore, Mar. 2007, pp. 353–355.

[37] L. Lai, S.-W. Ho, and H. V. Poor, "Privacy-security trade-offs in biometric security systems," *IEEE Trans. Inform. Forensics Security*, vol. 6, no. 1, pp. 122–151, Mar. 2011.

[38] L. Ballard, S. Kamara, and M. K. Reiter, "The practical subtleties of biometric key generation," in *Proc. 17th Conf. Security Symp.*, 2008, pp. 61–74.

[39] K. Simoens, B. Yang, X. Zhou, F. Beato, C. Busch, E. M. Newton, and B. Preneel, "Criteria towards metrics for benchmarking template protection algorithms," in *Proc. 5th IAPR Int. Conf. Biometrics*, Mar. 2012, pp. 498–505.

[40] T. Ignatenko and F. M. J. Willems, "Information leakage in fuzzy commitment schemes," *IEEE Trans. Inform. Forensics Security*, vol. 5, no. 2, pp. 337–348, June 2010.

[41] M. Inuma and A. Otsuka, "Relations among security metrics for template protection algorithms," in *Proc. IEEE 6th Int. Conf. Biometrics: Theory, Applications and Systems (BTAS)*, Sept. 2013, pp. 1–8.

[42] Y. Zhu, S. C. Dass, and A. K. Jain, "Statistical models for assessing the individuality of fingerprints," *IEEE Trans. Inform. Forensics Security*, vol. 2, no. 3, pp. 391–401, Sept. 2007.

[43] A. Adler, R. Youmaran, and S. Loyka, "Towards a measure of biometric feature information," *Pattern Anal. Applicat.*, vol. 12, no. 3, pp. 261–270, 2009.

[44] J. Daugman, "Probing the uniqueness and randomness of IrisCodes: Results from 200 billion iris pair comparisons," *Proc. IEEE*, vol. 94, no. 11, pp. 1927–1935, 2006.

[45] K. Takahashi and T. Murakami, "A measure of information gained through biometric systems," *Image Vis. Comput.*, vol. 32, no. 12, pp. 1194–1203, Dec. 2014.

[46] B. Fu, S. X. Yang, J. Li, and D. Hu, "Multibiometric cryptosystem: Model structure and performance analysis," *IEEE Trans. Inform. Forensics Security*, vol. 4, no. 4, pp. 867–882, Dec. 2009.

[47] X. Boyen, "Reusable cryptographic fuzzy extractors," in *Proc. ACM Conf. Computer and Communications Security*, Washington, DC, Oct. 2004, pp. 82–91.

[48] E. J. C. Kelkboom, J. Breebaart, T. A. M. Kevenaar, I. Buhan, and R. N. J. Veldhuis, "Preventing the decodability attack based cross-matching in a fuzzy commitment scheme," *IEEE Trans. Inform. Forensics Security*, vol. 6, no. 1, pp. 107–121, Mar. 2011.

[49] X. Boyen, Y. Dodis, J. Katz, R. Ostrovsky, and A. Smith, "Secure remote authentication using biometric data," in *Proc. Advances in Cryptology–EUROCRYPT* 2005, Aarhus, Denmark, May 2005, pp. 147–163.

[50] W. J. Scheirer, B. Bishop, and T. E. Boult, "Beyond PKI: The biocryptographic key infrastructure," in *Proc. IEEE Int. Workshop on Information Forensics and Security*, Dec. 2010.

[51] E. Barker, W. Barker, W. Burr, W. Polk, and M. Smid, "Recommendation for key management," NIST 800-57, July 2012.

[SP]

John Bustard

The Impact of EU Privacy Legislation on Biometric System Deployment

Protecting citizens but constraining applications

©ISTOCKPHOTO.COM/GREYFEBRUARY

Biometrics Security and Privacy Protection

Biometric systems provide a valuable service in helping to identify individuals from their stored personal details. Unfortunately, with the rapidly increasing use of such systems [1], there is a growing concern about the possible misuse of that information. To counteract the threat, the European Union (EU) has introduced comprehensive legislation [2] that seeks to regulate data collection and help strengthen an individual's right to privacy. This article looks at the implications of the legislation for biometric system deployment. After an initial consideration of current privacy concerns, the definition of "personal data" and its protection is examined in legislative terms. Also covered are the issues surrounding the storage of biometric data, including its accuracy, its security, and justification for what is collected. Finally, the privacy issues are illustrated through three biometric use cases: border security, online bank access control, and customer profiling in stores.

Digital Object Identifier 10.1109/MSP.2015.2426682
Date of publication: 13 August 2015

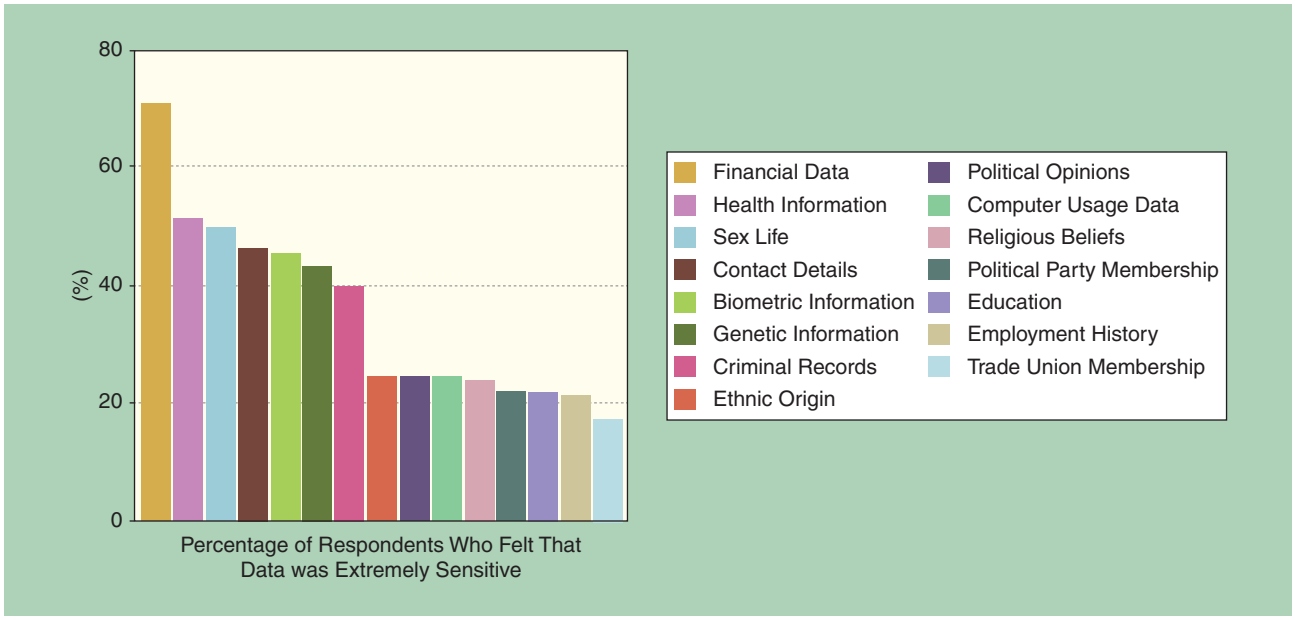

[FIG1] **Sensitivity ranking of personal data [4].**

PRIVACY CONCERNS WITH BIOMETRICS

Many people are currently concerned about the possible misuse of biometric data [3], [4] (see "Key Privacy Questions"). In 2006, for example, a telephone survey by the U.K. Information Commissioner's Office (Figure 1) revealed that more than 45% of respondents viewed biometric data as "extremely sensitive" [4]. This was a higher percentage than for other forms of personal data that already carry strong legal protections, such as ethnic origin, political opinions, religious beliefs, and trade union membership.

Some privacy issues and concerns specific to biometrics are:

- Biometric systems could be used to reveal medical conditions.
- Biometric use makes it easier to gather personal information, including the ability to do so covertly. For example, recent developments in biometrics at a distance [5] (Figure 2) have increased the accuracy with which individuals can be identified remotely. Such technology is starting to be deployed commercially in security [6] and customer profiling applications [7].

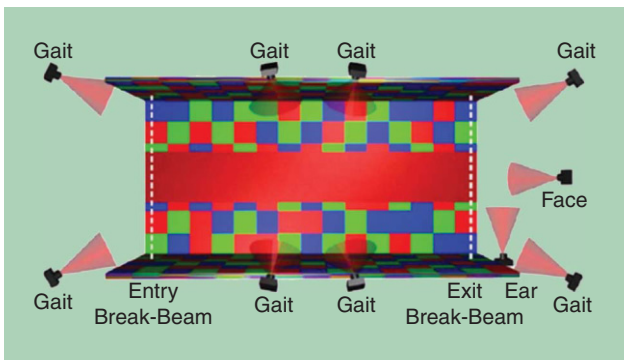

[FIG2] **An image from the Southampton Multi-biometric Tunnel [8]. The tunnel automatically recognizes individuals passing through it using three-dimensional gait, ear, and face recognition. Such systems can have many applications but also raise significant privacy concerns as they have the potential to be deployed covertly.**

KEY PRIVACY QUESTIONS

- What biometric data is being gathered and by whom?
- Is data being used solely for the purpose for which it was gathered?
- Is data accurate?
- Is data held securely?
- Is everyone operating within legal regulations?
- Are legal regulations sufficient?
- Are legal regulations proportionate to the threat posed to privacy?

- Biometrics could be used to link databases that have been anonymized yet still contain images of the individuals concerned. This is not necessarily an argument against the use of biometrics for identification, as much as a legitimate concern that de-anonymization techniques should not be applied to subvert citizens' attempts to maintain their privacy.

There are also psychological objections to biometric use, with some suggesting that measurements of a person's body are inherently more personal than other data about them [9]. Also in psychological terms, public resistance to the adoption of biometric technology is perhaps more a reflection of an understandable resistance to change rather than any substantial harm involved.

For example, this is illustrated in recent discussions about the use of biometrics in schools, where there was concern raised that such use could lead to "desensitization" [10].

Concerns cover both public and private use of biometrics. Despite legal regulations on how personal data, including biometrics, can be used, there remain doubts over whether organizations can be trusted to follow such regulations. Moreover, national security services are typically exempt from these controls, provided internal governmental oversight committees agree their actions are proportional to the threat involved. In light of recent revelations about data collection by some security services [11], however, there are understandable doubts that such oversight is sufficient. Indeed, even when organizations do not actively attempt to abuse personal data, it is often difficult in practice for them to ensure its privacy, as illustrated by some of the well-publicized breaches of security that have occurred [12].

The concerns of the public are further heightened by the fact that biometrics are often used in situations where there is a significant asymmetry of power between those deploying the technology and those who will be monitored by it. Examples range from employers monitoring the time keeping of employees [13] to governments monitoring those entering and leaving their country.

LEGAL CONTEXT

The growing concern over citizens' privacy has led to a number of changes in government legislation that will directly affect how biometric systems are deployed. In particular, the EU is in the process of introducing new data protection legislation [2] that will strengthen and unify existing laws in European member states. Significantly, the legislation also subjects companies outside the EU to the same data protection regulations if they offer services to EU citizens, or monitor their behavior.

THE EUROPEAN UNION HAS INTRODUCED COMPREHENSIVE LEGISLATION THAT SEEKS TO REGULATE DATA COLLECTION AND HELP STRENGTHEN AN INDIVIDUAL'S RIGHT TO PRIVACY.

DEFINITION OF PERSONAL DATA

When personal data is gathered by biometric systems it is subject to data protection legislation. The new European Data Protection Regulation defines personal data as:

> any information relating to an identified or identifiable natural person (data subject); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity.

In other words, if any data can be linked to an individual, it is "personal." This is intentionally broad and includes data such as the IP address of computers when such information can uniquely identify users [14]. Biometric data, both raw images and biometric templates, would clearly fall into this category, as they are inherently linked to a specific individual.

CCTV

Because the EU classifies facial images as sensitive personal data, this raises questions about the legitimacy of CCTV use, which frequently captures facial images without explicit consent. In a recent case, a Belgian court dealt with this issue by claiming that the data gathering itself is not processing [17]. However, this is inconsistent with the privacy concerns on which the legislation is based. Sensitive data is protected because it could be used for discrimination. While gathering CCTV imagery is not necessarily discriminatory, hackers or "feature creep" could lead to discriminatory applications in the future. If so, then the Belgian ruling may well be challenged at some stage.

The EU also makes a distinction between personal data and sensitive personal data, which is information that relates to health, sex life, racial or ethnic origin, political opinions, religious or philosophical beliefs, and even trade-union membership. Because of the close connection between biometrics and the physical body, ethnic origin and a number of medical conditions can be inferred from some biometric data, making it "sensitive" [15]. In particular, EU legislation explicitly mentions facial images as a form of sensitive personal data [16] (see "CCTV").

CONSENT

In general, the processing, storage, or transmission of sensitive personal data is not permitted. One important exception, however, is when explicit, free consent is given. This is convenient, for example, in applications such as unlocking a mobile phone. However, the use of such applications is still conditional on 1) sufficient data security being applied, 2) the data not being used for other purposes or shared with third parties, and 3) provision made for users to revoke their consent at any time.

Workplaces and commercial businesses are not typically required to obtain explicit, free consent for technology they deploy. This is because it is argued that employees and customers can leave organizations when they are uncomfortable with their working practices, although in reality, some may have little choice.

Current methods of explicit consent often take the form of complex legal terms and conditions that are typically not understood fully by the person giving consent. Also, such terms often do not reflect actual privacy preferences but are simply accepted because the person giving approval believes that there is no reasonable alternative [18].

For many biometric applications, there will be an explicit enrolment stage where biometric features are recorded in a controlled way. This stage may be the appropriate point at which to obtain explicit consent. Biometrics technology can also be used to identify

whether someone has agreed to biometric identification, as long as all biometric information is discarded if consent is not given [19].

PROTECTION THROUGH ANONYMITY

One general approach to overcoming the limitations imposed by data protection legislation is to anonymize data. However, this is not an option for biometric systems. As noted in a report by a data protection committee for the council of Europe:

> ... with regard to biometric data, the option of making the data anonymous is not available as biometric data by their very nature, form an instrument to identify individuals, particularly when they are automatically processed [20].

For EU law, the definition of *identifiable* is so broad that data can be considered personal if the data controller has any way of identifying the persons behind the data [21].

There are also obligations to implement data protection by "design and by default" (Article 23) [2]. These design principles mean that biometric system designers are obliged to minimize the quantity of personal data that is collected and processed. They must also restrict the time that data is held and keep the number of individuals who have access to the data to a minimum. Existing analysis of default practices [22] indicate that most people will accept default settings. As a result, requiring explicit consent is likely to result in the substantially reduced adoption of new biometric technologies.

> THE REVERSAL MOVES THE FOCUS OF BIOMETRIC TECHNOLOGY USE FROM SITUATIONS IN WHICH IT MAY BE BENEFICIAL, TO SITUATIONS IN WHICH IT IS EVIDENTLY NEEDED.

One possible technical approach to anonymity is to use encryption methods to separate the storage of biometric templates from the system performing the verification [23]. This is done to ensure that the organization that stores the templates is unaware of which verification transactions are occurring and, in turn, that the organization verifying an identity cannot access the personal biometric template. Such an approach would not necessarily avoid the necessity for consent as the initial storage of biometrics would require user permissions, as would any processing performed using such data. However, users may find that such an approach is more acceptable to them than trusting a single organization with all of their data. Such anonymization techniques are still at a research stage, however, and so are unlikely to form a legal requirement. However, once practical commercial implementations become available, data protection authorities may interpret them as "data protection by design" requirements.

PROTECTION THROUGH AGGREGATE STATISTICS

Biometric technology can also be used to create aggregated statistics as, for example, in recognizing the number of unique visitors to a store. In this way, biometric systems can help automate business intelligence gathering that has historically been performed manually. This aggregated usage data is typically anonymous, referring only to total numbers of unique individuals rather than individual usage patterns. This does not address the privacy issues of using biometric information itself but does limit how much information is linked to a specific individual.

It is currently permissible in the EU to obtain categorization information about a person, as long as that information is not itself personally identifiable and the information is not combined in a way that can make it personally identifiable. However, the new regulation includes restrictions on the use of categorization data for profiling purposes:

> Every natural person shall have the right not to be subject to a measure which produces legal effects concerning this natural person or significantly affects this natural person, and which is based solely on automated processing intended to evaluate certain personal aspects relating to this natural person or to analyse or predict in particular the natural person's performance at work, economic situation, location, health, personal preferences, reliability or behaviour (Article 20) [2].

This would restrict the use of so-called soft biometrics [24], which identify broad features of an individual, such as their age, sex, or race. The systems involved do not gather uniquely identifiable biometric signatures and so could provide demographic information about customers without identifying them. However, there are situations where soft biometric data may be sufficient to identify an individual uniquely and so may also be problematic in relation to the legislation.

BIOMETRIC DATA RETENTION ISSUES

There are a number of data protection issues associated with the storage of personal data. In particular, the ensured accuracy, security, control, and proportionality of that storage are especially important.

DATA ACCURACY

Stored personal data must be accurate. The European Data Protection Regulation states in Article 5 that personal data must be:

> (d) accurate and kept up to date; every reasonable step must be taken to ensure that personal data that are inaccurate, having regard to the purposes for which they are processed, are erased or rectified without delay (Article 5) [2].

Biometric signatures can change and thus any biometric system needs a means of updating biometric templates. In particular, ageing has a significant effect on many biometrics [25].

DATA SECURITY

The General European Data Protection Regulation states that:

> The controller and the processor shall implement appropriate technical and organizational measures to ensure a level of security appropriate to the risks represented by the processing and the nature of the personal data to be protected, having regard to the state of the art and the costs of their implementation (Article 23) [2].

The regulation also states that further acts may be passed for the purpose of specifying the criteria for achieving these standards.

Organizations such as EuroPriSe [26] provide certification of products and IT services to ensure that a sufficiently high standard of data security is in place. Such standards have a strong emphasis on internal organizational measures, including an assurance of the physical security of stored data, and providing authentication and logging facilities to ensure that only authorized processing is performed. However, the numerous data breaches that have occurred suggest that either these measures are insufficient within large organizations or that they are difficult to enforce in practice.

In addition to these general data security measures a number of technologies designed specifically for securing biometric templates have been developed and this continues to be an active area of research [27]. The precise methods required for securing biometrics have not been made explicit within the law. In practice they will be determined by judgements based on advice from experts. The subsequent rulings will then provide a precedent for what security measures are required.

IT IS POSSIBLE TO PROCESS PERSONAL DATA WHERE IT IS IN THE SUBSTANTIAL PUBLIC INTEREST AND WHERE REQUIRING CONSENT WOULD UNDERMINE THE EFFECTIVENESS OF ITS USE.

DATA CONTROL

The new EU regulation emphasizes the rights of individuals to control the information that is stored about them. Those gathering personal data, including data that could be used for biometric analysis, must clearly inform those affected that the data is being collected and explain how it will be used (Article 5) [2]. They must also provide a means to identify the information already stored and enable those affected to adjust that information if it is inaccurate. In addition, citizens have the right to object to such data processing, requiring it to cease unless organizations can demonstrate "compelling legitimate grounds" (Article 19) [2]. This is a significant change of emphasis from previous legislation where processing was permissible unless citizens could find a legitimate reason for it to stop. The reversal moves the focus of biometric technology use from situations in which it may be beneficial, to situations in which it is evidently needed.

DATA COLLECTION PROPORTIONALITY

The new Data Protection Regulation allows for the use of biometrics without consent provided certain conditions are met. In particular, it is possible to process personal data where it is in the substantial public interest and where requiring consent would undermine the effectiveness of its use. For example, this includes the prevention or detection of crime (Article 2) [2] and journalistic investigation (Article 80) [2]. This means that for many security applications, biometric use would still be possible. However, in such cases the data processing must be:

> (c) adequate, relevant, and limited to the minimum necessary in relation to the purposes for which they are processed (Article 5) [2].

In addition, for it to be used legitimately, biometrics must be judged proportional to the application [28]. For example, the widespread use of biometric systems in schools within the United Kingdom was challenged by the EU commission on grounds of proportionality, which resulted in a requirement for parental consent and alternative identification methods being made available [29].

One consideration in assessing proportionality is whether a less invasive alternative approach could be used. As biometrics is considered a potential threat to privacy, this ruling, in effect, means that biometrics is only proportional when no reasonable alternative identification method exists, which imposes a significant bias against the use of biometric technology. This restriction seems out of step with the treatment of other workplace practices, such as the use of time-keeping machines, which may have similar negative associations but are not limited by legislation in the same way. In many cases, biometrics is used as a convenient alternative to a door key or identity card, and so perhaps should be treated in a similar way.

USE CASES FOR BIOMETRICS

Biometrics can be applied in a variety of different circumstances and each brings with it different concerns and legislative constraints. This section aims to highlight these differences with three example use cases: border security, online bank account access control, and customer profiling. The discussion covers consent for data collection, the security measures for its protection and, where appropriate, a consideration of the proportionality of data collection and the accuracy of that data.

BORDER SECURITY

All countries carefully monitor the identity of individuals passing through their borders. These checks identify suspected security threats and help prevent illegal immigration. Biometrics technology offers a means to automate this process as well as potentially increasing the accuracy of identifying those claiming a false identity.

From an ethical perspective, it is important that any such automated system be suitable for the diverse range of users it is likely to process, from babies to wheelchair users. No section of the population should experience undue inconvenience or unjustified discrimination because of their specific needs.

In addition to automating routine identification checks, biometric systems can also be used to monitor a "watch list" of individuals who are a particularly high security risk. Such individuals are likely to be traveling with false identification papers that may pass existing inspection methods.

- *Consent*: Where biometrics are applied as a cost-saving automation measure, biometric use may be optional. This often takes the form of a "fast-track" route through border

[FIG3] Biometric passports are being widely adopted throughout the EU.

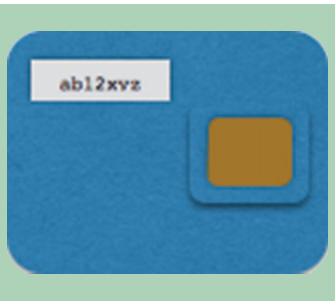

[FIG4] Small portable devices with fingerprint readers can be used to provide time-linked passwords to secure online services.

control. However, where biometrics are applied as a means to improve the accuracy of identification, such as with the U.S.-VISIT program [30], biometric checks will be mandatory.

- *Accuracy*: When biometric systems are used to automate passport control, an alternative verification method is needed. This is due to the potential for a false negative verification match as well as the likelihood of processing citizens who have missing or damaged biometric features. For example, fingerprints can be obscured by manual work.
- *Security*: Public bodies have significant oversight and, in some cases, have freedom of information legislation that would facilitate investigation of abuses of data protection. This is generally not possible in the private sector, where such investigations would reveal commercially sensitive information. However, in the case of border control, there are significant security issues. As a result, the operation of the technology is likely to be secret and thus it falls to whistle-blowers to reveal potential abuses by government.

Another security factor crucial to the privacy of users is whether biometric information is held on a centralized database or carried with the user, such as on a biometric passport (Figure 3). Each additional link in processing or data storage carries with it an increased risk that it may be compromised by hackers or that feature creep by one of the organizations involved will lead to further invasions of privacy. Similar concerns can arise if biometric data is transmitted to a third party to perform verification tests. To some extent, modern encryption methods can mitigate these concerns, but they do not remove them as all solutions rely on some degree of trust.

However, even without a centralized biometric database, the introduction of identity papers with biometric information can potentially introduce significant privacy issues. If the biometric data is accessible via a remote wireless connection, there is a risk that passports could be compromised by a hacker with a nearby sensor. Likewise, such passports would require an enrollment system, which itself may involve a number of third parties, each of which could be compromised or could introduce privacy invading features in the future.

- *Proportionality:* Border security focuses on preventing serious criminal and terrorist activity and, as a result, it is considered legitimate to partially invade individual privacy if doing so preserves the higher priority of preventing harm. However, if biometrics are used in watch-list applications there are further concerns. In particular, such applications raise the question of proportionality. Specifically, on what grounds should border security be permitted to automatically identify an individual and subject them to increased scrutiny? Also, because of the potential seriousness of watch-list false matches, the accuracy of biometric identification needs to be considered—particularly in light of the case of Brandon Mayfield, who was held for over two weeks on terrorism charges, partly because of a single false match to a fingerprint obtained from bomb parts [31]. In addition, preventing an individual from leaving a country is a serious restriction on their freedom and a common abuse of governmental powers against critics [32].

ONLINE BANK ACCOUNT ACCESS CONTROL

Bank transactions are increasingly being performed using online applications that enable the monitoring of accounts and the transfer of funds. However, there is a significant risk that criminals may use these systems to steal from the accounts involved. Biometrics is one way to improve the security of online banking. Specifically, in conjunction with existing security systems, biometrics can be used to provide two-factor identification. This can take the form of a combination of something that is known, say a password, with someone's biometric signature, based on a physical feature. Other factors can also be used to further enhance security such as the media access control address of the user's PC.

THE NEW EU REGULATION EMPHASIZES THE RIGHTS OF INDIVIDUALS TO CONTROL THE INFORMATION THAT IS STORED ABOUT THEM.

- *Consent*: It seems reasonable to assume implied consent where customers have the option of moving to another bank that doesn't require biometric security. Under the new Data Protection Regulation, however, explicit, free consent require-

ments mean that implied consent is insufficient.

- *Security*: As with passport control, the privacy of the system is affected by whether biometric information can be kept locally. This is possible, for example, by using a fingerprint scanner (Figure 4) to unlock a device owned by the customer that produces a secure, time-linked password for accessing a remote banking Web site. Some biometrics, such as voice, however, may require biometric templates to be stored on a server.
- *Proportionality*: Under EU law, verification applications, where a user claims an identity that is verified, are considered less invasive than recognition applications where a user is compared against a large database to determine identity. However, unless free consent is provided, such technology may well be viewed as disproportionate if alternative security methods are available. An individual bank may view biometrics as a more secure alternative, but the final decision would rest with the courts.

CUSTOMER PROFILING

Although the use of biometrics has traditionally been associated with security applications, there are many other circumstances where the automated recognition of individuals is valuable. One commercially important area is in tracking customers while they shop, primarily to help understand their interests and hence identify ways to sell them more goods and services (Figure 5). The current technology used for customer tracking in physical spaces is similar to the initial tracking capabilities of Web analytics companies, focusing on counting the number of unique visitors and identifying statistics of where customers travel within stores.

- *Consent*: One form of consent is through the use of a loyalty card, which already tracks customer behavior through monitoring their purchases. However, not everyone uses a loyalty card and some are concerned about being monitored in this way. Another form of monitoring is through the tracking of the unique identifier transmitted by a customer's smart phone. Here companies typically try to obtain consent by using an opt-out policy, posting signs to inform customers that they are being monitored [33]. This approach is controversial, however [34], and would no longer be permissible under the new EU regulations if any of the gathered data were categorized as personal.
- *Security*: To facilitate tracking across different locations it may be necessary to share biometrics between different sites. If a centralized database is used, this will increase concerns about the security of the data. There is also likely to be a market for user profile information, similar to how such data is used for online profiling of customers. Current legislation would require strong contractual constraints on such data, particularly if any of it is identified as being sensitive.

[FIG5] **Using soft biometrics, advertising billboards can detect the numbers, genders, and age groups of viewers in an area. This helps retailers understand how shoppers are affected by advertising and promotions.**

THE NEW EU REGULATION WILL SIGNIFICANTLY INCREASE THE PROTECTION OF EACH CITIZEN'S PRIVACY. HOWEVER, IT IS ALSO LIKELY TO LIMIT THE ADOPTION OF BIOMETRIC TECHNOLOGY, PARTICULARLY IN WORKPLACES AND IN COMMERCIAL ORGANIZATIONS.

It is likely that there would also be a commercial advantage in extending monitoring to provide similar levels of information to that available via online profiling. Such profiling includes the acquisition of demographic information, such as age and gender [35]. Recent developments in soft biometrics [24] could be used to estimate some of this additional information but the new regulation is likely to greatly restrict these applications unless explicit free consent has been given.

CONCLUSIONS

This article began by acknowledging public concern about the possible abuse of the personal data that biometric systems collect and store. This set the context for identifying the main measures introduced by the new European Data Protection Regulation to control data collection and help strengthen a citizen's rights to privacy and data protection. The discussion first clarified what is meant by "personal data" in the legislation, before considering its use in protecting that data. Protection included a consideration of the legislation relating to the accuracy and security of the data held, justification for what is collected, and the option of individuals providing "consent" to data use. The privacy issues associated with biometrics were illustrated through a consideration of three

biometric use cases: border security, online banking, and customer tracking in stores.

Biometrics is frequently given as an example of a technology that raises privacy concerns and, as has been shown in this article, there are significant legislative restrictions applied to its use. In general, it is desirable to limit the complexity and application of legal restrictions as they will consume valuable time and resources. Also, increased complexity of the law further isolates citizens from the legal process and can create a situation where only those who can afford specialized legal services can understand when they are acting legitimately. EU data protection regulation can be interpreted as focusing on minimizing personal data collection. Further work is needed to identify if the harms that such legislation prevent are adequately balanced against the potential gains made possible by the new technology.

The overall conclusion is that the new EU regulation will significantly increase the protection of each citizen's privacy. However, it is also likely to limit the adoption of biometric technology, particularly in workplaces and in commercial organizations.

AUTHOR

John Bustard (j.bustard@qub.ac.uk) is a lecturer in the Speech, Image, and Vision Systems Research Group at Queen's University Belfast, United Kingdom. He has more than seven years research experience in biometrics, including work on the Tabula Rasa European Framework 7 project investigating the vulnerability of biometrics to spoofing attacks. He is currently a U.K. expert for the International Organization for Standardization group contributing to the WD5 30107 standard on presentation attack detection. He also has an interest in ethics, which has led to invited presentations on the ethics of biometrics and surveillance for the Centre for Science, Society, and Citizenship and for the National Endowment for Science, Technology, and the Arts. He is currently the chair of an ethical advisory group for the European Framework 7 research projects ISAR+ and SOTERIA. He is a co-investigator on the Centre for Secure Information Technologies Phase 2 EPSRC project, which is concerned with securing the Internet of Things.

REFERENCES

[1] Article 29 Data Protection Working Party, "Opinion 03/2012 on developments in biometric technologies," WP 193, Apr. 2012.

[2] European Commission. (2012). Proposal for a regulation of the European parliament and of the council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (general data protection regulation). [Online]. Available: http://eur-lex.europa.eu/legal-content/en/ALL/?uri=CELEX:52012PC0011

[3] A. Krupp, C. Rathgeb, and C. Busch, "Social acceptance of biometric technologies in Germany: A survey," in *Proc. Int. Conf. Biometrics Special Interest Group (BIOSIG)*, Sept. 2013, pp. 1–5.

[4] K. McCullagh, "Data sensitivity: Proposals for resolving the conundrum," *J. Int. Commercial Law Technol.*, vol. 2, no. 4, pp. 190–201, 2007.

[5] M. Tistarelli, S. Z. Li, and R. Chellappa, *Handbook of Remote Biometrics: For Surveillance and Security*, 1st ed. New York: Springer, 2009.

[6] Progeny Systems Corporation. (2014). Surveillance, persistent observation and target recognition (spotr). [Online]. Available: http://www.spotrtech.com/

[7] S. Harris. (2013). Computer to shopstaff: VIP approaching. [Online]. Available: http://www.thesundaytimes.co.uk/sto/news/uknews/Tech/article1287590.ece

[8] R. D. Seely, S. Samangooei, M. Lee, J. N. Carter, and M. S. Nixon, "The University of Southampton multi-biometric tunnel and introducing a novel 3D gait dataset," in *Proc. 2nd IEEE Int. Conf. Biometrics: Theory, Applications and Systems*, Sept. 2008, pp. 1–6.

[9] J. D. Woodward, K. W. Webb, E. M. Newton, M. A. Bradley, D. Rubenson, K. Larson, J. Lilly, K. Smythe, et al., "Army biometric applications: Identifying and addressing sociocultural concerns," MR-1237-A, Rand Corp.

[10] A. Ramasastry. (2012). Biometrics in the school lunch line: Why parents should be concerned about the privacy implications of this trend. [Online]. Available: https://verdict.justia.com/2012/10/09/biometrics-in-the-school-lunch-line

[11] S. Landau, "Making sense from Snowden: What's significant in the NSA surveillance revelations," *IEEE Security Privacy*, vol. 11, no. 4, pp. 54–63, 2013.

[12] PriceWaterhouseCoopers. (2014). The global state of information security survey. [Online] Available: http://www.pwc.com/gx/en/consulting-services/information-security-survey/download.jhtml

[13] Allday Time Systems Ltd. (2015). Biometric attendance systems. [Online]. Available: http://www.alldaytime.co.uk/

[14] Scarlet Extended SA v Société belge des auteurs, compositeurs et éditeurs SCRL (SABAM) Case C 70/10, 24 Nov. 2011.

[15] E. Mordini, "Biometrics, human body, and medicine: A controversial history," *Ethical, Legal, and Social Issues in Medical Informatics, in Medical Information Science Reference*, 2008, pp. 249–272.

[16] Project Group on Data Protection, "Progress report on the application of the principles of Convention 108 to the collection and processing of biometric data," 2005.

[17] P. De Hert, O. De Schutter, and S. Gutwirth, "Pour une réglementation de la vidéosurveillance," *Journal des Tribunaux*, vol. 115, pp. 569–579, Sept. 1996.

[18] E. Luger, S. Moran, and T. Rodden, "Consent for all: Revealing the hidden complexity of terms and conditions," in *Proc. SIGCHI Conf. Human Factors in Computing Systems*, 2013, pp. 2687–2696.

[19] Article 29 Data Protection Working Party, "Opinion 02/2012 on facial recognition in online and mobile services," WP 192, 2012.

[20] Y. Liu, "Identifying legal concerns in the biometric context," *J. Int. Commercial Law Technol.*, vol. 3, no. 1, pp. 45–54, 2008.

[21] P. de Hert, "Biometrics: Legal issues and implications," in *Background Paper for the Institute of Prospective Technological Studies, DG JRC–Sevilla*, European Commission, 2005.

[22] C. W. Park, S. Y. Jun, and D. J. MacInnis, "Choosing what I want versus rejecting what I do not want: An application of decision framing to product option choice decisions," *J. Market. Res.*, vol. 37, no. 2, pp. 187–202, 2000.

[23] N. D. Sarier, "A survey of distributed biometric authentication systems," in *Proc. Special Interest Group on Biometrics and Electronic Signatures*, pp. 129–140 Sept. 2009.

[24] D. Reid, S. Samangooei, C. Chen, M. Nixon, and A. Ross, "Soft biometrics for surveillance: An overview," *Machine Learning: Theory and Applications*, Handbook of Statistics, Vol. 31. Elsevier, 2013, pp. 327–352.

[25] M. Fairhurst, "Age factors in biometric processing," *IET Digital*, Professional Applications of Computing. Institution of Engineering and Technology.

[26] EuroPriSe. [Online]. Available: www.european-privacy-seal.eu

[27] A. K. Jain, K. Nanakumar, and A. Nagar, "Biometric template security," *EURASIP J. Adv. Signal Process.*, vol. 2008, pp. 1–17, 2008.

[28] Office of the Privacy Commissioner of Canada. (2011, Feb.). Data at your fingertips: Biometrics and the challenges to privacy. [Online]. Available: https://www.priv.gc.ca/information/pub/gd_bio_201102_e.asp

[29] UK Department for Education. (2012). Protection of biometric information of children in schools. [Online]. Available: https://www.gov.uk/government/publications/protection-of-biometric-information-of-children-in-schools

[30] L. Amoore, "Biometric borders: Governing mobilities in the war on terror," *Political Geogr.*, vol. 25, no. 3, pp. 336–351, 2006.

[31] Mayfield v. US, No. 07-35865. FindLaw. Dec. 2009.

[32] C. Harvey and R. P. Barnidge. "Human rights, free movement, and the right to leave in international law," *Int. J. Refugee Law*, vol. 19, no. 1, pp. 1–21, 2007.

[33] Future of Privacy Forum. (2013). Mobile location analytics code of conduct. [Online]. Available: http://www.futureofprivacy.org/wp-content/uploads/10.22.13-FINAL-MLA-Code.pdf

[34] P. Cohan. (2013). How Nordstrom uses WiFi to spy on shoppers. [Online]. Available: http://www.forbes.com/sites/petercohan/2013/05/09/how-nordstrom-and-home-depot-use-wifi-to-spy-on-shoppers/

[35] Google. (2014). Overview of demographics and interests reports. [Online]. Available: https://support.google.com/analytics/answer/2799357

[SP]

Carlos Óscar S. Sorzano

sp EDUCATION

Undergraduate Students Compete in the IEEE Signal Processing Cup: Part 2

This is the second part to [1], which summarized the first two editions of the IEEE Signal Processing Cup. The perspectives of the competition organizers and the student finalists are given in this column.

SHAPING THE FIRST TOPIC

In May 2013, Dimitri van de Ville, the chair of the IEEE Bioimaging and Signal Processing Technical Committee (BISP TC) [2] forwarded an e-mail to the members of the committee calling for proposals for the IEEE Signal Processing Society's (SPS's) First Signal Processing Cup. The idea was to promote signal processing, particularly its applications to solve real-world problems, among undergraduate students.

I work at the Instruct Image Processing Center [3] in Madrid, Spain, a reference center for image processing in structural biology. The call for proposals was a perfect match between the need to achieve higher resolution in the three-dimensional (3-D) reconstructions of biological macromolecules and the possibility to explore the recently introduced ideas in the areas of superresolution, image restoration, and denoising. The possibility of having students all over the world trying to tackle this problem sounded attractive and challenging. So, with the support of the BISP TC, the Instruct Image Processing Center submitted a proposal for enhancing the resolution achieved by electron microscopes in the elucidation of the structure of single particles. Students could be as creative as they wished as long as the resolution limit imposed by the microscope could be pushed toward atomic resolution.

In August 2013, the proposal was approved by the SPS and, in November, we already had training and testing data sets available on the Internet. We also set up an online system [4] that allowed participants to self-evaluate their performances: they only had to upload their results to a web page and the system returned immediate feedback about the quality of their 3-D reconstructions. This system has been most valuable to empower the challengers and let them explore multiple directions while pursuing the most promising algorithms. The website activity, as well as the interaction of students with the organizers of the challenge, was quite high in the following months.

ALGORITHM FINALISTS

The objective of the challenge was to enhance the resolution of macromolecular structures reconstructed from images taken by an electron microscope (EM) (see Figure 1).

The best algorithm (Algorithm 1) consisted of three steps: 1) estimating the point spread function from the molecule background, 2) deconvolving with the Richardson–Lucy iterative algorithm, and 3) automatically detecting the location of the macromolecule and masking the output to produce a background-free volume.

The second-best algorithm (Algorithm 2) was an exemplar-based approach to enhance the degraded EM images. The key observation is that the protein data bank (PDB) data share a high degree of similarity within protein structures. The algorithm exploits PDB data as prior information to help recover the degraded input EM data.

The third-place algorithm (Algorithm 3) took advantage of the availability of multiple realizations of the same

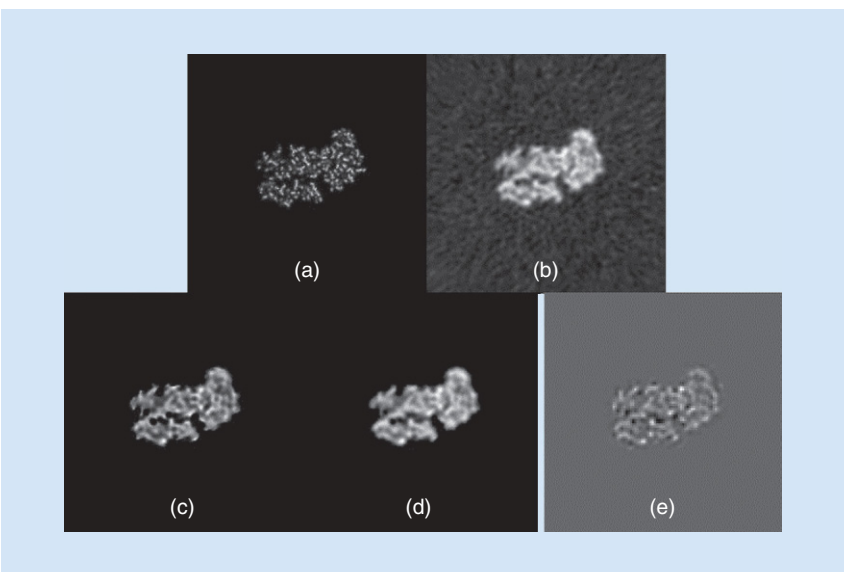

[FIG1] (a) The central slice of one of the ground-truth volumes. (b) The corresponding slice in the challenging volume. (c)–(e) Slices produced by Algorithms 1, 2, and 3, respectively.

Digital Object Identifier 10.1109/MSP.2015.2435816
Date of publication: 13 August 2015

molecule. Their superresolution approach is based on the idea that if one has multiple low-resolution images that are spatially shifted a bit, then a high-resolution image can be estimated by taking all of them into consideration.

THE WINNING TEAMS

In May 2014, during the final presentation of the shortlisted algorithms at the International Conference on Acoustics, Speech, and Signal Processing Conference in Florence, Italy, the jury members were impressed by the technical level of the presentations and the algorithms developed (see the algorithms summary in the "Algorithm Finalists" section). The presentations certainly covered a wide range of different image and signal processing techniques, showing the success of the challenge to promote creativity, learning in the area of signal processing, and engagement of undergraduate students to solve real-world problems. And not only that, students actually dived into the details of the application domain for which they were trying to provide innovative solutions. As the person who set up the problem, I was astounded when, five minutes before the final presentations in Florence, I heard a couple of students in different groups discussing how images were recorded and processed in electron microscopy (EM)—and I realized that most of the details they mentioned were technically correct. This was quite impressive for people who haven't been in a structural biology lab! In the final competition, teams EPOCH (Bangladesh University of Engineering and Technology), NtUeLsA (National Taiwan University), and Uchihas (Bangladesh University of Engineering and Technology) won first, second, and third place, respectively.

Enough of my babbling—I now give the floor to the three finalist teams of the Signal Processing Cup 2014 (see "Team EPOCH," "Team NtUeLsA," and "Team Uchihas").

REFERENCES

[1] K.-M. Lam, C. Ó. S. Sorzano, Z. Zhang, and P. Campisi, "Undergraduate students compete in the IEEE Signal Processing Cup: Part 1," *IEEE Signal Processing Mag.*, vol 32, no. 4, pp. 123–125, July 2015.

[2] SPS Technical Committee on Bioimaging and Signal Processing (BISP TC). [Online]. Available: http://www.signalprocessingsociety.org/technical-committees/list/bisp-tc/

[3] Instruct Image Processing Center (I2PC). [Online]. Available: http://i2pc.cnb.csic.es/

[4] 3DEM Benchmark. [Online]. Available: http://i2pc.cnb.csic.es/3dembenchmark/LoadSubTasks.htm?subtaskId=3

TEAM EPOCH—*FROM BANGLADESH UNIVERSITY OF ENGINEERING AND TECHNOLOGY*

We learned of the IEEE Signal Processing Cup 2014 from the IEEE Signal Processing Society's website. As we were highly interested in working on real-world problems in signal processing, we immediately decided to take part in the competition. However, the theme of the competition, "Image Restoration/Superresolution for Single Particle Analysis" was completely new to us. So we, four undergraduate students from the Department of Electrical and Electronic Engineering, contacted one of our faculty members, Dr. Mohammad Ariful Haque. We asked him to be our supervisor in the SP Cup. After reviewing the problem, he agreed to supervise our team.

We met three times a week to check our progress with Dr. Haque. At first, we did not know much about image processing, let alone single-particle analysis. Our supervisor suggested we read numerous review papers and explore the recent state of the art in image enhancement and restoration techniques. We started our journey with the book *Digital Image Processing* by Gonzalez and Woods. At the time of our review, we had realized that the main problem was to estimate the blur function in 3-D electronic microscopic volume. So, we concentrated our efforts on estimating the blur function from the noisy and low-resolution 3-D volumes. Without any prior information about the latent structure, estimating the blur function was really difficult. Our supervisor therefore suggested we exploit the noise cue that was generated from the amorphous ice (in which the macromolecules are embedded). After a long series of experiments, we could successfully estimate the blur function. We then applied the standard Lucy–Richardson (LR) deconvolution algorithm to restore the high-resolution volume. With the submission deadline quickly approaching, we started to work day and night to test, revise, and refine our algorithm. At this stage, we were able to improve the performance of our algorithm by applying constrained LR deconvolution as well as reducing the spatial noise outside of the main molecules using the breadth-first search technique.

Our joy knew no bounds when we won first place in the inaugural competition of the SP Cup. It was really interesting to work on a real-world signal processing problem. We are grateful to Mathworks: MATLAB made our work much easier and better. Without MATLAB, we would have had to divert much more effort to writing code rather than developing the algorithm. The technical group of the competition was really helpful. Last but not least, we feel very fortunate to work under a great supervisor. Thanks to our participation to the SP Cup, we could get a taste of real-world signal processing work, and it has been a lifetime experience for us.

—Anik Khan, Forsad Al Hossain, Tawab Ullash, and Abu Raihan, Undergrads

—Mohammad Ariful Haque, Supervisor

TEAM NtUeLsA—*FROM NATIONAL TAIWAN UNIVERSITY*

We heard about the IEEE Signal Processing Cup 2014 from Prof. Shao-Yi Chien. Thanks to his encouragement, a team of nine members—eight from the Department of Electrical Engineering and one from the Department of Computer Science and Information Engineering—was put together in a short time. Prof. Chien also asked his Ph.D. student, Wei-Chih Tu, whose field of research is image processing and computational photography, to help coach the team.

We were quite excited about the fact that the challenge was coming from a real-world problem, although we did not have any background in EM imaging. The first step was literature survey. Under the guidance of our supervisors, we reviewed the advances of image superresolution for natural images and adopted some of these ideas in our work. We categorized superresolution algorithms into signal processing methods and learning-based methods, and we finally decided to focus on the learning-based methods.

Indeed, we found out that the data in the Signal Processing Cup 2014 was not only about down-sampling and that more prior information could be considered in the learning-based methods to further improve performances. The second step involved algorithm development, implementation, and evaluation, with the latter part being facilitated by the online self-evaluation system provided by the Signal Processing Cup. In our final proposal, the 3-D input data is treated as a stack of low-resolution images, a mask is generated to highlight the region of interest for shortening processing time, and the concept of learning-based superresolution is adopted. The input image is split into patches, and their nearest neighbors are searched from the training set. With low-resolution data and their high-resolution counterparts in the training set, we directly take the high-resolution parts of the nearest neighbor as the high-resolution part of the target low-resolution input, and the overlapped pixels are averaged for smoothness. Eventually, we had to prepare the presentation. Although only one member was on the stage for the final competition, all team members spent a lot of time rehearsing for the presentation.

The Signal Processing Cup has been an unforgettable journey for us all. We were motivated by the real-world nature of the problem, which is quite a different experience from doing homework in textbooks. The online self-evaluation system providing immediate feedback was quite helpful for algorithm development. Moreover, we would like to thank Mathworks. With MATLAB, our ideas could be prototyped quickly for evaluation, and the quick feedback always stimulated our creativity. We learned a lot about how signal processing technology can solve real-world problems, and our problem-solving skills became more complete, including literature survey, algorithm development, and presentation. This wonderful experience motivated all of the team members to go for advanced studies, and we will definitely encourage other students to participate in the next Signal Processing Cup.

—Kai-Wen Liang, Yen-Chen Wu, Guan-Lin Chao, Kuan-Hao Huang, Shao-Hua Sun, Ming-Jen Yang, Po-Wen Hsiao, Ti-Fen Pan, and Yi-Ching Chiu, Undergrads

—Wei-Chih Tu, Graduate student
—Shao-Yi Chien, Supervisor

TEAM UCHIHAS—*FROM BANGLADESH UNIVERSITY OF ENGINEERING AND TECHNOLOGY*

We were all in our third year of undergraduate studies when we first heard about this competition. At that time, we only had a handful of baseline course experience on continuous signals and linear systems, digital signal processing, random signal processing, etc. However, we had no preliminary background on EM, which made the problem description quite technical and daunting for us but, even so, we still wanted to test ourselves. We eventually signed up as a team along with our supervisor, Dr. Md. Kamrul Hasan.

He encouraged and guided us through the whole adventure. At the very beginning, he used to ask for slide presentations on topics like single-particle analysis and superresolution…and we had just learned about these topics from simple Google and Wikipedia searches. In a couple of weeks' time, we had a basic understanding of SP and a more refined understanding of the problem statement. We then tried to think of plausible solutions that might work. Since the problem was essentially a 3-D image restoration problem, we looked at state-of-the-art algorithms for the two-dimensional case. We reviewed a lot of literature, and this was the point where our research truly began. We were literally overwhelmed with notions such as blur estimation, motion registration, regularization, dictionary-based-learning, etc. Nevertheless, with the help of our supervisor, we were able to cope and to make sense of the whole image processing landscape. We then sought ways to incorporate our ideas to the problem at hand. At this stage, we had to spend more time with MATLAB but, even here, we were facing a lot of technical difficulties. Thankfully, we had already gotten used to dealing with highly complex problems. We simply had to go ahead, and our progress slowly started to emerge. We also had to cope with academic pressure in parallel, and the competition vibe was really pressing. After months of trying various things such as Laplacian, total variation, etc., we went one step forward and tried to invent our own regularization term. Eventually, we settled on a wavelet-based sparsity constraint, which gave better results than our previous endeavors. By that time, the submission deadline had approached, so we started to wrap up our research in a paper format.

Even now, when we glance at the paper, we feel inspired as it bears testament to the adventurous research experience we had with the Signal Processing Cup 2014.

—Emroz Khan, Shiekh Zia Uddin, and Mukhlasur Rahman Tanvir, Undergrads

—Md. Kamrul Hasan, Supervisor

[SP]

H. Joel Trussell and Dror Baron

Creating Analytic Online Homework for Digital Signal Processing

An article by W.L. Everitt in the 1962 50th anniversary issue of *Proceedings of the IEEE*, "Engineering Education–Circa 2012 A.D.," was one of many predictive articles that appeared in that issue [1]. One of Everitt's observations was the distinction between training and education. He then predicted that, in the future, training will be done primarily with computers, remarking, "Relieved of the necessity of spending most of their time on the training function, devoted teachers will be able to concentrate their efforts on 'education'."

The distinction Everitt makes can be briefly presented to define what we are trying to accomplish. Training is concerned with learning methods, vocabularies, computational skills, and mathematical manipulations. It is reinforced by solving problems with definite solutions or measures of performance. Education is broader and is concerned with creativity, understanding, and the ability to apply the trained skills to meet new situations.

We are now reaching the point of Everitt's prediction, where we can create and grade problems on the computer that are similar to those training problems used in engineering texts. With the increasing use of the Internet for course content delivery, both for standard on-campus courses and massive open online courses, this capability becomes critically important.

For engineering, there are several online course/learning management systems that are in use. Included among these are Moodle, WebAssign [2], LON-CAPA [3], WeBWorK [4], MasteringEngineering (Pearson Publishing) [5], and Connect Engineering (McGraw-Hill Publishing) [6]. All of these systems manage the administrative aspects of classes, including maintaining grade books, reporting grades, posting assignments and messages, tracking students' progress, and presenting and grading problems. Currently, only WeBWorK, WebAssign, and LON-CAPA allow the creation of the analytical problems that are required for engineering courses.

Of the three systems that can create the required problems, WebAssign is the most widely used at approximately 2,300 institutions. WeBWork and LON-CAPA are used at about 300 and 150 institutions, respectively. It is unclear how many instructors use which features of the systems. None of the systems have found extensive application to electrical and computer engineering education. We will give a brief comparison of the systems later.

We used WeBWorK to help realize Everitt's prediction. This system was available at North Carolina State University (NCSU) and was free to use for both students and instructors. WebAssign is available at NCSU, but at a cost to the students. Since WeBWorK was used by our Department of Mathematics, there was good local experience with problems similar to those for electrical and computer engineering. LON-CAPA is not available at NCSU. Since the administrative capabilities of the systems are about the same, we limit our discussion to introducing the types of problems that can be presented to electrical and computer engineering students in general, and signal processing students in particular.

REALISTIC PROBLEMS

All of the most widely used online homework and testing packages allow the common types of questions, including true/false or yes/no, simple calculated numerical answer, multiple choice or multiple select (several possible correct answers), matching, fill-in-the-blank, etc. However, a problem form of major interest to most engineers is one that asks the student to find a mathematical relationship between the input and output of a system. This relationship is easily written as a formula or equation.

In the Fall 2013 semester, we used WeBWorK extensively for a sophomore electrical and computer engineering class at NCSU, called ECE220. The course covers basic math, beyond calculus, needed for electrical and computer engineering (ECE). This online homework approach was used in subsequent semesters of this course with quite positive evaluations from the students. To follow on this success, we developed problems for our senior-level digital signal processing course, ECE421.

To illustrate the type of problems that we created, let us consider the solution to second-order differential equations, which could represent a variety of circuit, control, or signal processing problems. We often want to see the various steps that a student makes to obtain a solution. For example, with given initial conditions, we ask the student for the form of the complementary solution, the particular solution, and the total solution. All of these are functional forms and the complementary solution has unknown constants that are determined only in the final stages of solving for the total solution. Such a differential equation problem is shown in Figure 1.

The student enters the answers using a format that resembles MATLAB. Figure 2 shows the student's input and its interpretation by the parser. The package

Digital Object Identifier 10.1109/MSP.2015.2438992
Date of publication: 13 August 2015

This problem is related to problems 7.10–7.18 in the text.

Instructions for forms of answers in differential equation problems.

For second-order DEs, the roots of the characteristic equation may be real or complex. If the roots are real, the complementary solution is the weighted sum of real exponentials. Use C1 and C2 for the weights, where C1 is associated with the root with smaller magnitude. If the roots are complex, the complementary solution is the weighted sum of complex conjugate exponentials, which can be written as a constant times a decaying exponential times a cosine with phase. Use C1 for the constant and Phi for the phase. (Note: Some equations in the text give the constant multiplying the decaying exponential as 2C1. This was done for the derivation. The constant for this problem should be C1 alone.)

All numerical angles (phases) should be given in radian angles (not degrees).

Given the differential equation $y'' + 12y' + 40y = 5cos(3t + 0.785398)u(t)$.

a. Write the functional form of the complementary solution, $y_c(t)$.

$y_c(t) =$ help (formulas)

b. Find the particular solution, $y_p(t)$.

$y_p(t) =$ help (formulas)

c. Find the total solution, $y(t)$ for the initial condition $y(0) = 5$ and $y'(0) = 5$.

$y(t) =$ help (formulas)

Note: You can earn partial credit on this problem.

Get a New Version of This Problem

Edit2 Edit3

Show correct answers

Preview Answers Check Answers

You have attempted this problem 0 times.

This homework set is closed.

Show Past Answers

E-mail instructor

[FIG1] The WeBWorK problem format.

will give hints as to problems with the format, e.g., unbalanced parentheses, unknown variables.

Problems can be created in a tutorial mode that gives the student step-by-step instructions on solving the problems. Alternatively, they can be written to check only the final answer. It is common that instructors wish to evaluate the students' written solutions to determine their understanding and pinpoint where mistakes are made. The automated approach can check the steps of the solution, if desired, as in the case of Figure 1. It cannot indicate what error the student made if he/she missed a part of the problem. However, we have found that by giving the students immediate feedback that answers are incorrect, students are usually able to determine their mistakes and proceed to the correct answer. This would seem preferable to giving detailed feedback on the error several days after the student turns in a written solution.

EXAMPLES AND PRACTICE

The digital signal processing (DSP) course at NCSU uses the text by Proakis and Manolakis [7]. We developed problems based on those in the text and have referenced the text problem number in the problem description. This will also allow the reader to observe our translation from text problems to realistic online problems. For examples of a typical z-transform problem, discrete Fourier transform problem, and a problem that has a purely symbolic answer, see Figures 3, 4, and 5, respectively.

In these problems, the various parameters are randomized but guaranteed to give reasonable formulations and answers.

Entered	Answer Preview
C1*exp(–6*t)*cos(2*t+Phi)*u(t)	$C_1 \exp(-6t)\cos(2t + \text{Phi})u(t)$
0.1052*cos(3*t–0.07449)*u(t)	$0.1052\cos(3t - 0.07449)\,u(t)$
[17.857*exp(–6*t)*cos(2*t – 1.2931) + 0.1052*cos(3*t – 0.07449)]*u(t)	$(17.857\exp(-6t)\cos(2t - 1.2931) + 0.1052\cos(3t - 0.07449))u(t)$

[FIG2] The WeBWorK answer format.

This problem is related to Problem 3.7 (page 215) in the text.
Consider the signal given by

$$x(n) = \begin{cases} \left(\frac{1}{2}\right)^n & n \geq 0 \\ \left(\frac{2}{8}\right)^{-n} & n < 0 \end{cases}$$

and the impulse response

$$h(n) = \begin{cases} \left(\frac{6}{9}\right)^n & n \geq 0 \\ 0 & n < 0 \end{cases}$$

Compute the convolution of the signal with the impulse response, $y(n) = x(n) * h(n)$.

For $n \geq 0$, $y(n) =$ [ ] help (formulas)

For $n < 0$, $y(n) =$ [ ] help (formulas)

Determine its region of convergence (ROC). Write the ROC as an interval.

ROC = [ ] help (intervals)

[FIG3] The *z*-transform problem.

This problem is related to Problem 5.23 (page 370) in the text.
This frequency response of an ideal bandpass filter is given by

$$H(\omega) = \begin{cases} 0 & |\omega| \leq \frac{4\pi}{16} \\ 1 & \frac{4\pi}{16} < |\omega| \leq \frac{10\pi}{16} \\ 0 & \frac{10\pi}{16} < |\omega| \leq \pi \end{cases}$$

Determine its impulse response $h(n)$

$h(n) =$ [ ] help (formulas)

Show this impulse response can be expressed as the product of the impulse response of an ideal low-pass filter, $h1(n)$ and $\cos(n14\pi/(2 * 16))$.

$h1(n) =$ [ ] help (formulas)

Note: You can earn partial credit on this problem.

Get a New Version of This Problem

[FIG4] The discrete Fourier transform problem.

The randomization of the parameters has the advantage of presenting each student with a different problem. The randomization also allows the students to use the same problem for additional practice. Note the "Get a new version of this problem" button at the bottom of Figure 4. This button appears after the student has submitted his or her answer to the assignment. The online solutions for the problems are written to use the parameters appropriate to each student. This allows students to see exactly how their problems are solved without having to translate from a common example in the text.

Each answer is noted as correct or not. The instructor sets the number of attempts that the student may use prior to submission of the answer for grading. The ability to check answers before submission has been well received by the students and helps them to identify their own errors. The students can use the "preview" button (seen in Figure 1) to check their answers for typos.

At the bottom of the problem (see Figure 1), there is a button to "e-mail the instructor." This allows the student who is stuck or has questions to ask the instructor for help. The student explains his problem in the e-mail, which is sent to the instructor and support group (teaching assistants or other faculty). The e-mail contains a link to the student's problem. This allows the instructor to see exactly what the student has entered and to formulate an answer to the student's question.

GRAPHICS

Problems for DSP homework and tests frequently refer to figures, such as block diagrams, time-domain signals, frequency-domain spectra, pole-zero plots, etc. As is common for most packages, WeBWorK allows the insertion of static graphics that are generated offline. In addition, WeBWorK allows dynamic graphics where the graph depends on the random parameters for a specific problem. This enables many more types of realistic problems that ask the student to obtain information from the graph.

A frequent assignment is to sketch a graph of a function that has been obtained.

Plots might be waveforms, frequency responses, regions of convergence, and the like. Currently, there is not a simple way for the student to enter a sketch of the function that can be evaluated for correctness. There is an app that can be embedded in WeBWorK that allows the student to input a curve in a graph with labeled axes. The student enters the curve using the stylus with a tablet or using the mouse of a standard computer. The correctness of the student's sketch is measured by a combination of maximum pointwise error and cumulative error. An example is shown in Figure 6, where the student is asked to sketch the derivative of the function in the top graph. The differences that determine the error are shown in the green and yellow areas. Note the "smooth" button allows entering rough sketches, followed by smoothing to assist in producing better curves. While this indicates a promising capability, we have not yet incorporated it into our problems.

We note that both WebAssign and LON-CAPA allow student-generated graphics. The examples that we have seen are

This problem is related to 4.17a in the text.

Let $x(n)$ be an arbitrary signal, not necessarily real valued, with discrete Fourier transform, $X(\omega)$, express the discrete Fourier transform of the following signal in terms of $X(w)$, Use w for ω(omega) to make typing easy.

$$y(n) = 5x(n) + 8x(n - 9),$$

$Y(\omega) =$ [ ]

[FIG5] The purely symbolic answer problem.

Entered	Answer Preview	Correct	Result
1	1	1	correct

The answer above is correct.

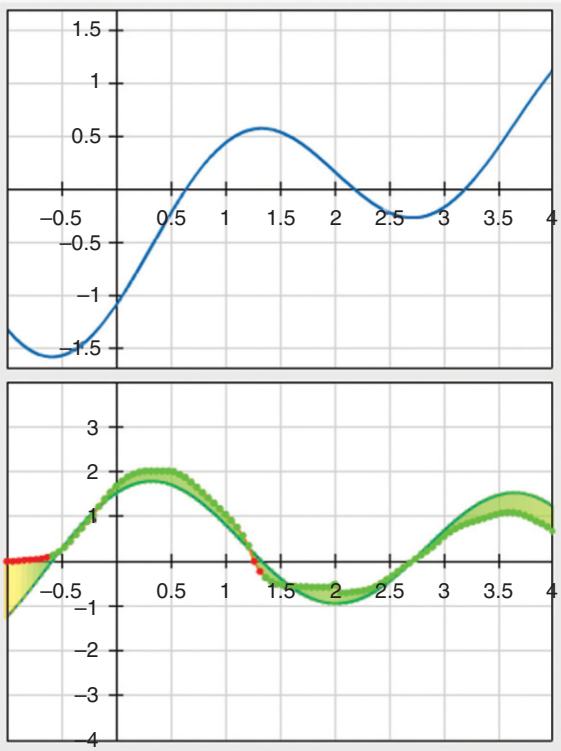

Derivative Sketching

Instructions: A function f(x) is shown in the upper graph. Sketch an accurate graph of the derivative of this function on the axes below. When you are satisfied with your sketch, click on the "Submit" button below.

Press and hold the "smooth" button to gradually smooth out your graph. Click on the "auto" button to turn on or off automatic smoothing.

Score:
signs (pos/neg): 90
%accuracy: 87%

Try Again

smooth
auto

This applet makes use of graphing classes developed by B. Kaskosz and D. Ensley, published at flashandmath.com.

This work was supported in part by the National Science Foundation under grant DUE-0941388.

[FIG6] The student sketches graph.

limited to very simple manipulations. These would have difficulty sketching a sinusoid or exponential. This ability will likely improve for all three systems.

PROGRAMMING THE PROBLEMS

Although the new problem generation and answer checking packages are able to handle a large variety of practical problems associated with undergraduate electrical and computer engineering, a major hurdle for the current packages is the steep learning curve required to generate online problems for a particular course.

Our experience in developing problems is with WeBWorK. This package is based on the Perl computer language and has been greatly augmented by the WeBWorK developers to handle the many cases required for mathematical problem development. While using templates from the large online library helped, we note there are many characteristics of engineering problems that are sufficiently different from the available math-based problems. Notations and problem motivations were typical of these. It required moderate editing to reformulate the problems for engineers. Fortunately, the developers had already solved the "$j = \sqrt{-1}$" versus "$i = \sqrt{-1}$" notation problem. We needed to insert the change into all of our problems.

The difficulty of developing problems is somewhat alleviated by the libraries of problems available from those who have learned the systems and developed their own problem sets. Thus, we often used problems from the WeBWorK Open Library as a starting point for our problems. An example of the code for Figure 4 is given in Figure 7. The language is basically Perl. The "#" symbol denotes a comment.

COMPARISON OF SYSTEMS AND RESOURCES

In general, WeBWork, WebAssign, and LON-CAPA have the ability to create similar problems that can be used for signal processing courses. We have extensive experience with WeBWorK, as mentioned earlier. This allows us to indicate the type of problems that can be created and are useful for our courses. We have verified that most of the features that we use in WeBWork are available in the other systems. The major features that we use include control of randomization of parameters, analytic formulas as answers, dynamic graphics, complex mathematics and easy input of complicated mathematical formulas. It is important that the instructor can interact with the students having questions and can see the exact problem and answer of the student.

There are major differences in the "look and feel" of the systems. Having extensive experience with WeBWorK and limited experience with the other two, it seems that all require much learning to use effectively. Both WebAssign and LON-CAPA have graphical user interfaces (GUIs) that allow the novice user to more easily create simple problems. However, using the GUIs may slow the experienced user when creating more complicated problems. We would liken the difference to people writing engineering papers in Microsoft Word and Latex. Once you learn the language, it is more efficient to write technical papers in Latex.

SINCE IT IS EASY TO START CODING PROBLEMS BY USING A SIMILAR PROBLEM AS A TEMPLATE, THE AVAILABILITY OF USER LIBRARIES IS VERY HELPFUL.

The languages of the three systems are different. WeBWork is Perl based. WebAssign and LON-CAPA use HTML or XML. For WebAssign and LON-CAPA, coding in these languages is almost required for complicated problems and associated answer checking. All systems allow the user to create answer checkers for specific problems with unique requirements.

It is possible to place restrictions on the form of the answers in all three systems. For example, in our course, we require an answer that contains a sinusoid to be entered in the form of a cosine with phase. We can disallow the use of the sine-plus-cosine form or the sum of complex exponentials form. This is relatively easy in WeBWorK by disabling a function with a simple command.

Documentation is important for all software systems. Since WebAssign and LON-CAPA are commercial systems, their documentations are very good. There is much available on the web. WeBWorK can be frustrating because of poor documentation. There are websites for many topics, and users can get additional information from the user forums. One approach to overcome the documentation problem is to use a search engine with the input: webwork topic. Since WeBWorK is open source, this works pretty well. This may not work as well with the commercial packages.

WebAssign has a large collection of texts that are keyed to its use. They cover a range of technical areas in the sciences and engineering, although engineering is represented by texts in only materials, mechanics, and circuits. Mathematics has over 400 texts and thousands of problems ranging from elementary to differential equations. We know of no texts that are keyed to WeBWorK or LON-CAPA.

Since it is easy to start coding problems by using a similar problem as a template, the availability of user libraries is very helpful. WeBWorK has an open problem library that is primarily mathematics, as is expected, since the package is supported by the Mathematical Association of America. There are small problem sets in physics, electric circuits, and mechanics. Many of the problem sets are keyed to specific texts like our examples of DSP problems. There is current work in developing problems for electrical and computer engineering at Louisiana Tech University and Northern Arizona University. Because of the nature of the commercial systems, access to previously coded problems is more difficult for WebAssign or LON-CAPA.

Finally, we should mention something about cost. WeBWorK is open source and freely distributed. There is, of course, a small but real cost since the institution must allocate resources to install and maintain the software on their servers. Use of the package is available for a US$200–300 cost via the Mathematical Association of America. LON-CAPA is also open source, similar to WeBWorK. It is

```
# DESCRIPTION
# Problem from 'Digital Signal Processing' Proakis and Manolakis, 4th ed.
# WeBWorK problem written by Joel Trussell, <hjt@ncsu.edu>
# ENDDESCRIPTION

## DBsubject(Digital Signal Processing)
## DBsection(Problems)
## Institution(North Carolina State University)
## Author(H. J. Trussell)
## TitleText1('Digital Signal Processing: Principles, Algorithms and Applications')
## AuthorText1('Proakis and Manolakis)
## EditionText1('4')
## Problem1('5.23)
##############################
#  Initialization

DOCUMENT();
# load macro functions for Webwork
loadMacros(
"PGstandard.pl",
"MathObjects.pl",
"AnswerFormatHelp.pl",
"PGcourse.pl"
);

TEXT(beginproblem());
############################
Context("Numeric");  # webwork context for math evaluations
Context()->variables->are( n=>"Real" );
# set function tolerance parameters
# always use absolute tolerance with signal processing problems
# since the random test points may result in values near zero
Context()->flags->set(  tolerance => 0.001,  tolType => "absolute");

$d1 = 16;  # base denominator, the following are numerators.
$n1 = random(2, 8, 2);
$n2 = random($n1+2, 14, 2);
$n4 = ($n2-$n1);
$n3 = ($n2+$n1);

# describe frequency response to take the inverse DFT of
#  formula  int (-b)^b exp(j*w*n) dw -  int (-a)^a exp(j*w*n) dw
#   (2/n)*2*sin((b-a)*n/2)*cos((b+a)*n/2)
#  a = n1*pi/d1; b = n2*pi/d1
#  write using fractions

$text_formula = "2*sin(($n2-$n1)*pi*n/(2*$d1))*cos(($n2+$n1)*pi*n/(2*$d1))/(pi*n)";
$answera = Formula($text_formula)->reduce;
# the reduce feature eliminates terms with zero multipliers and
# a coefficient of unity - among other things
$answerb = Formula("2*sin(($n2-$n1)*pi*n/(2*$d1))/(pi*n)")->reduce;

############################
#  Main text1

Context()->texStrings;
BEGIN_TEXT
This problem is related to Problem 5.23 (page 370) in the text.
$PAR
The frequency response of an ideal bandpass filter is given by
 \[ H(\omega)  = \left\lbrace \begin{array}{ l l }
 0 & |\omega | \leq \frac{$n1 \pi}{$d1} \\
1 & \frac{$n1 \pi}{$d1} <  |\omega |  \leq  \frac{$n2 \pi}{$d1} \\
0 &  \frac{$n2 \pi}{$d1} <  |\omega |  \leq   \pi  \end{array} \right. \]
$PAR Determine its impulse response \(h(n)\)
$BR
\( h(n) = \)  \{ ans_rule(60) \} \{ AnswerFormatHelp("formulas") \}

$PAR Show this impulse response can be expressed as the product
of the impulse response of an ideal low-pass filter, \(h1(n)\) and
\(\cos(n $n3 \pi / (2*$d1) ).\)
$BR
 \( h1(n) \) = \{ ans_rule(60) \}  \{ AnswerFormatHelp("formulas") \}
END_TEXT
Context()->normalStrings;
############################
#  Answer evaluation1

$showPartialCorrectAnswers = 1; # allows students partial credit

ANS( $answera->cmp() );  # answer checker for first blank
ANS( $answerb->cmp() );   # answer checker for second  blank

ENDDOCUMENT();
```

[FIG7] An example of WeBWork code to produce Figure 4.

the driver for the commercial venture, CourseWeaver [8]. We are not familiar with the licenses for the commercial systems WebAssign and CourseWeaver and must refer the reader to their websites.

STUDENT EVALUATIONS AND COMMENTS

In the first semester that we used WeBWorK problems for ECE220, all homework was submitted online. The students were very positive and gave the system a 4.1/5 rating. The course was offered with two sections the following spring and during the summer session. While no formal questions were put on the course evaluation survey, the response via comments to the instructors was quite positive. The evaluation for Fall 2014 yielded a 4.0/5 rating. A trial run was made in ECE421 for the Fall 2014 class. Student comments were again positive. Related problems were used for BME311 (Linear Systems in Biomedical Engineering) in Fall 2014. WeBWorK was expanded to a more tutorial environment, where if students submitted a few wrong answers, they were directed to series of more elementary tutorial problems, before returning to the main problem thread. These students provided overwhelmingly positive feedback, giving 85% strongly agree or agree responses to questions on the effectiveness of the tutorial.

The students like the immediate feedback of correctness. They usually are able to determine their errors and submit a correct answer in a few attempts. We allow multiple attempts for analytic problems, since typographical errors are common, as are arithmetic errors. In the cases where the students have problems, they appreciate the easy access to the instructor via the "e-mail" button. We have found the e-mail question load to be quite small. With almost 200 students enrolled in ECE220 for Spring 2014, we averaged under nine e-mails per week. Students were able to use the problems for review for the final exam, since the randomization offered practice problems.

Students did not like long problems that had only a single answer. They preferred multiple step problems, where intermediate results were required, similar to Figure 1. Even in the multiple step problems, a similar problem was that each submission counted toward the total number for the problem. This means that taking five attempts for part one reduces the number of attempts for part two by five. Our approach has been to allow proportionately more attempts for multiple step problems. WebAssign allows the instructor to set the number of attempts for each part, but this has its own problems.

WE HOPE THAT READERS WILL RECOGNIZE THE POSSIBILITIES AND JOIN IN THE CREATION OF LIBRARIES THAT CAN BE EASILY USED BY THE MAJORITY OF SIGNAL PROCESSING EDUCATORS.

While we informed students of the required accuracy of the solutions, students found using four decimal places for answers annoying. In all of the systems, the instructor can assign the type of accuracy required, relative or absolute. For signal processing problems, which have signals and functions that often have values near zero, absolute accuracy is preferred.

A LINK TO AN ONLINE DEMO

We have collected problems used in both ECE220 and ECE421 classes at NCSU in a WeBWorK problem set that can be accessed with a guest login. This will allow readers to get a better idea of the range of problems currently available, as well as ideas about how they may be modified to fit their particular needs. We have set the permissions to allow the problem code to be viewed for only a few problems. There is also an orientation problem set, which contains examples of how to use WeBWork, e.g., input answers, previewing answers, contact the instructor, etc. The link is http://webwork-jitar.math.ncsu.edu/webwork2/DemoCourse/.

CONCLUSIONS

Our goal is to demonstrate the ability to create online problems to train engineers in the mathematics of linear systems and signal processing. Such training requires that the students see realistic problems that illustrate the skills to be learned, they receive feedback on the performance of those skills, and they have the ability to continue practicing those skills. We hope that readers will recognize the possibilities and join in the creation of libraries that can be easily used by the majority of signal processing educators.

AUTHORS

H. Joel Trussell (hjt@ncsu.edu) received the B.S. degree from the Georgia Institute of Technology, the M.S. degree from Florida State University, and the Ph.D. degree from the University of New Mexico. He worked with the Los Alamos National Lab from 1969 to 1980 and is currently a professor in the Electrical and Computer Engineering Department at North Carolina State University, Raleigh.

Dror Baron (barondror@ncsu.edu) received the B.S. and M.S. degrees from the Technion–Israel Institute of Technology, and the Ph.D. degree from the University of Illinois at Urbana-Champaign, all in electrical engineering. He joined the Electrical and Computer Engineering Department at North Carolina State University in 2010, where he is an assistant professor.

REFERENCES

[1] W. L. Everitt, "Engineering education - circa 2012 A.D.," *Proc. IEEE*, vol. 50, no. 5, pp. 571–572, May 1962.

[2] WebAssign. [Online]. Available: http://webassign.com

[3] LON-CAPA Hosting and Support. [Online]. Available: http://www.courseweaver.com/lc.php

[4] WeBWorK. [Online]. Available: http://WeBWorK.maa.org

[5] MasteringEngineering. [Online]. Available: http://www.pearsonmylabandmastering.com/northamerica/masteringengineering/

[6] ConnectEngineering. [Online]. Available: http://connect.customer.mcgraw-hill.com/

[7] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms and Applications*, 4th ed. Prentice Hall: Upper Saddle River, NJ, 2007.

[8] [Online]. Available:http://www.courseweaver.com/ CourseWeaver East Lansing, MI

[SP]

3rd IEEE Global Conference on Signal & Infomation Processing

Orlando, Florida, USA December 14-16 2015

IEEE

CALL FOR PARTICIPATIONS

GlobalSIP'15 General Chairs: Jose Moura and Dapeng Oliver Wu
Technical Program Chairs: Mihaela van der Schaar, Xiaodong Wang, and Hsiao-Chun Wu
URL: http://2015.ieeeglobalsip.org/

The IEEE Global Conference on Signal and Information Processing (GlobalSIP) is a recently launched flagship conference of the *IEEE Signal Processing Society*. GlobalSIP'15 will be held in Orlando, Florida, USA, 14-16 December 2015. The conference will focus broadly on signal and information processing with an emphasis on up-and-coming signal processing themes. The conference will feature world-class speakers, tutorials, exhibits, and technical sessions consisting of poster or oral presentations. GlobalSIP'15 technical program will be comprised of a main program (General Symposium) and several collocated symposia on special topics. Technical papers submitted to GlobalSIP'15 address a number of topics:

- Signal processing in communications and networks, including green communication and Signal processing in optical communication
- Image and video processing
- Selective topics in speech and language processing
- Signal processing in security applications
- Signal processing in energy and power systems
- Signal processing in genomics and bioengineering (physiological, pharmacological and behavioral)
- Signal processing for social media networks
- Neural signal processing
- Seismic signal processing
- Hardware and real-time implementations
- Other novel and significant Applications of selected areas of signal processing

Symposia:

- General Symposium
- Symposium on Signal Processing on Graphics Processing Units and Multicores
- Symposium on Signal Processing in Mobile Multimedia Communication Systems
- Symposium on 3GPP EVS and beyond
- Symposium on Signal and Information Processing for Optimizing Future Energy Systems
- Symposium on Signal Processing Challenges in Human Brain Connectomics
- Symposium on Real-Time Signal Processing for Low-Cost and Low-Power Smart Devices
- Symposium on Signal Processing for Optical Wireless Communications
- Symposium on Signal and Info. Processing for Software-Defined Ecosystems, and Green Computing
- Symposium on Signal Processing Applications in Smart Buildings
- Symposium on Signal Processing and Mathematical Modeling of Biological Processes with Applications to Cyber-Physical Systems for Precise Medicine
- Symposium on Massive MIMO and Full-Dimension MIMO (FD-MIMO) Communications

Important Deadlines for Participations:

Author Registration (required for accepted papers)	5 September 2015
Advance Registration (discounted rate)	30 October 2015

Full registrations and one-day registrations at full rate cover lunches, breaks, and receptions.

Digital Object Identifier 10.1109/MSP.2015.2464391

lecture NOTES

A. Enis Cetin and
Mohammad Tofighi

Projection-Based Wavelet Denoising

In this lecture note, we describe a wavelet domain denoising method consisting of making orthogonal projections of wavelet (subbands) signals of the noisy signal onto an upside down pyramid-shaped region in a multidimensional space. Each horizontal slice of the upside down pyramid is a diamond shaped region and it is called an ℓ_1-ball. The upside down pyramid is called the epigraph set of the ℓ_1-norm cost function. We show that the method leads to soft-thresholding as in standard wavelet denoising methods. Orthogonal projection operations automatically determine the soft-threshold values of the wavelet signals.

PREREQUISITES

Prerequisites for understanding the material of this article are linear algebra, discrete-time signal processing, and wavelets. Orthogonal projection of a vector onto a hyperplane is the key mathematical operation used in this lecture note. Let w_o be a vector in $\mathbb{R}^K$. The orthogonal projection w_{po} of w_o onto the hyperplane $h = \mathrm{a}^T\mathrm{w} = \sum_{n=1}^{K} \mathrm{a}[n]\mathrm{w}[n]$ is given by

$$\mathrm{w}_{po}[n] = \mathrm{w}_o[n] + \frac{h - \sum_{n=1}^{K} \mathrm{a}[n]\mathrm{w}_o[n]}{\|\mathrm{a}\|_2^2}\,\mathrm{a}[n] \quad n = 1,2,\ldots,K, \tag{1}$$

where $\mathrm{w}_o[n]$, $\mathrm{w}_{po}[n]$, and $\mathrm{a}[n]$ are the nth entries of the vectors w_o, w_{po}, and a, respectively, and $\|\mathrm{a}\|_2$ is the Euclidean length (norm) of the vector a.

In this lecture note, orthogonal projections onto an upside down-shaped pyramid are computed. Each face of the upside down pyramid is a wedge-shaped subset of a hyperplane. Therefore, we can make an orthogonal projection onto an upside down pyramid by performing an orthogonal projection onto a face of the pyramid.

Orthogonal projection onto a hyperplane is also routinely used in the well-known normalized least mean squares (NLMS) adaptive filtering algorithm and many online learning algorithms [1].

PROBLEM STATEMENT

Denoising refers to the process of reducing noise in a given signal, image, and video. In standard wavelet denoising, a signal corrupted by additive noise is transformed to the wavelet domain and the resulting wavelet signals are soft- or hard-thresholded. After this step, the denoised signal is reconstructed from the thresholded wavelet signals [2], [3]. Thresholding wavelet coefficients intuitively makes sense because wavelet signals obtained from an orthogonal or biorthogonal wavelet filter bank exhibit large amplitude coefficients only around edges or jumps of the original signal. The assumption is that other small amplitude wavelet coefficients should be due to noise. Signals that can be represented with a small number of coefficients are called sparse signals and it turns out that most natural signals are sparse in some transfer domain [4], [5]. A wide range of wavelet denoising methods that take advantage of the sparse nature of practical signals in wavelet domain are developed using this baseline denoising idea by Donoho and Johnstone; see, e.g., [2]–[4] and [6]–[9].

Consider the following basic denoising framework. Let v be a discrete-time signal and x be its noisy version, i.e., $\mathrm{x}[n] = \mathrm{v}[n] + \boldsymbol{\xi}[n], n = 1,2,\ldots,N$, where $\boldsymbol{\xi}$ is some additive, independent and identically distributed (i.i.d.), zero-mean, white Gaussian noise with variance σ^2. An L-level discrete wavelet transform of x is computed and the lowband signal x_L and wavelet signals $\mathrm{w}_1, \mathrm{w}_2, \ldots, \mathrm{w}_L$ are obtained. After this step, wavelet signals are soft-thresholded as shown in Figure 1. The soft-threshold value, θ, can be selected in many ways using statistical methods [3], [4], [6], [10]. Donoho proposed the following threshold for all wavelet signals:

$$\theta = \gamma.\sigma.\sqrt{2\log(N)/N}, \tag{2}$$

where N is the number of samples of the signal x, and the constant γ is defined in [3]. In (2), the noise variance σ^2 has to be known or properly estimated from the observations, x, which may be difficult to achieve in practice. In [3], a single threshold is used for all wavelet signals. We refer the reader to [3], [4], [6], and [10] for many ways of estimating the parameters γ and σ in Donoho's method.

It is possible to define a soft-threshold θ_i for each wavelet signal w_i. Here we present how to estimate a soft-threshold value θ_i for each wavelet signal w_i using a deterministic approach based on linear algebra and orthogonal projections. In this approach, there is no need to estimate the variance σ^2. Thresholds are automatically determined by orthogonal projections onto an upside-down pyramid shaped region, which is the epigraph set of the ℓ_1-norm cost function.

WAVELET SIGNALS DENOISING WITH PROJECTIONS ONTO ℓ_1-BALLS

Let us first study the projection of wavelet signals $\mathrm{w}_1, \mathrm{w}_2, \ldots, \mathrm{w}_L$ onto ℓ_1-balls, which

Digital Object Identifier 10.1109/MSP.2015.2440051
Date of publication: 13 August 2015

1053-5888/15©2015IEEE

we will use to describe the projection onto the epigraph set of ℓ_1-norm cost function. We will use the term vector and signal in an interchangeable manner from now on. An ℓ_1-ball C_i, with size d_i is defined as follows:

$$C_i = \{\mathrm{w} \in \mathbb{R}^N : \sum_n |\mathrm{w}[n]| \le d_i\}, \quad (3)$$

where $\mathrm{w}[n]$ is the nth component of the vector w and d_i is the size of the ℓ_1-ball. In other words, an ℓ_1-ball is the set of vectors characterized by the fact that the sum of the magnitude of its components is lower than some specified value. Geometrically, such an ℓ_1-ball is a diamond shaped region bounded by a collection of hyperplanes as depicted in Figure 2. The orthogonal projection of a wavelet vector w_i onto an ℓ_1-ball is mathematically defined as follows:

$$\mathrm{w}_{pi} = \arg\min \|\mathrm{w}_i - \mathrm{w}\|_2^2$$
$$\text{such that } \|\mathrm{w}_i\|_1 = \sum_n |\mathrm{w}_i[n]| \le d_i, \quad (4)$$

where w_i is the ith wavelet signal, $\|\,.\,\|_2$ is the Euclidean norm, and $\|\,.\,\|_1$ is the ℓ_1-norm. The orthogonal projection operation onto an ℓ_1-ball is graphically shown in Figure 2. When $\|\mathrm{w}_i\|_1 \le d_i$ is satisfied, the wavelet signal is inside the ball, the projection has no effect and $\mathrm{w}_{pi} = \mathrm{w}_i$. In general, it can be shown that the orthogonal projection operation soft-thresholds each wavelet coefficient $\mathrm{w}_i[n]$ as follows:

$$\mathrm{w}_{pi}[n] = \mathrm{sign}(\mathrm{w}_i[n]) \max\{(|\mathrm{w}_i[n]| - \theta_i), 0\}, \quad (5)$$

where $\mathrm{sign}(\mathrm{w}_i[n])$ is the sign of $\mathrm{w}_i[n]$, and θ_i is a soft-thresholding constant whose value is determined according to the size of the ℓ_1-ball, d_i [11]. Algorithm 1 is an example of a method to solve the minimization problem (4) and thereby provide the constant θ_i for a given d_i value [11].

Projection of a wavelet signal onto an ℓ_1-ball reduces the amplitude of the wavelet coefficients of the input vector and eliminates the small valued wavelet coefficients, which are lower than the threshold θ_i. As a result, wavelet coefficients, which are probably due to noise, are removed by the projection operation. Projection operation onto an ℓ_1-ball retains the edges and sharp variation regions of the original signal because wavelet signals have large amplitude valued coefficients corresponding to edges [2] and they are not significantly affected by soft-thresholding. In standard wavelet denoising methods, the low-band signal x_L is not processed because x_L is a low resolution version of the original signal containing large amplitude coefficients almost for all n for most practical signals and images.

The next step is the estimation of the size of the ℓ_1-ball, d_i. We estimate the size of the ℓ_1-ball, d_i, by projecting w_i onto the epigraph set of the ℓ_1-norm cost function, which is an upside-down pyramid in $\mathbb{R}^{N+1}$ as shown in Figure 3. An upside-down pyramid is constructed by a family of ℓ_1-balls or diamond-shaped regions with different sizes ranging from 0 to $d_{\max,i} = \sum_n |\mathrm{w}_i[n]|$, whose value is the ℓ_1-norm of w_i. When we orthogonally project w_i onto the upside down pyramid, we not only estimate the size of the ℓ_1-ball, but also soft-threshold the wavelet signal w_i as discussed in the following section.

ESTIMATION OF DENOISING THRESHOLDS

The epigraph set of the ℓ_1-norm cost function is an upside-down pyramid shaped region as shown in Figure 3. Each horizontal slice of the upside down pyramid is an ℓ_1-ball. The smallest value of the ℓ_1-ball is 0, which is at the bottom of the pyramid. The largest value of the ℓ_1-ball in the upside-down pyramid is $d_{\max,i} = \|\mathrm{w}_i\|_1$, which is determined by the boundary of the ℓ_1-ball touching the wavelet signal w_i, i.e., the wavelet signal w_i is on one of the boundary hyperplanes of the ℓ_1-ball.

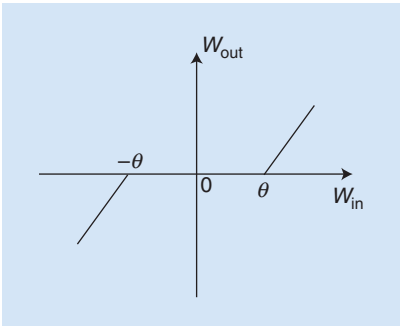

[FIG1] Soft-thresholding: $\mathrm{w}_{out}[n] = \mathrm{sign}(\mathrm{w}_{in}[n]) \max\{(|\mathrm{w}_{in}[n]| - \theta), 0\}$.

Orthogonal projection of w_i onto an ℓ_1-ball with $d = 0$ produces an all-zero result. Projection of w_i onto an ℓ_1-ball with size $d_{\max,i}$, does not change w_i because w_i is on the boundary of the ℓ_1-ball. Therefore, for meaningful results, the size of the ℓ_1-ball, $d_i = z_{pi}$, must satisfy the inequality $0 < z_{pi} < d_{\max,i}$, for denoising. This condition can be expressed as follows:

$$\|\mathrm{w}_i\|_1 = \sum_{k=1}^{K} |\mathrm{w}_i[k]| \le z_{pi}, \quad (6)$$

where K is the length of the wavelet vector $\mathrm{w} = [\mathrm{w}[1], \mathrm{w}[2], \ldots, \mathrm{w}[K]]^T \in \mathbb{R}^K$. The condition (6) corresponds to the epigraph set C of the ℓ_1-norm cost function in $\mathbb{R}^{K+1}$, which is graphically illustrated

---

Algorithm 1: Order $(K\log(K))$ algorithm implementing projection onto the ℓ_1-ball with size d_i.

---

1): **Inputs:**

A vector $\mathrm{w}_i = [\mathrm{w}_i[1], \ldots, \mathrm{w}_i[K]]$ and a scalar $d_i > 0$

2): **Initialize:**

Sort $|\mathrm{w}_i[n]|$ for $n = 1, \ldots, K$ and obtain the rank ordered sequence $\mu_1 \ge \mu_2 \ge, \ldots, \ge \mu_K$. The soft-threshold value, θ_i, is given by

$$\theta_i = \frac{1}{\rho}\left(\sum_{n=1}^{\rho} \mu_n - d_i\right) \text{ such that } \rho = \max\{j \in \{1,2,\ldots,K\} : \mu_j - \frac{1}{j}\left(\sum_{r=1}^{j} \mu_r - d_i\right) > 0\}$$

3): **Output:**

$$\mathrm{w}_{pi}[n] = \mathrm{sign}(\mathrm{w}_i[n]) \max\{|\mathrm{w}_i[n]| - \theta_i, 0\}, \; n = 1,2,\ldots,K$$

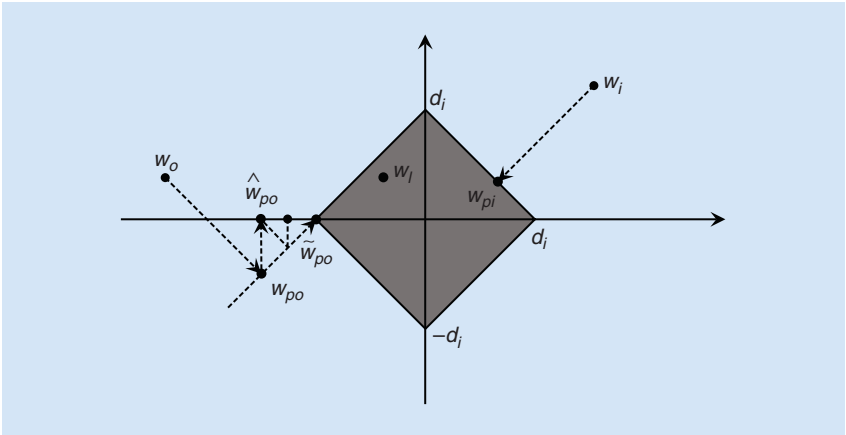

[FIG2] A graphical illustration of s projection onto an ℓ_1-ball with size d_i: Vectors w_{pi} and $\tilde{w}_{po}$ are orthogonal projections of w_i and w_o onto an ℓ_1-ball with size d_i, respectively. The vector w_l is inside the ball, $\|w_l\|_1 \le d_i$, and projection has no effect: $w_{pl} = w_l$.

in Figure 3 for $w_i \in \mathbb{R}^2$ [8], [9]. The epigraph set C is defined in $\mathbb{R}^{N+1}$ as follows:

$$C = \{\underline{w}_i = [w_i^T, z_{pi}]^T \in \mathbb{R}^{K+1} : \|w_i\|_1 \le z_{pi},\ z_{pi} \le d_{\max,i}\}, \quad (7)$$

which represents a family of ℓ_1-balls for $0 < z_{pi} \le d_{\max,i}$ in $\mathbb{R}^{K+1}$. In (7) there are $K+1$ variables: $w_i[1], \ldots, w_i[K]$, and z_{pi}. Since the space is now $K+1$ dimensional, we increase the size of wavelet signals by one:

$$\underline{w}_i = [w_i^T, 0]^T = [w_i[1], w_i[2], \ldots, w_i[K], 0]^T, \quad (8)$$

where $\underline{w}_i \in \mathbb{R}^{K+1}$. The signal $\underline{w}_i$ is the $K+1$ dimensional version of vector $w_i \in \mathbb{R}^K$. From now on, we underline vectors in $\mathbb{R}^{K+1}$ to distinguish them from K-dimensional vectors.

The extended wavelet vector $\underline{w}_i$ can be projected onto the epigraph set C to determine the vector $\underline{w}_{pi} = [w_{pi}[1], \ldots, w_{pi}[K], z_{pi}]^T$ as graphically illustrated in Figure 3. This projection is unique and is the closest vector on the epigraph set to $\underline{w}_i = [w_i^T, 0]^T$. The baseline mathematical operation is an orthogonal projection onto a hyperplane which is the face (boundary) of the epigraph set C in the quadrant of the $\underline{w}_i$. The orthogonal projection $\underline{w}_{pi}$ of $\underline{w}_i$ is a denoised version of $\underline{w}_i$ because it is equivalent to the orthogonal projection of w_i onto the ℓ_1-ball with size z_{pi} in $\mathbb{R}^K$, as graphically illustrated in Figure 3.

Orthogonal projection onto the epigraph set C can be computed in two steps. In the first step, $[w_i^T, 0]^T$ is projected onto the boundary hyperplane of the epigraph set which is defined as:

$$\sum_{n=1}^{K} \text{sign}(w_i[n]) . w_i[n] - z_{pi} = 0, \quad (9)$$

where the coefficients of the above hyperplane are determined according to the signs of $w_i[n]$. This hyperplane determines the boundary of the epigraph set C facing the vector $\underline{w}_i$ as shown in Figure 3. The projection vector $\underline{w}_{pi}$ onto the hyperplane (9) in $\mathbb{R}^{K+1}$ is determined using (1):

$$w_{pi}[n] = w_i[n] - \frac{\sum_{n=1}^{K} |w_i[n]|}{K+1} \text{sign}(w_i[n]) \quad n = 1, 2, \ldots, K, \quad (10)$$

where $K+1 = \|[\text{sign}(w_i[1]), \ldots, \text{sign}(w_i[k]), -1]\|^2$ and the last component z_{pi} of $\underline{w}_{pi}$ is given by

$$z_{pi} = \frac{\sum_{n=1}^{K} \text{sign}(w_i[n]) w_i[n]}{K+1} = \frac{\sum_{n=1}^{K} |w_i[n]|}{K+1}. \quad (11)$$

As mentioned earlier above, this orthogonal projection operation also determines the size of the ℓ_1-ball, $d_i = z_{pi}$, which can be verified using geometry.

In general, the projection vector $\underline{w}_{pi}$ may or may not be the projection of $\underline{w}_i$ onto the epigraph set C. In Figures 2 and 3, it is. The ℓ_1-ball in Figure 2 can be interpreted as the projection of 3-D ℓ_1-ball onto 2-D plane (view from the top). The issue comes from the fact that projecting onto the ℓ_1-ball has been simplified to projecting onto a single hyperplane, which may not yield the desired result in some specific geometrical configurations. For instance, in Figure 2, the vector w_{po} is neither the orthogonal projection of w_o onto the ℓ_1-ball, nor to the epigraph set of the ℓ_1-ball, because w_{po} is not on the ℓ_1-ball. Such cases can easily be spotted by checking the signs of the entries of the projection vectors. If the signs of the entries $w_{pi}[n]$ of projection

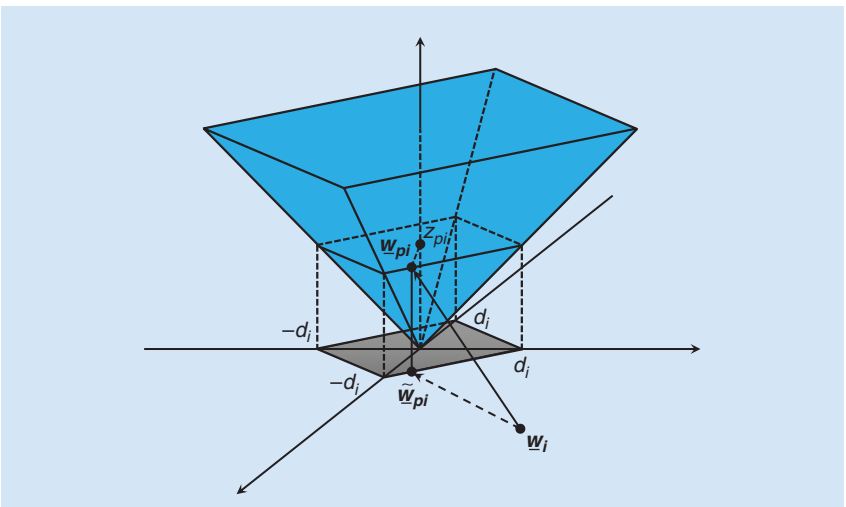

[FIG3] Projection of $w_i[n]$ onto the epigraph set of ℓ_1-norm cost function: $C = \{w : \sum_{n=0}^{K-1} |w[k]| \le z_{pi}\}$, gray-shaded region.

vector $\mathbf{w}_{pi}$ are the same as $\mathbf{w}_i[n]$ for all n, then the $\mathbf{w}_{pi}$ is on the epigraph set C, otherwise $\mathbf{w}_{pi}$ is not on the ℓ_1-ball. If $\mathbf{w}_{pi}$ is not on the ℓ_1-ball we can still project $\mathbf{w}_i$ onto the ℓ_1-ball using Algorithm 1 or Duchi et al's ℓ_1-ball projection algorithm [11] using the value of $d_i = z_{pi}$ determined in (11).

In summary, we have the following two steps: 1) project $\underline{\mathbf{w}}_i = [\mathbf{w}_i^T, 0]^T$ onto the boundary hyperplane of the epigraph set C and determine d_i using (11); and 2) if $\text{sign}(\mathbf{w}_i[n]) = \text{sign}(\mathbf{w}_{pi}[n])$ for all n, $\mathbf{w}_{pi}$ is the projection vector; otherwise, use $d_i = z_{pi}$ in Algorithm 1 to determine the final projection vector. Since there are $i = 1, 2, \ldots, L$ wavelet signals, each wavelet signal $\mathbf{w}_i$ should be projected onto possibly distinct ℓ_1-balls with sizes d_i. Notice that d_i is not the value of the soft-threshold, it is the size of the ℓ_1-ball. The value of the soft-threshold is determined using Algorithm 1.

In practice, we may further simplify step 2 in denoising applications. Our goal is to zero-out insignificant wavelet coefficients. Therefore, we compare signs of entries of $\mathbf{w}_{po}$ and $\mathbf{w}_o$. We can zero-out those entries whose signs change after the orthogonal projection. Step 2 is then becomes as is shown in (12) below.

This operation is also graphically illustrated in Figure 2. The vector $\mathbf{w}_o$ is projected onto the boundary hyperplane facing $\mathbf{w}_o$ to obtain $\mathbf{w}_{po}$, which then projected back to the quadrant of $\mathbf{w}_o$ to obtain the denoised version $\widehat{\mathbf{w}}_{po}$. This process can be iterated a couple of times to approach the orthogonal projection vector $\widetilde{\mathbf{w}}_{po}$ as shown in Figure 2.

Stronger denoising of the input vector is simply a matter of selecting a z_p value smaller than z_{pi} in (11). A z_p value closer to zero leads to a higher threshold and forces more wavelet coefficients to be zero after the projection operation.

IT IS ALSO POSSIBLE TO USE A PYRAMIDAL STRUCTURE FOR SIGNAL DECOMPOSITION INSTEAD OF THE WAVELET TRANSFORM.

HOW TO DETERMINE THE NUMBER OF WAVELET DECOMPOSITION LEVELS

There are many ways to estimate the number of wavelet decomposition levels [6]. It is also possible to use the Fourier transform of the noisy signal to approximately estimate the bandwidth of the signal. Once the bandwidth ω_0 of the original signal is approximately determined, it can be used to estimate the number of wavelet transform levels and the bandwidth of the low-band signal $\mathbf{x}_L$. In an L-level wavelet decomposition, the low-band signal $\mathbf{x}_L$ approximately comes from the $[0, (\pi/2^L)]$ frequency band of the signal $\mathbf{x}$. Therefore, $(\pi/2^L)$ must be comparable to ω_0 so that the actual signal components are not soft-thresholded. Only wavelet signals $\mathbf{w}_1, \ldots, \mathbf{w}_{L-1}, \mathbf{w}_L$, whose Fourier transforms approximately occupy the bands $[(\pi/2), \pi], \ldots, [(\pi/2^{L-1}), (\pi/2^{L-2})], [(\pi/2^L), (\pi/2^{L-1})]$, respectively, should be soft-thresholded in denoising. For example, consider the `cusp` signal defined in MATLAB. It is possible to estimate an approximate frequency value ω_0 for this signal. The `cusp` signal is corrupted by additive zero-mean white Gaussian noise with $\sigma = 20\%$ of the maximum amplitude of the original signal as shown in Figure 4. The magnitude of the Fourier transform of the `cusp` signal is shown in Figure 5. For this signal, an $L = 5$ level wavelet decomposition is suitable because the magnitude of the Fourier transform approaches the noise floor level at high frequencies after $\omega_0 \approx (\pi/46)$ as shown in Figure 5. Therefore, $L = 5$ $((\pi/2^5) > \omega_0)$ is selected as the number of wavelet decomposition levels.

It is also possible to use a pyramidal structure for signal decomposition instead of the wavelet transform. In this case, the noisy signal is filtered with a lowpass filter with a cut-off frequency of $(\pi/8)$ and the output $\mathbf{x}_{lp}$ is subtracted

$$\widehat{\mathbf{w}}_{po}[n] = \begin{cases} \mathbf{w}_{po}[n], & \text{if } \text{sign}(\mathbf{w}_{po}[n]) = \text{sign}(\mathbf{w}_o[n]) \\ 0, & \text{otherwise.} \end{cases} \tag{12}$$

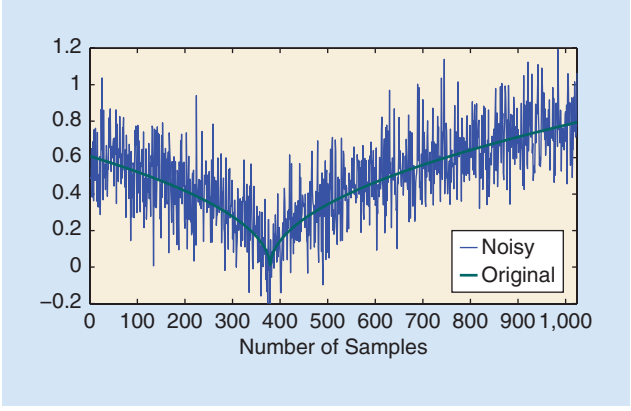

[FIG4] The `cusp` signal and its corrupted version with Gaussian noise with $\sigma = 20\%$ of maximum amplitude of the original signal.

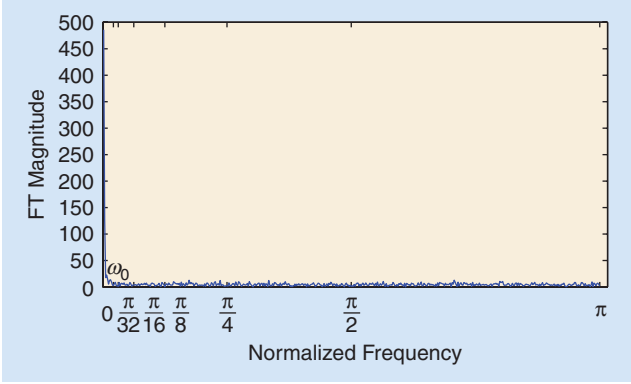

[FIG5] The discrete-time Fourier transform magnitude of `cusp` signal corrupted by noise. The wavelet decomposition level L is selected as five satisfying $(\pi/2^5) > \omega_0$, which is the approximate bandwidth of the signal.

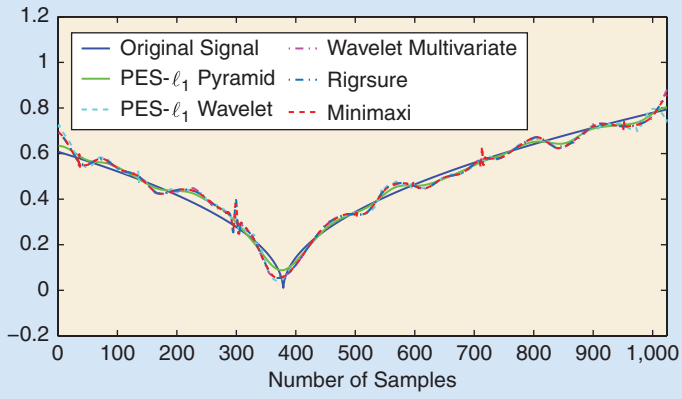

[FIG6] Original `cusp` signal (blue), denoised signal (green) using PES-ℓ_1-ball with pyramid; SNR = 28.26 dB, denoised signal (cyan) using PES-ℓ_1-ball with wavelet; SNR = 25.30 dB, denoised signal (magenta) using MATLAB wavelet multivariate method; SNR = 25.08 dB [6], denoised signal (petroleum blue) using wavelet denoising `rigrsure` algorithm [3]; SNR = 23.28 dB, and denoised signal (red) using wavelet denoising `minimaxi` algorithm [7]; SNR = 24.52 dB.

from the noisy signal x to obtain the high-pass signal $\mathbf{x}_{\text{hp}}$ as shown in [12]. The highpass signal $\mathbf{x}_{\text{hp}}$ is projected onto the epigraph of ℓ_1-norm cost function and the denoised signal $\mathbf{x}_{\text{hd}}$ is obtained. Projection onto the epigraph set of ℓ_1-ball (PES-ℓ_1) removes the noise by soft-thresholding. The pyramidal signal decomposition approach is similar to the undecimated wavelet transform. The denoised signal is reconstructed by adding $\mathbf{x}_{\text{hd}}$ and $\mathbf{x}_{\text{lp}}$.

In Figure 6, the signal is restored using PES-ℓ_1 with a pyramid structure, PES-ℓ_1 with wavelet, MATLAB's wavelet multivariate denoising algorithm [6], MATLAB's soft-thresholding denoising algorithms (`minimaxi` and `rigrsure` thresholds), and wavelet thresholding denoising method. The denoised signals have signal-to-noise (SNR) values equal to 28.26, 25.30, 25.08, 23.28, and 24.52 dB, respectively. In average, PES-ℓ_1 with pyramid and PES-ℓ_1 with wavelet method produce better denoising results than the other soft-thresholding methods. The SNR is calculated using the formula: $\text{SNR} = 20 \times \log_{10}(\|\mathbf{x}_{\text{orig}}\|/\|\mathbf{x}_{\text{orig}} - \mathbf{x}_{\text{rec}}\|)$. Extensive simulation results and the denoising software are available on the Internet [12].

CONCLUSIONS

PROS

Orthogonal projection-based denoising is computationally efficient because projection onto a boundary hyperplane of an ℓ_1-ball or the epigraph set can be implemented by performing only one division and $K+1$ additions and/or subtractions, and sign computations. Once the size of the ℓ_1-ball using (10) and (11) is determined, the orthogonal projection onto an ℓ_1-ball operation is an order (K) operation. Equations (10) and (11) only involve multiplications by ± 1.

CONS

It is not possible to incorporate any prior knowledge about the noise probability density function or any other statistical information to the orthogonal projection based denoising method. However, it produces good denoising results under additive white Gaussian noise. Most of the denoising methods available in MATLAB also assumes that the noise is additive, white Gaussian.

ACKNOWLEDGMENT

This work is funded by the Scientific and Technological Research Council of Turkey (TUBITAK) under project 113E069.

AUTHORS

A. Enis Cetin (cetin@bilkent.edu.tr) is a professor in the Department of Electrical and Electronics Engineering, Bilkent University, Ankara, Turkey. His main research interests are multimedia signal processing and its applications. He is a Fellow of the IEEE.

Mohammad Tofighi (tofighi@ee.bilkent.edu.tr) is an M.Sc. student in the Department of Electrical and Electronics Engineering, Bilkent University, Ankara, Turkey. His research interests include signal and image processing, inverse problems in signal processing, computer vision, pattern recognition, and machine learning. He is a Student Member of the IEEE.

REFERENCES

[1] K. Slavakis, S.-J. Kim, G. Mateos, and G. Giannakis, "Stochastic approximation vis-a-vis online learning for big data analytics," *IEEE Signal Processing Mag.*, vol. 31, no. 6, pp. 124–129, Nov. 2014.

[2] S. Mallat and W.-L. Hwang, "Singularity detection and processing with wavelets," *IEEE Trans. Inform. Theory*, vol. 38, no. 2, pp. 617–643, Mar. 1992.

[3] D. Donoho, "De-noising by soft-thresholding," *IEEE Trans. Inform. Theory*, vol. 41, no. 3, pp. 613–627, May 1995.

[4] S. Chang, B. Yu, and M. Vetterli, "Adaptive wavelet thresholding for image denoising and compression," *IEEE Trans. Image Processing*, vol. 9, no. 9, pp. 1532–1546, Sept. 2000.

[5] R. Baraniuk, "Compressive sensing," *IEEE Signal Processing Mag.*, vol. 24, no. 4, pp. 118–121, July 2007.

[6] M. Aminghafari, N. Cheze, and J.-M. Poggi, "Multivariate denoising using wavelets and principal component analysis," *Computat. Stat. Data Anal.*, vol. 50, no. 9, pp. 2381–2398, 2006.

[7] D. L. Donoho and I. M. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *J. Am. Stat. Assoc.*, vol. 90, no. 432, pp. 1200–1224, 1995.

[8] G. Chierchia, N. Pustelnik, J.-C. Pesquet, and B. Pesquet-Popescu, "Epigraphical projection and proximal tools for solving constrained convex optimization problems," *Signal, Image, Video Process.*, 2014, pp. 1–13.

[9] M. Tofighi, K. Kose, and A. E. Cetin, "Denoising using projections onto the epigraph set of convex cost functions," in *Proc. 2014 IEEE Int. Conf. Image Processing (ICIP)*, Oct. 2014, pp. 2709–2713.

[10] J. Fowler, "The redundant discrete wavelet transform and additive noise," *IEEE Signal Processing Lett.*, vol. 12, no. 9, pp. 629–632, Sept. 2005.

[11] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, "Efficient projections onto the l1-ball for learning in high dimensions," in *Proc. 25th Int. Conf. Machine Learning* (*ICML'08*). New York: ACM, 2008, pp. 272–279.

[12] (2014). PES-L1 Denoising Software. [Online]. Available: http://signal.ee.bilkent.edu.tr/1DDenoisingSoftware.html

SP

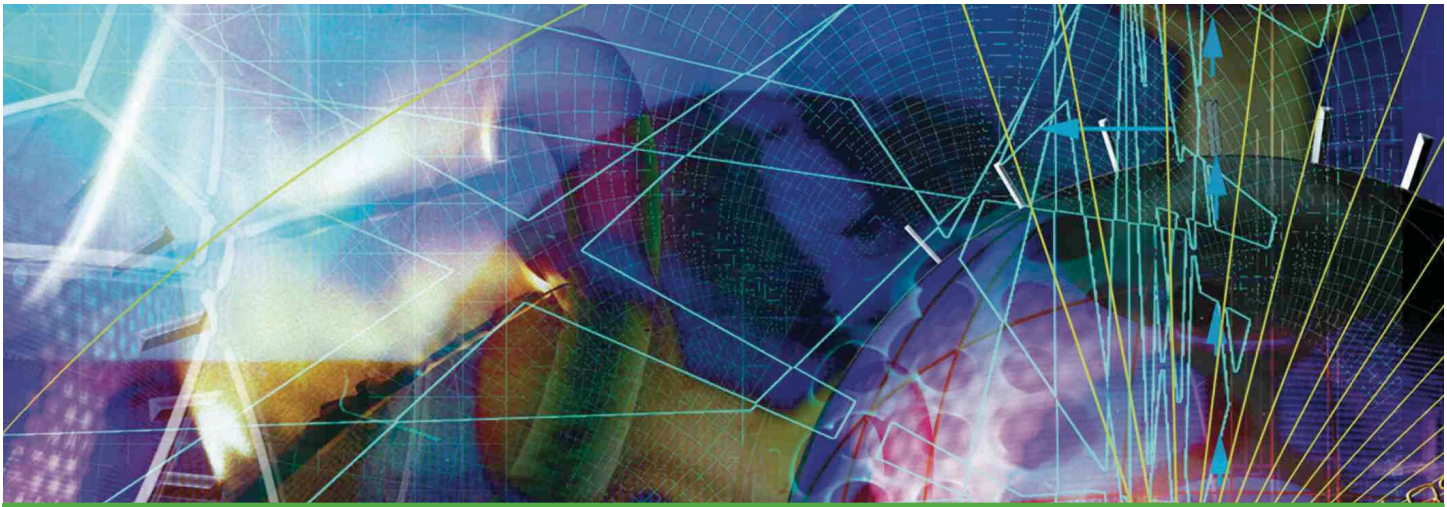

IEEE Signal Processing Society

Online Video Resources

SPS YouTube Channel

goo.gl/mUeOo9

SigView Technical Tutorials

SigView.org

SigPort Document Repository

SigPort.org

Social Links of Signal Processing Magazine

linkd.in/1aEgGXd

- **What is Signal Processing?**

  A perfect 2-minute start point to share with students, family, and friends

- **Signal Processing and Machine Learning**

  NEW! From brain signal processing to real-time speech translation, see cutting-edge technologies in action in a Hollywood style

- **Numerous Technical Tutorials on SigView**

  Missed ICASSP plenaries? Want to learn compressed sensing or big data? See them all FREE with your SPS member login

- **Highlights of Documents on SigPort**

  First ever global document repositories on signal and information processing

IEEE Signal Processing Society

IEEE

Digital Object Identifier 10.1109/MSP.2015.2464393

new

IEEE TRANSACTIONS ON

SIGNAL AND INFORMATION PROCESSING OVER NETWORKS

Now accepting paper submissions

The new *IEEE Transactions on Signal and Information Processing over Networks* publishes high-quality papers that extend the classical notions of processing of signals defined over vector spaces (e.g. time and space) to processing of signals and information (data) defined over networks, potentially dynamically varying. In signal processing over networks, the topology of the network may define structural relationships in the data, or may constrain processing of the data. Topics of interest include, but are not limited to the following:

Adaptation, Detection, Estimation, and Learning
- Distributed detection and estimation
- Distributed adaptation over networks
- Distributed learning over networks
- Distributed target tracking
- Bayesian learning; Bayesian signal processing
- Sequential learning over networks
- Decision making over networks
- Distributed dictionary learning
- Distributed game theoretic strategies
- Distributed information processing
- Graphical and kernel methods
- Consensus over network systems
- Optimization over network systems

Communications, Networking, and Sensing
- Distributed monitoring and sensing
- Signal processing for distributed communications and networking
- Signal processing for cooperative networking
- Signal processing for network security
- Optimal network signal processing and resource allocation

Modeling and Analysis
- Performance and bounds of methods
- Robustness and vulnerability
- Network modeling and identification

Modeling and Analysis (cont.)
- Simulations of networked information processing systems
- Social learning
- Bio-inspired network signal processing
- Epidemics and diffusion in populations

Imaging and Media Applications
- Image and video processing over networks
- Media cloud computing and communication
- Multimedia streaming and transport
- Social media computing and networking
- Signal processing for cyber-physical systems
- Wireless/mobile multimedia

Data Analysis
- Processing, analysis, and visualization of big data
- Signal and information processing for crowd computing
- Signal and information processing for the Internet of Things
- Emergence of behavior

Emerging topics and applications
- Emerging topics
- Applications in life sciences, ecology, energy, social networks, economic networks, finance, social sciences, smart grids, wireless health, robotics, transportation, and other areas of science and engineering

Editor-in-Chief: Petar M. Djurić, Stony Brook University (USA)

To submit a paper, go to: https://mc.manuscriptcentral.com/tsipn-ieee

IEEE COMMUNICATIONS SOCIETY

IEEE Signal Processing Society

IEEE COMPUTER SOCIETY

Digital Object Identifier 10.1109/MSP.2015.2464394

advertisers INDEX

The Advertisers Index contained in this issue is compiled as a service to our readers and advertisers: the publisher is not liable for errors or omissions although every effort is made to ensure its accuracy. Be sure to let our advertisers know you found them through *IEEE Signal Processing Magazine*.

ADVERTISER	PAGE	URL	PHONE
EPFL – Ecole Polytechnique Fererale de Lausanne	9	http://FoundationsOfSignalProcessing.org	
IEEE Marketing Dept.	6	innovate.ieee.org	
IEEE MDL/Marketing	3	www.ieee.org/go/trymdl	
IEEE USA	11	www.ieeeusa.org/policy/govfel	+1 202 530 8347
Mathworks	CVR 4	www.mathworks.com/ltc	+1 508 647 7040
Mini-Circuits	CVR 2, 5, CVR 3	www.minicircuits.com	+1 718 934 4500

advertisers SALES OFFICES

James A. Vick
Sr. Director, Advertising
Phone: +1 212 419 7767;
Fax: +1 212 419 7589
jv.ieeemedia@ieee.org

Marion Delaney
Advertising Sales Director
Phone: +1 415 863 4717;
Fax: +1 415 863 4717
md.ieeemedia@ieee.org

Mark David
Sr. Manager Advertising & Business Development
Phone: +1 732 465 6473
Fax: +1 732 981 1855
m.david@ieee.org

Mindy Belfer
Advertising Sales Coordinator
Phone: +1 732 562 3937
Fax: +1 732 981 1855
m.belfer@ieee.org

Product Advertising

MIDATLANTIC
Lisa Rinaldo
Phone: +1 732 772 0160;
Fax: +1 732 772 0164
lr.ieeemedia@ieee.org
NY, NJ, PA, DE, MD, DC, KY, WV

NEW ENGLAND/SOUTH CENTRAL/ EASTERN CANADA
Jody Estabrook
Phone: +1 774 283 4528;
Fax: +1 774 283 4527
je.ieeemedia@ieee.org
ME, VT, NH, MA, RI, CT, AR, LA, OK, TX
Canada: Quebec, Nova Scotia, Newfoundland, Prince Edward Island, New Brunswick

SOUTHEAST
Cathy Flynn
Phone: +1 770 645 2944;
Fax: +1 770 993 4423
cf.ieeemedia@ieee.org
VA, NC, SC, GA, FL, AL, MS, TN

MIDWEST/CENTRAL CANADA
Dave Jones
Phone: +1 708 442 5633;
Fax: +1 708 442 7620
dj.ieeemedia@ieee.org
IL, IA, KS, MN, MO, NE, ND, SD, WI, OH
Canada: Manitoba, Saskatchewan, Alberta

MIDWEST/ ONTARIO, CANADA
Will Hamilton
Phone: +1 269 381 2156;
Fax: +1 269 381 2556
wh.ieeemedia@ieee.org
IN, MI. Canada: Ontario

WEST COAST/MOUNTAIN STATES/ WESTERN CANADA
Marshall Rubin
Phone: +1 818 888 2407;
Fax: +1 818 888 4907
mr.ieeemedia@ieee.org
AZ, CO, HI, NM, NV, UT, AK, ID, MT, WY, OR, WA, CA. Canada: British Columbia

EUROPE/AFRICA/MIDDLE EAST ASIA/FAR EAST/PACIFIC RIM
Louise Smith
Phone: +44 1875 825 700;
Fax: +44 1875 825 701
les.ieeemedia@ieee.org
Europe, Africa, Middle East
Asia, Far East, Pacific Rim, Australia, New Zealand

Recruitment Advertising

MIDATLANTIC
Lisa Rinaldo
Phone: +1 732 772 0160;
Fax: +1 732 772 0164
lr.ieeemedia@ieee.org
NY, NJ, CT, PA, DE, MD, DC, KY, WV

NEW ENGLAND/EASTERN CANADA
Liza Reich
Phone: +1 212 419 7578;
Fax: +1 212 419 7589
e.reich@ieee.org
ME, VT, NH, MA, RI. Canada: Quebec, Nova Scotia, Prince Edward Island, Newfoundland, New Brunswick

SOUTHEAST
Cathy Flynn
Phone: +1 770 645 2944;
Fax: +1 770 993 4423
cf.ieeemedia@ieee.org
VA, NC, SC, GA, FL, AL, MS, TN

MIDWEST/SOUTH CENTRAL/ CENTRAL CANADA
Darcy Giovingo
Phone: +224 616 3034;
Fax: +1 847 729 4269
dg.ieeemedia@ieee.org;
AR, IL, IN, IA, KS, LA, MI, MN, MO, NE, ND, SD, OH, OK, TX, WI. Canada: Ontario, Manitoba, Saskatchewan, Alberta

WEST COAST/SOUTHWEST/ MOUNTAIN STATES/ASIA
Tim Matteson
Phone: +1 310 836 4064;
Fax: +1 310 836 4067
tm.ieeemedia@ieee.org
AZ, CO, HI, NV, NM, UT, CA, AK, ID, MT, WY, OR, WA. Canada: British Columbia

EUROPE/AFRICA/MIDDLE EAST
Louise Smith
Phone: +44 1875 825 700;
Fax: +44 1875 825 701
les.ieeemedia@ieee.org
Europe, Africa, Middle East

Digital Object Identifier 10.1109/MSP.2015.2388980

dates AHEAD

Please send calendar submissions to:
Dates Ahead, c/o Jessica Barragué
IEEE Signal Processing Magazine
445 Hoes Lane
Piscataway, NJ 08855 USA
e-mail: j.barrague@ieee.org

2015

AUGUST

12th IEEE International Conference on Advanced Video- and Signal-Based Surveillance (AVSS)
25–28 August, Karlsruhe, Germany.
General Chairs: Jürgen Beyerer and Rainer Stiefelhagen
URL: http://avss2015.org

2015 23rd European Signal Processing Conference (EUSIPCO)
31 August–4 September, Nice, France.
General Chairs: Jean-Luc Dugelay and Dirk Slock
URL: http://www.eusipco2015.org

SEPTEMBER

Ninth International Symposium on Image and Signal Processing Anaylsis (ISPA)
7–9 September, Zagreb, Croatia.
Honorary Cochairs: Sanjit K. Mitra and Tariq Durrani
General Cochairs: Sven Lončaric and Dick Lerski
URL: http://www.isispa.org/

IEEE Signal Processing Society Italy Chapter Summer School on Signal Processing (S3P)
7–11 September, Brescia, Italy.

Sensor Signal Processing for Defence (SSPD)
9–10 September, Edinburgh, United Kingdom.
General Chairs: Mike Davies, Jonathon Chambers, and Paul Thomas
URL: http://www.sspdconference.org

IEEE International Workshop on Machine Learning for Signal Processing (MLSP)
17–20 September, Boston, Massachusetts, United States.
General Chair: Deniz Erdogmus
URL: http://mlsp2015.conwiz.dk/home.htm

IEEE International Conference on Image Processing (ICIP)
28 September–1 October, Quebec City, Quebec, Canada.
General Chairs: Jean-Luc Dugelay and André Morin
URL: http://www.icip2015.org/

OCTOBER

IEEE International Conference on Ubiquitous Wireless Broadband (ICUWB)
4–7 October, Montreal, Canada.
Contact: info@icuwb2015.org
URL: http://www.icuwb2015.org/index.html

IEEE Workshop on Signal Processing Systems (SiPS)
14–16 October, Hangzhou, China.
General Chairs: Chaitali Chakrabarti and Nam Ling
URL: http://www.sips2015.org/

IEEE International Workshop on Multimedia Signal Processing (MMSP)
19–21 October, Xiamen, China.
General Chairs: Xiao-Ping Zhang, Oscar C. Au, and Jonathan Li
URL: http://www.mmsp2015.org/

IEEE International Conference on Signal and Image Processing Applications (ICSIPA)
19–21 October, Kuala Lumpur, Malaysia.
General Chair: Syed Khaleel
URL: http://spsocmalaysia.org/icsipa2015/

NOVEMBER

49th Asilomar Conference on Signals, Systems, and Computers (ACSSC)
8–11 November, Pacific Grove, California, United States.
General Chair: Erik G. Larsson
URL: http://www.asilomarsscconf.org/

Seventh IEEE International Workshop on Information Forensics and Security (WIFS)
16–19 November, Rome, Italy.
General Chairs: Patrizio Campisi and Nasir Memon
URL: http://www.wifs2015.org/

DECEMBER

IEEE 6th International Workshop on Computational Advances in Multisensor Adaptive Processing (CAMSAP)
13–16 December, Cancun, Mexico.
URL: http://inspire.rutgers.edu/camsap2015/

IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)
13–17 December, Scottsdale, Arizona, United States.
URL: http://www.asru2015.org/

International Conference on 3D Imaging (IC3D)
14–15 December, Liege, Belgium.
Contact: alain@3dstereomedia.eu
URL: http://www.3dstereomedia.eu/ic3d

IEEE Global Conference on Signal and Information Processing (GlobalSIP)
14–16 December, Orlando, Florida, United States.
General Chairs: José M.F. Moura and Dapeng Oliver Wu
URL: http://2015.ieeeglobalsip.org/

IEEE Second World Forum on Internet of Things (WF-IoT)
14–16 December, Milan, Italy.
Conference Chair: Latif Ladid
URL: http://sites.ieee.org/wf-iot/

Asia-Pacific Signal and Information Processing Association Annial Summit and Conference (APSIPA)
16–19 December, Hong Kong.
Honorary General Chair: Wan-Chi Siu
General Cochairs: Kenneth Lam, Helen Meng, and Oscar Au
URL: http://www.apsipa2015.org/

2016

MARCH

41st IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)
21–25 March, Shanghai, China.
General Chairs: Zhi Ding, Zhi-Quan Luo, and Wenjun Zhang
URL: http://icassp2016.org

Digital Object Identifier 10.1109/MSP.2015.2448373
Date of publication: 13 August 2015

SP

0.5-8GHz

New

Ultra-Wideband Low Noise Amplifiers

Low noise, high dynamic range, high output power, and flat gain from 0.5 to 8 GHz, all in a single model! And it's very easy to use: a simple fixed inductor at the input and output provide matching over the entire band!* It's ideal for sensitive, high-dynamic-range receivers, instrumentation, defense systems, LTE, WiFi, S-band and C-band radar, satcom and more! It operates on a single 5V supply and comes in a tiny 3x3mm MCLP package for excellent manufacturability. It's available off the shelf for a great value, so go to minicircuits.com and place your order today for delivery as soon as tomorrow!

*See datasheet for suggested application circuit.
†Flatness specified over 0.5 to 7 GHz.

Only $6.95 (qty. 1000)

FEATURES

Low Noise, 1.3 dB
High Gain, 21 dB
Excellent gain flatness, ±0.7 dB†
High IP3, +35 dBm
High POUT, +23.2 dBm
Tiny size, 3x3 mm

Mini-Circuits®

www.minicircuits.com P.O. Box 350166, Brooklyn, NY 11235-0003 (718) 934-4500 sales@minicircuits.com

541 rev org

IEEE SIGNAL PROCESSING SOCIETY

CONTENT GAZETTE

[ISSN 2167-5023]

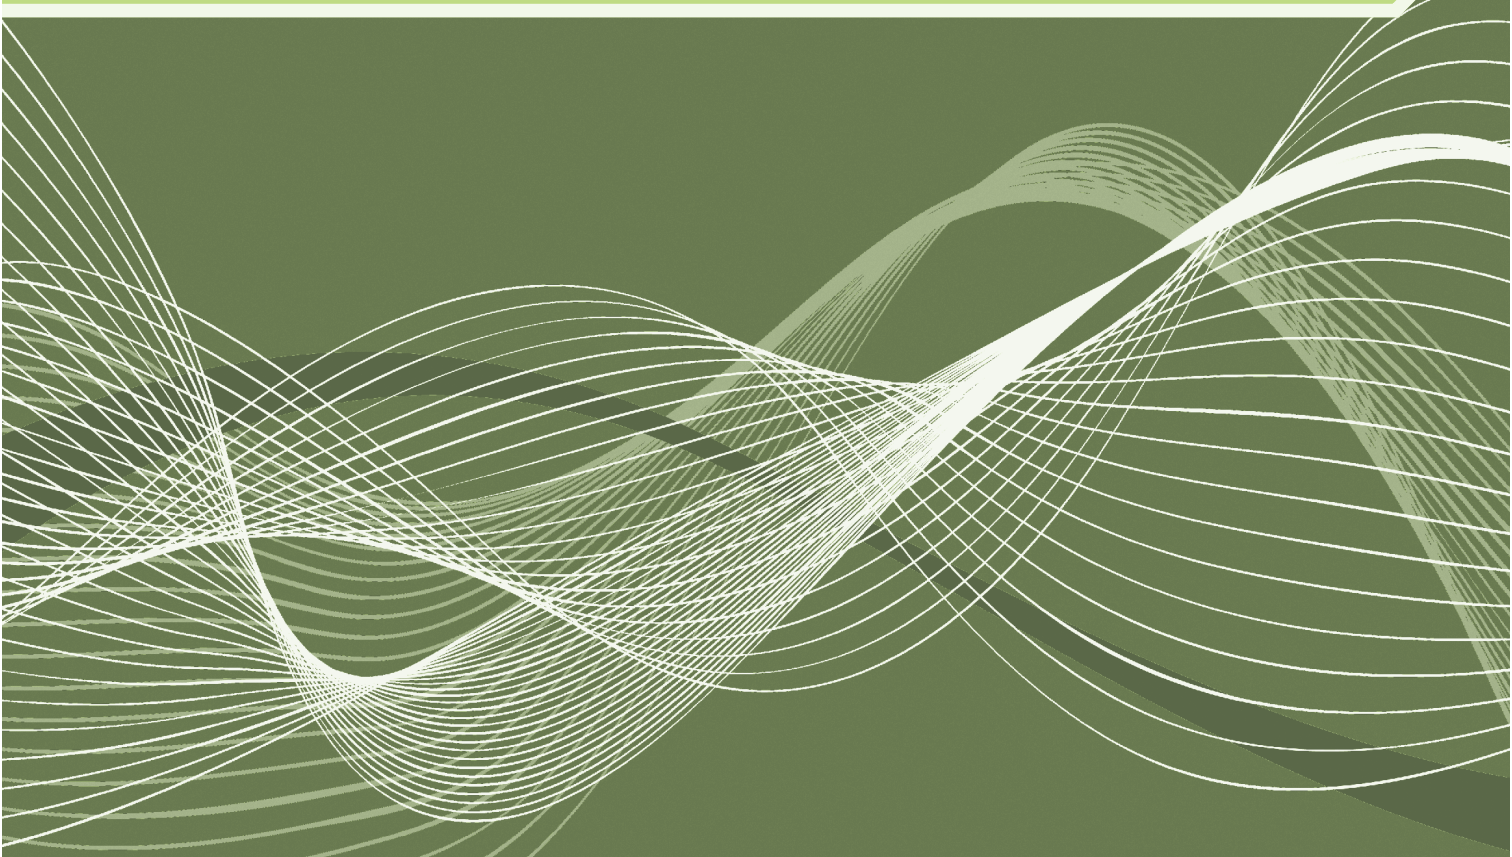

SEPTEMBER 2015

IEEE Signal Processing Society

IEEE

SignalProcessing Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page Qmags THE WORLD'S NEWSSTAND®

IEEE International Symposium on Biomedical Imaging

April 13-16, 2016, Prague, Czech Republic

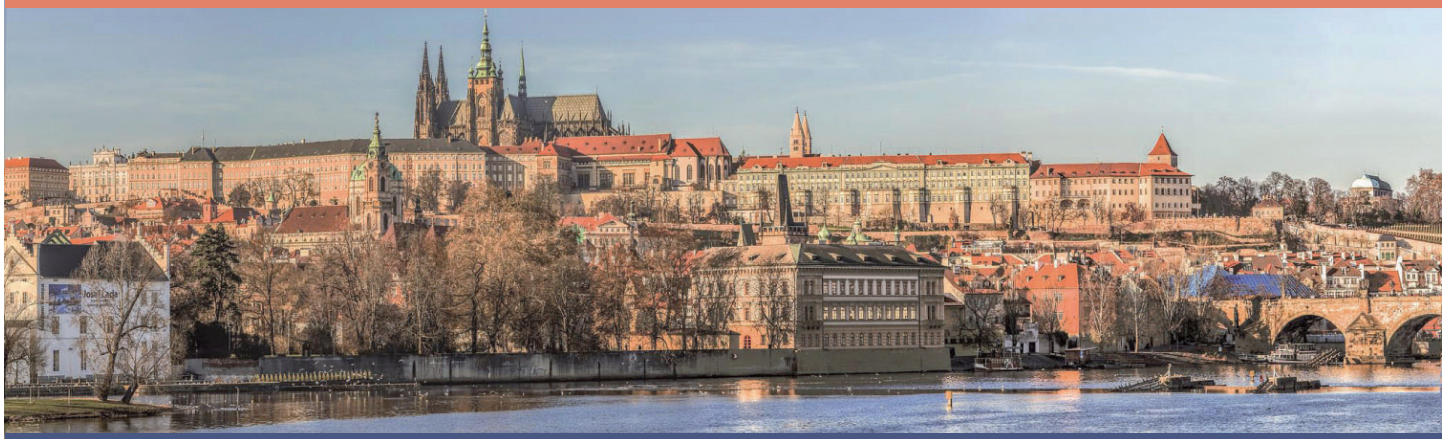

CALL FOR PAPERS

The IEEE International Symposium on Biomedical Imaging (ISBI) is a scientific conference dedicated to mathematical, algorithmic, and computational aspects of biomedical imaging, across all scales of observation. It fosters knowledge transfer among different imaging communities and contributes to an integrative approach to biomedical imaging.

ISBI is a joint initiative from the IEEE Signal Processing Society (SPS) and the IEEE Engineering in Medicine and Biology Society (EMBS). The 2016 meeting will include tutorials, and a scientific program composed of plenary talks, invited special sessions, challenges, as well as oral and poster presentations of peer-reviewed papers.

High-quality papers are requested containing original contributions to the topics of interest including image formation and reconstruction, computational and statistical image processing and analysis, dynamic imaging, visualization, image quality assessment, and physical, biological, and statistical modeling. Accepted 4-page regular papers will be published in the symposium proceedings published by IEEE and included in IEEE Xplore.

To encourage attendance by a broader audience of imaging scientists and offer additional presentation opportunities, ISBI 2016 will continue to propose a second track featuring posters selected from 1-page abstract submissions without subsequent archival publication.

Venue:

ISBI 2016 will be held in the 4-star Clarion Congress Hotel Prague, one of the most modern congress hotels in Prague, Czech Republic. The hotel has space for up to 2500 delegates and a corresponding accommodation capacity, including a wellness and fitness center. It takes less than 15 minutes by Underground to reach the historical center of Prague, a UNESCO Heritage Site. Do not miss Prague's world-famous Old Town Square, Charles Bridge, and Prague Castle.

IEEE | IEEE Signal Processing Society | EMB

Important Dates:

4-page paper submission
August 3rd - October 26th, 2015

Author Notification for 4-page papers
December 23rd, 2015

Final version of 4-page papers & registration
January 11th, 2016

Conference Chairs

Jan Kybic
Czech Technical University in Prague

Milan Sonka
The University of Iowa

Program Chairs

Karl Rohr
University of Heidelberg

Boudewijn Lelieveldt
Leiden University Medical Center, Netherlands

Organizing Committee

Joe Reinhardt, Bram van Ginneken, Franjo Pernus, Punam Saha, Mathews Jacob, Arrate Munoz-Barrutia, Zoltan Szabo, Radim Krupicka, Michal Kozubek, Ipek Oguz, Tomaz Vrtovec, Jiri Jan, Jiri Janacek, Lucie Kubinova, Pavel Tomancak, Eduardo Romero, Juan David Garcia, Jan Petr, Michal Sofka

Contact

Janice Sandler j.sandler@ieee.org

http://biomedicalimaging.org/2016

IEEE TRANSACTIONS ON SIGNAL PROCESSING

A PUBLICATION OF THE IEEE SIGNAL PROCESSING SOCIETY

IEEE Signal Processing Society

www.signalprocessingsociety.org

Indexed in PubMed® and MEDLINE®, products of the United States National Library of Medicine

PubMed MEDLINE U.S. National Library of Medicine

AUGUST 1, 2015 VOLUME 63 NUMBER 15 ITPRED (ISSN 1053-587X)

REGULAR PAPERS

IEEE TRANSACTIONS ON SIGNAL PROCESSING (ISSN 1053-587X) is published semimonthly by the Institute of Electrical and Electronics Engineers, Inc. Responsibility for the contents rests upon the authors and not upon the IEEE, the Society/Council, or its members. **IEEE Corporate Office:** 3 Park Avenue, 17th Floor, New York, NY 10016-5997. **IEEE Operations Center:** 445 Hoes Lane, Piscataway, NJ 08854-4141. **NJ Telephone:** +1 732 981 0060. **Price/Publication Information**: Individual copies: IEEE Members $20.00 (first copy only), nonmembers $602.50 per copy. (Note: Postage and handling charge not included.) Member and nonmember subscription prices available upon request. **Copyright and Reprint Permissions:** Abstracting is permitted with credit to the source. Libraries are permitted to photocopy for private use of patrons, provided the per-copy fee of $31.00 is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923. For all other copying, reprint, or republication permission, write to Copyrights and Permissions Department, IEEE Publications Administration, 445 Hoes Lane, Piscataway, NJ 08854-4141.

Postmaster: Send address changes to IEEE TRANSACTIONS ON SIGNAL PROCESSING, IEEE, 445 Hoes Lane, Piscataway, NJ 08854-4141. GST Registration No. 125634188. CPC Sales Agreement #40013087. Return undeliverable Canada addresses to: Pitney Bowes IMEX, P.O. Box 4332, Stanton Rd., Toronto, ON M5W 3J4, Canada. IEEE prohibits discrimination, harassment and bullying. For more information visit http://www.ieee.org/nondiscrimination. Printed in U.S.A.

IEEE

IEEE TRANSACTIONS ON SIGNAL PROCESSING

A PUBLICATION OF THE IEEE SIGNAL PROCESSING SOCIETY

IEEE Signal Processing Society

www.signalprocessingsociety.org

Indexed in PubMed® and MEDLINE®, products of the United States National Library of Medicine

PubMed MEDLINE U.S. National Library of Medicine

AUGUST 15, 2015 VOLUME 63 NUMBER 16 ITPRED (ISSN 1053-587X)

REGULAR PAPERS

IEEE TRANSACTIONS ON SIGNAL PROCESSING (ISSN 1053-587X) is published semimonthly by the Institute of Electrical and Electronics Engineers, Inc. Responsibility for the contents rests upon the authors and not upon the IEEE, the Society/Council, or its members. **IEEE Corporate Office:** 3 Park Avenue, 17th Floor, New York, NY 10016-5997. **IEEE Operations Center:** 445 Hoes Lane, Piscataway, NJ 08854-4141. **NJ Telephone:** +1 732 981 0060. **Price/Publication Information**: Individual copies: IEEE Members $20.00 (first copy only), nonmembers $602.50 per copy. (Note: Postage and handling charge not included.) Member and nonmember subscription prices available upon request. **Copyright and Reprint Permissions:** Abstracting is permitted with credit to the source. Libraries are permitted to photocopy for private use of patrons, provided the per-copy fee of $31.00 is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923. For all other copying, reprint, or republication permission, write to Copyrights and Permissions Department, IEEE Publications Administration, 445 Hoes Lane, Piscataway, NJ 08854-4141. **Postmaster:** Send address changes to IEEE TRANSACTIONS ON SIGNAL PROCESSING, IEEE, 445 Hoes Lane, Piscataway, NJ 08854-4141. GST Registration No. 125634188. CPC Sales Agreement #40013087. Return undeliverable Canada addresses to: Pitney Bowes IMEX, P.O. Box 4332, Stanton Rd., Toronto, ON M5W 3J4, Canada. IEEE prohibits discrimination, harassment and bullying. For more information visit http://www.ieee.org/nondiscrimination. Printed in U.S.A.

IEEE

new

IEEE TRANSACTIONS ON

SIGNAL AND INFORMATION PROCESSING OVER NETWORKS

Now accepting paper submissions

The new *IEEE Transactions on Signal and Information Processing over Networks* publishes high-quality papers that extend the classical notions of processing of signals defined over vector spaces (e.g. time and space) to processing of signals and information (data) defined over networks, potentially dynamically varying. In signal processing over networks, the topology of the network may define structural relationships in the data, or may constrain processing of the data. Topics of interest include, but are not limited to the following:

Adaptation, Detection, Estimation, and Learning
- Distributed detection and estimation
- Distributed adaptation over networks
- Distributed learning over networks
- Distributed target tracking
- Bayesian learning; Bayesian signal processing
- Sequential learning over networks
- Decision making over networks
- Distributed dictionary learning
- Distributed game theoretic strategies
- Distributed information processing
- Graphical and kernel methods
- Consensus over network systems
- Optimization over network systems

Communications, Networking, and Sensing
- Distributed monitoring and sensing
- Signal processing for distributed communications and networking
- Signal processing for cooperative networking
- Signal processing for network security
- Optimal network signal processing and resource allocation

Modeling and Analysis
- Performance and bounds of methods
- Robustness and vulnerability
- Network modeling and identification

Modeling and Analysis (cont.)
- Simulations of networked information processing systems
- Social learning
- Bio-inspired network signal processing
- Epidemics and diffusion in populations

Imaging and Media Applications
- Image and video processing over networks
- Media cloud computing and communication
- Multimedia streaming and transport
- Social media computing and networking
- Signal processing for cyber-physical systems
- Wireless/mobile multimedia

Data Analysis
- Processing, analysis, and visualization of big data
- Signal and information processing for crowd computing
- Signal and information processing for the Internet of Things
- Emergence of behavior

Emerging topics and applications
- Emerging topics
- Applications in life sciences, ecology, energy, social networks, economic networks, finance, social sciences, smart grids, wireless health, robotics, transportation, and other areas of science and engineering

Editor-in-Chief: Petar M. Djurić, Stony Brook University (USA)

To submit a paper, go to: https://mc.manuscriptcentral.com/tsipn-ieee

IEEE COMMUNICATIONS SOCIETY

IEEE Signal Processing Society

IEEE COMPUTER SOCIETY

NEW PUBLICATION:

Transactions on Signal and Information Processing over Networks (T-SIPN)*

http://www.signalprocessingsociety.org/publications/periodicals/tsipn/

>>We are accepting paper submissions: please submit a manuscript here<<

There has been an explosion of research in network systems of various types, including physical, engineered, biological and social systems. Its aim is to find answers to fundamental questions about the systems and with them be able to understand, predict, and control them better. To that end, a core area of work is signal and information processing over networks.

Network systems represent a growing research field encompassing numerous disciplines in science and engineering. Their complexity is reflected in the diversity and the interconnectivity of their elements, which have the capacity to adapt and learn from experience. Applications of network systems are wide and include communications (wireless sensor networks, peer-to-peer networks, pervasive mobile networks, the Internet of Things), the electric power grid, biology, the Internet, the stock market, ecology, and in animal and human societies.

The Transactions on Signal and Information Processing over Networks (T-SIPN) publishes timely peer-reviewed technical articles on advances in the theory, methods, and algorithms for signal and information processing, inference, and learning in network systems. The following core topics define the scope of the Transaction:

Adaptation, Detection, Estimation, and Learning (ADEL)

- Distributed detection and estimation (ADEL-DDE)
- Distributed adaptation over networks (ADEL-DAN)
- Distributed learning over networks (ADEL-DLN)
- Distributed target tracking (ADEL-DTT)
- Bayesian learning; Bayesian signal processing (ADEL-BLSP)
- Sequential learning over networks (ADEL-SLN)
- Decision making over networks (ADEL-DMN)
- Distributed dictionary learning (ADEL-DDL)
- Distributed game theoretic strategies (ADEL-DGTS)
- Distributed information processing (ADEL-DIP)
- Graphical and kernel methods (ADEL-GKM)
- Consensus over network systems (ADEL-CNS)
- Optimization over network systems (ADEL-ONS)

Communications, Networking, and Sensing (CNS)

- Distributed monitoring and sensing (CNS-DMS)
- Signal processing for distributed communications and networking (CNS-SPDCN)
- Signal processing for cooperative networking (CNS-SPCN)
- Signal processing for network security (CNS-SPNS)
- Optimal network signal processing and resource allocation (CNS-NSPRA)

(continued on next page)

Modeling and Analysis (MA)

- Performance and bounds of methods (MA-PBM)
- Robustness and vulnerability (MA-RV)
- Network modeling and identification (MA-NMI)
- Simulations of networked information processing systems (MA-SNIPS)
- Social learning (MA-SL)
- Bio-inspired network signal processing (MA-BNSP)
- Epidemics and diffusion in populations (MA-EDP)

Imaging and Media Applications (IMA)

- Image and video processing over networks (IMA-IVPN)
- Media cloud computing and communication (IMA-MCCC)
- Multimedia streaming and transport (IMA-MST)
- Social media computing and networking (IMA-SMCN)
- Signal processing for cyber-physical systems (IMA-SPCPS)
- Wireless/mobile multimedia (IMA-WMM)

Data Analysis (DA)

- Processing, analysis, and visualization of big data (DA-BD)
- Signal and information processing for crowd computing (DA-CC)
- Signal and information processing for the Internet of Things (DA-IOT)
- Emergence of behavior (DA-EB)

Emerging topics and applications (ETA)

- Emerging topics (ETA-ET)
- Applications in life sciences, ecology, energy, social networks, economic networks, finance, social sciences etc. smart grids, wireless health, robotics, transportation, and other areas of science and engineering (ETA-APP)

>>We are accepting paper submissions: please submit a manuscript here<<

***T-SIPN is co-sponsored by the Signal Processing, Communications and Computer societies**

IEEE/ACM TRANSACTIONS ON

AUDIO, SPEECH, AND LANGUAGE PROCESSING

A PUBLICATION OF THE IEEE SIGNAL PROCESSING SOCIETY

IEEE Signal Processing Society | acm Association for Computing Machinery | www.signalprocessingsociety.org

Indexed in PubMed® and MEDLINE®, products of the United States National Library of Medicine

PubMed | MEDLINE U.S. National Library of Medicine

AUGUST 2015 | VOLUME 23 | NUMBER 8 | ITASFA | (ISSN 2329-9290)

REGULAR PAPERS

IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE PROCESSING (ISSN 2329-9290) is published bimonthly in print and monthly online by the Institute of Electrical and Electronics Engineers, Inc. Responsibility for the contents rests upon the authors and not upon the IEEE, the Society/Council, or its members. **IEEE Corporate Office:** 3 Park Avenue, 17th Floor, New York, NY 10016-5997. **IEEE Operations Center:** 445 Hoes Lane, Piscataway, NJ 08854-4141. **NJ Telephone:** +1 732 981 0060. **Price/Publication Information:** Individual copies: IEEE Members $20.00 (first copy only), nonmembers $339.00 per copy. (Note: Postage and handling charge not included.) Member and nonmember subscription prices available upon request. **Postmaster:** Send address changes to IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE PROCESSING, IEEE, 445 Hoes Lane, Piscataway, NJ 08854-4141. GST Registration No. 125634188. CPC Sales Agreement #40013087. Return undeliverable Canada addresses to: Pitney Bowes IMEX, P.O. Box 4332, Stanton Rd., Toronto, ON M5W 3J4, Canada. IEEE prohibits discrimination, harassment and bullying. For more information visit http://www.ieee.org/nondiscrimination. Printed in U.S.A.

IEEE

The Ninth IEEE Sensor Array and Multichannel Signal Processing Workshop

10th-13th July 2016,Rio de Janeiro, Brazil

IEEE Signal Processing Society

Call for Papers

General Chairs
Rodrigo C. de Lamare,
PUC-Rio, Brazil and University of York, United Kingdom

Martin Haardt,
TU Ilmenau, Germany

Technical Chairs
Aleksandar Dogandzic,
Iowa State University, USA

Vítor Nascimento,
University of São Paulo, Brazil

Special Sessions Chair
Cédric Richard,
University of Nice, France

Publicity Chair
Maria Sabrina Greco,
Universityof Pisa, Italy

Important Dates

Special Session Proposals
29th January , 2016

Submission of Papers
26th February, 2016

Notification of Acceptance
29th April , 2016

Final Manuscript Submission
16th May, 2016

Advance Registration
16th May, 2016

Technical Program
The SAM Workshop is an important IEEE Signal Processing Society event dedicated to sensor array and multichannel signal processing. The organizing committee invites the international community to contribute with state-of-the-art developments in the field. SAM 2016 will feature plenary talks by leading researchers in the field as well as poster and oral sessions with presentations by the participants.

Welcome to Rio de Janeiro! – The workshop will be held at the Pontifical Catholic University of Rio de Janeiro, located in Gávea, in a superb area surrounded by beaches, mountains and the Tijuca National Forest, the world's largest urban forest. Rio de Janeiro is a world renowned city for its culture, beautiful landscapes, numerous tourist attractions and international cuisine. The workshop will take place during the first half of July about a month before the 2016 Summer Olympic Games when Rio will offer plenty of cultural activities and festivities, which will make SAM 2016 a memorable experience.

Research Areas
Authors are invited to submit contributions in the following areas:
- Adaptive beamforming
- Array processing for biomedical applications
- Array processing for communications
- Blind source separation and channel identification
- Computational and optimization techniques
- Compressive sensing and sparsity-based signal processing
- Detection and estimation
- Direction-of-arrival estimation
- Distributed and adaptive signal processing
- Intelligentsystems and knowledge-based signal processing
- Microphone and loudspeaker array applications
- MIMO radar
- Multi-antenna systems: multiuser MMO, massive MIMO and space-time coding
- Multi-channel imaging and hyperspectral processing
- Multi-sensor processing for smart grid and energy
- Non-Gaussian, nonlinear, and non-stationary models
- Performance evaluations with experimental data
- Radar and sonar array processing
- Sensor networks
- Source Localization, Classification and Tracking
- Synthetic aperture techniques
- Space-time adaptive processing
- Statistical modelling for sensor arrays
- Waveform diverse sensors and systems

Submission of papers – Full-length four-page papers will be accepted only electronically.

Special session proposals – They should be submitted by e-mail to the Technical Program Chairs and the Special Sessions Chair and include a topical title, rationale, session outline, contact information, and list of invited speakers.

IEEE TRANSACTIONS ON

IMAGE PROCESSING

A PUBLICATION OF THE IEEE SIGNAL PROCESSING SOCIETY

IEEE Signal Processing Society

www.signalprocessingsociety.org

Indexed in PubMed® and MEDLINE®, products of the United States National Library of Medicine

PubMed MEDLINE U.S. National Library of Medicine

AUGUST 2015 VOLUME 24 NUMBER 8 IIPRE4 (ISSN 1057-7149)

PAPERS

IEEE

IEEE TRANSACTIONS ON

IMAGE PROCESSING

A PUBLICATION OF THE IEEE SIGNAL PROCESSING SOCIETY

IEEE Signal Processing Society

www.signalprocessingsociety.org

Indexed in PubMed® and MEDLINE®, products of the United States National Library of Medicine

PubMed MEDLINE U.S. National Library of Medicine

SEPTEMBER 2015 VOLUME 24 NUMBER 9 IIPRE4 (ISSN 1057-7149)

PAPERS

IEEE

NEW!

IEEE TRANSACTIONS ON COMPUTATIONAL IMAGING

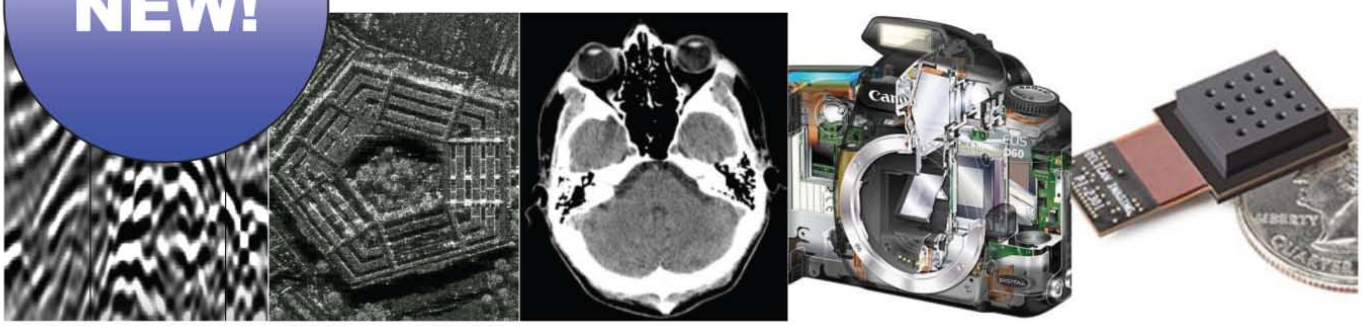

The new IEEE Transactions on Computational Imaging seeks original manuscripts for publication. This new journal will publish research results where computation plays an integral role in the image formation process. All areas of computational imaging are appropriate, ranging from the principles and theory of computational imaging, to modeling paradigms for computational imaging, to image formation methods, to the latest innovative computational imaging system designs. Topics of interest include, but are not limited to the following:

Imaging Models and Representation

- Statistical-model based methods
- System and image prior models
- Noise models
- Graphical and tree-based models
- Perceptual models

Computational Sensing

- Coded source methods
- Structured light
- Coded aperture methods
- Compressed sensing
- Light-field sensing
- Plenoptic imaging
- Hardware and software systems

Computational Image Creation

- Sparsity-based methods
- Statistically-based inversion methods, Bayesian regularization
- Super-resolution, multi-image fusion
- Learning-based methods, Dictionary-based methods
- Optimization-based methods; proximal iterative methods, ADMM

Computational Photography

- Non-classical image capture, Generalized illumination
- Time-of-flight imaging
- High dynamic range imaging
- Focal stacks

Computational Consumer Imaging

- Cell phone imaging
- Camera-array systems
- Depth cameras

Computational Acoustic Imaging

- Multi-static ultrasound imaging
- Photo-acoustic imaging
- Acoustic tomography

Computational Microscopic Imaging

- Holographic microscopy
- Quantitative phase imaging
- Multi-illumination microscopy
- Lensless microscopy

Tomographic Imaging

- X-ray CT
- PET
- SPECT

Magnetic Resonance Imaging

- Diffusion tensor imaging
- Fast acquisition

Radar Imaging

- Synthetic aperture imaging
- Inverse synthetic imaging
- Terahertz imaging

Geophysical Imaging

- Multi-spectral imaging
- Ground penetrating radar
- Seismic tomography

Multi-spectral Imaging

- Multi-spectral imaging
- Hyper-spectral imaging
- Spectroscopic imaging

Editor-in-Chief: W. Clem Karl, Boston University.
To submit a paper go to: https://mc.manuscriptcentral.com/tci-ieee

IEEE Signal Processing Society | EMB | GRSS | IEEE Consumer Electronics Society

The 23rd IEEE International Conference on Image Processing (ICIP) will be held in the Phoenix Convention Centre, Phoenix, Arizona, USA, on September 25 - 28, 2016. ICIP is the world's largest and most comprehensive technical conference focused on image and video processing and computer vision. In addition to the Technical Program, ICIP 2016 will feature **an Innovation Program focused on innovative vision technologies and fostering innovation, entrepreneurship, and networking**. The conference will feature world-class speakers, tutorials, exhibits, and a vision technology showcase.

Topics in the ICIP 2016 Technical Program include but are not limited to the following:

Filtering, Transforms, Multi-Resolution Processing
Restoration, Enhancement, Super-Resolution
Computer Vision Algorithms and Technologies
Compression, Transmission, Storage, Retrieval
Computational Imaging
Color and Multispectral Processing
Multi-View and Stereoscopic Processing
Multi-Temporal and Spatio-Temporal Processing
Video Processing and Analytics
Authentication and Biometrics
Biological and Perceptual-based Processing
Visual Quality Assessment
Scanning, Display, and Printing
Document and Synthetic Visual Processing
Applications to various fields (e.g., biomedical, Advanced Driving Assist Systems, assistive living, security, learning, health and environmental monitoring, manufacturing, consumer electronics)

The ICIP 2016 innovation program will feature a **vision technology showcase** of state-of-the-art vision technologies, innovation challenges, talks by innovation leaders and entrepreneurs, tutorials, and networking.

Paper Submission: Prospective authors are invited to submit full-length papers at the conference website, with up to four pages for technical content including figures and references, and with one additional optional 5th page for references only. Submission instructions, templates for the required paper format, and information on "no show" policy are available at www.icip2016.com.

Tutorials and Special Sessions Proposals: Tutorials will be held on September 25, 2016. Tutorial proposals should be submitted to *tutorials@icip2016.com* and must include title, outline, contact information, biography and selected publications for the presenter(s), and a description of the tutorial and material to be distributed to participants. Special Sessions proposals should be submitted to *specialsessions@icip2016.com* and must include a topical title, rationale, session outline, contact information, and a list of invited papers. For detailed submission guidelines, please refer the ICIP 2016 website at www.icip2016.com.

Important Deadlines:
Special Session and Tutorial Proposals: November 16, 2015
Notification of Special Session and Tutorial Acceptance: December 18, 2015

Paper Submissions: January 25, 2016
Notification of Paper Acceptance: April 30, 2016

Visual Technology Innovator Award Nomination: March 30, 2016

Revised Paper Upload Deadline: May 30, 2016
Authors' Registration Deadline: May 30, 2016

http://www.facebook.com/icip2016
https://twitter.com/icip2016/
https://www.linkedin.com/groups/ICIP-2016-6940658

IEEE TRANSACTIONS ON
INFORMATION FORENSICS AND SECURITY

A PUBLICATION OF THE IEEE SIGNAL PROCESSING SOCIETY

IEEE Signal Processing Society

www.signalprocessingsociety.org

AUGUST 2015 VOLUME 10 NUMBER 8 ITIFA6 (ISSN 1556-6013)

PAPERS

ANNOUNCEMENTS

IEEE TRANSACTIONS ON MULTIMEDIA

A PUBLICATION OF
THE IEEE SIGNAL PROCESSING SOCIETY
THE IEEE CIRCUITS AND SYSTEMS SOCIETY
THE IEEE COMMUNICATIONS SOCIETY

IEEE Signal Processing Society | CAS | IEEE COMMUNICATIONS SOCIETY

http://www.signalprocessingsociety.org/tmm/

TECHNICALLY COSPONSORED BY THE IEEE COMPUTER SOCIETY

IEEE computer society

AUGUST 2015 **VOLUME 17** **NUMBER 8** **ITMUF8** **(ISSN 1520-9210)**

PAPERS

IEEE

IEEE JOURNAL OF
SELECTED TOPICS IN SIGNAL PROCESSING

IEEE Signal Processing Society

www.ieee.org/sp/index.html

AUGUST 2015 VOLUME 9 NUMBER 5 IJSTGY (ISSN 1932-4553)

ISSUE ON SPATIAL AUDIO

IEEE

Call for Papers
IEEE Signal Processing Society
IEEE Journal of Selected Topics in Signal Processing

Special Issue on Financial Signal Processing and Machine Learning for Electronic Trading

The financial sector has been historically served by experts in finance, quantitative finance, risk management, and electronic trading. However, it still presents a very rich and diverse application area for signal processing methods and technologies. These techniques span across multiple specialties of signal processing field and change dramatically depending on the trading frequency, ranging from covariance modeling to short-term prediction based on market microstructure. Moreover, high performance computing and DSP technologies (FPGA, GPU, others) have already transformed the financial industry by facilitating the implementation of computationally demanding analysis and modeling of high frequency market data and others in real-time. The currently available computational power has brought in once hard to implement machine learning tools for the use of financial applications. The term Big Data Finance is already coined in the field.

This special issue aims to compile relevant research contributions from the disciplines of finance, mathematics, data science and engineering to facilitate scientific cross-fertilization. It will also serve the signal processing community to be exposed to the state of the art in mathematical finance, financial engineering, financial signal processing and electronic trading, and to foster future research in this emerging area.

Topics of interest include but are not limited to

- Big Data Finance
- High Performance DSP (FPGA, GPU, others) for Finance and Electronic Trading
- Machine Learning Methods for Financial Applications and Trading
- Signal Processing Algorithms for Electronic Trading
- Multiresolution Techniques for Multi-frequency Investment and Trading
- High Frequency Trading (HFT) and Order Routing
- Market Microstructure Modeling (price behavior and discovery, limit order book, etc.)
- Theory of Games and Auctions in Financial Models
- Financial News and Social Media Analysis for Intelligent Portfolio Management
- Portfolio Optimization and Management
- Risk Analysis and Models (risk and correlation measures, estimation techniques)

Submission Procedure:

Information for prospective authors can be found at: http://www.signalprocessingsociety.org/publications/periodicals/jstsp/. Manuscripts should be submitted through Manuscript Central system at http://mc.manuscriptcentral.com/jstsp-ieee. Manuscripts will be peer reviewed according to the IEEE standards.

Manuscript submission due:	Oct 1, 2015	First review completed:	Dec 15, 2015
Revised manuscript due:	Jan 31, 2016	Second review completed:	March 15, 2016
Final manuscript due:	May 1, 2016		

Prospective authors may contact Prof. Ali N. Akansu at Akansu@NJIT.edu with inquiries.

Guest editors:

Ali N. Akansu, NJIT, USA
Dmitry Malioutov, IBM Research, USA
Daniel P. Palomar, HKUST, Hong Kong
Emmanuelle Jay, QAMLab, France
Danilo P. Mandic, Imperial College London, UK

Call for Papers - IEEE Journal on Selected Topics in Signal Processing
Special Issue on Person-Centered Signal Processing for Assistive, Rehabilitative and Wearable Health Technologies

Human-centered computing (HCC) has emerged as a major interdisciplinary subfield of engineering that puts the human at the center of research activities and places emphasis on understanding human behavior, needs, adaptation, and societal and cultural differences to design better technologies. Person-centered computing and signal processing allows HCC to focus on an *individual* user's needs and behaviors while maintaining broad applicability to the wider population through built-in flexibility and the process of co-adaptation. Co-adaptation is the bidirectional process of a human and machine both learning and adapting over time through continual use and experience. The onus of adaptation in a person-centered design lies more with the system and the modus of interaction. The complexity is mainly due to human behavior being multimodal and complex, motivated by needs that are individualized, always changing, and often implicit. Multimodal sensing is commonly targeted at the visual and auditory channels, but there are many other complementary modalities including movement, touch, vital signs, physiological response, and brain-computer interfaces. At the core of every person-centered computing system is a network of sensors. This paradigm has created a need for research to develop and validate models for person-centered systems based on intelligent, reliable, robust and adaptive sensor networks. We invite authors to submit articles representing the cutting edge in signal processing topics including (but not limited to) those listed below. Topics should be approached from a person-centered perspective, considering individualized yet generalizable designs and co-adaptation.

Applications - assistive technology: Computer vision for navigation aids, shopping assistants, social interaction assistants, tactile-vision substitution systems, and general accessibility for individuals who are blind; Audio and acoustic signal processing for speech synthesis and sensory substitution (e.g., tactile-audio) to assist individuals with disabilities in communication and computer access; Signal processing and robust classification techniques for brain-computer interfaces to assist individuals with disabilities in communication and computer access.
Applications – rehabilitation: Signal processing, feature extraction and pattern recognition techniques toward understanding and analyzing motion data from position/inertial body worn sensors, computer vision and depth information to support physical rehabilitation and therapeutic exercise.
Applications – wearable health: Signal processing, machine learning, predictive modeling and gesture/activity recognition for wearable health technology devices including physiological sensors, health monitors, and vital signs trackers; and Gait signal processing, machine learning and activity recognition for gait monitoring including step detection, stride length estimation and event detection (e.g., shuffling, freezing of gait, falls).
Models: Learning and inference tools and models adapted to person-centered signal processing and computing including alternative classification techniques; Signal processing and data fusion methods for multimodal sensor analytics; and Signal processing methods for Wireless Body Area Sensor Networks communication and data fusion.

Prospective authors should visit the IEEE signal processing website for information on paper submission. Manuscripts should be submitted at http://mc.manuscriptcentral.com/jstsp-ieee.

Important Dates

- Manuscript submission due: September 1, 2015
- First review completed: November 15, 2015
- Revised manuscript due: December 31, 2015
- Second review completed: February 15, 2016
- Final manuscript due: April 1, 2016
- Publication date: August 2016

Guest Editors

- Sethuraman Panchanathan, Arizona State University, USA, panch@asu.edu
- Diane Cook, Washington State University, USA, cook@eecs.wsu.edu
- Noel O'Connor, Dublin City University, Ireland, noeloconnor@dcu.edu.ie
- Mrinal Mandal, University of Alberta, Canada, mmandal@ualberta.ca
- Mohan Kankanhalli, National University of Singapore, Singapore, mohan@comp.nus.edu.sg

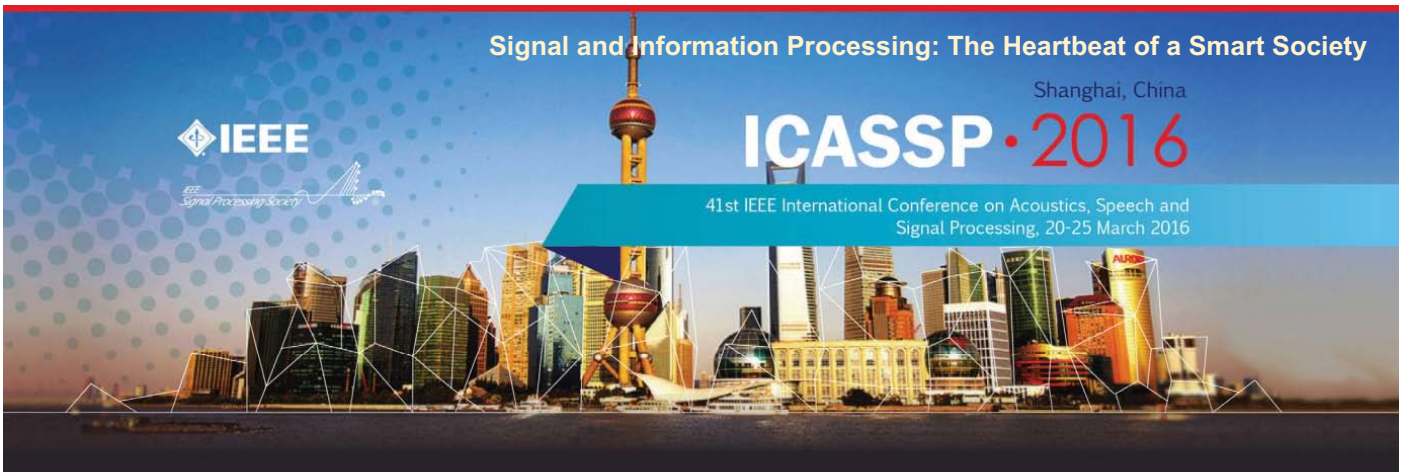

General Chairs
Zhi Ding, Univ. of California, Davis, USA
Zhi-Quan Luo, Univ. of Minnesota, USA
Wenjun Zhang, Shanghai Jiao Tong Univ., China

Technical Program Chairs
P. C. Ching, Chinese Univ. of Hong Kong, Hong Kong
Dominic K.C. Ho, Univ. of Missouri, USA

Finance Chairs
Shuguang Cui, Texas A&M Univ., USA
Rong Xie, Shanghai Jiao Tong Univ., China

Plenaries Chairs
Zhi-Pei Liang, UIUC, USA
Björn Ottersten, Univ. of Luxembourg, Luxembourg

Special Sessions Chairs
Tim Davidson, McMaster Univ., Canada
Jianguo Huang, Northwestern Polytech. Univ., China

Tutorials Chairs
Jian Li, Univ. of Florida, USA
Jose Principe, Univ. of Florida, USA

Student Session Chair
Wei Zhang, Univ. of New South Wales, Australia

Registration Chairs
Tongtong Li, Michigan State Univ., USA
Xiaojun Yuan, ShanghaiTech Univ., China

Publicity Chairs
Xiaokang Yang, Shanghai Jiao Tong Univ., China
Mounir Ghogho, Leeds Univ., UK
Ignacio Santamaria, Univ. of Cantabria, Spain

Publication Chairs
Min Dong, Univ. of Ontario Inst. of Tech., Canada
Thomas Fang Zheng, Tsinghua Univ., China

Industrial & Exhibit Chairs
Li Deng, Microsoft, USA
Jinyu Li, Microsoft, USA
Cathy Wicks, Texas Instruments, USA

Local Arrangement Chairs
Ning Liu, Shanghai Jiao Tong Univ., China
Meixia Tao, Shanghai Jiao Tong Univ., China

Webmaster
Yi Xu, Shanghai Jiao Tong Univ., China

Workshop Chairs
Jianguo Huang, Northwestern Polytech. Univ., China
Jiwu Huang, Sun Yat-sen Univ., China

ICASSP2016: Signal and information processing is the driving heartbeat in the development of technologies that enrich our lives and advance our society. The 41st International Conference on Acoustics, Speech, and Signal Processing (ICASSP) will be held in the Shanghai International Convention Center, Shanghai, China between March 20 and 25, 2016. The conference provides, both for researchers and developers, an engaging forum to exchange ideas and propel new developments in this field. The 2016 conference will showcase world-class presentations by internationally renowned speakers and will facilitate a fantastic opportunity to network with like-minded professionals from around the world. Topics include but are not limited to:

Audio and acoustic signal processing
Bio-imaging and biomedical signal processing
Signal processing education
Speech processing
Industry technology tracks
Information forensics and security
Machine learning for signal processing
Signal processing for Big Data
Multimedia signal processing
Sensor array & multichannel signal processing
Design & implementation of signal processing systems
Signal processing for communications & networking
Image, video & multidimensional signal processing
Signal processing theory & methods
Spoken language processing
Signal processing for the Internet of Things

Shanghai: Shanghai is the most populous city in China and one of the most populous cities in the world. A global city, Shanghai exerts influence over global commerce, finance, culture, art, fashion, research and entertainment. The city is located in the middle portion of the Chinese coast, and sits at the mouth of the Yangtze River. The city is a tourist destination renowned for its historical landmarks, such as the Bund and City God Temple, and its modern and ever-expanding Pudong skyline including the Oriental Pearl Tower. Today, Shanghai is the largest center of commerce and finance in mainland China, and has been described as the "showpiece" of the world's fastest-growing major economy.

Submission of Papers: Prospective authors are invited to submit full-length papers, with up to four pages for technical content including figures and possible references, and with one additional optional 5th page containing only references. A selection of best student papers will be made by the ICASSP 2016 committee upon recommendations from the Technical Committees.

Tutorial and Special Session Proposals: Tutorials will be held on March 20 and 21, 2016. Tutorial proposals must include title, outline, contact information, biography and selected publications for the presenter(s), and a description of the tutorial and the material to be distributed to participants. Special session proposals must include a topical title, rationale, session outline, contact information, and a list of invited speakers. Additional information can be found at the ICASSP 2016 website.

Signal Processing Letters: Authors of IEEE Signal Processing Letters (SPL) papers will be given the opportunity to present their work at ICASSP 2016, subject to space availability and approval by the ICASSP Technical Program Chairs. SPL papers published between January 1, 2015 and December 31, 2015 are eligible for presentation at ICASSP 2016.

Show and Tell: S&T offers a perfect stage to showcase innovative ideas in all technical areas of interest at ICASSP. S&T sessions contain demos that are highly interactive and visible. Please refer to the ICASSP 2016 website for additional information regarding demo submission.

Important Deadlines:

Special session & tutorial proposals	August 3, 2015
Notification of special session & tutorial acceptance	September 11, 2015
Submission of regular papers	September 25, 2015
Signal processing letters	December 16, 2015
Notification of paper acceptance	December 21, 2015
Revised paper upload	January 22, 2016
Author registration	January 22, 2016

IEEE

SIGNAL PROCESSING LETTERS

A PUBLICATION OF THE IEEE SIGNAL PROCESSING SOCIETY

IEEE Signal Processing Society

www.ieee.org/sp/index.html

OCTOBER 2015 VOLUME 22 NUMBER 10 ISPLEM (ISSN 1070-9908)

LETTERS

IEEE SIGNAL PROCESSING LETTERS (ISSN 1070–9908) is published quarterly in print and monthly online by the Institute of Electrical and Electronics Engineers, Inc. Responsibility for the contents rests upon the authors and not upon the IEEE, the Society/Council, or its members. **IEEE Corporate Office:** 3 Park Avenue, 17th Floor, New York, NY 10016-5997. **IEEE Operations Center:** 445 Hoes Lane, Piscataway, NJ 08854-4141. **NJ Telephone:** +1 732 981 0060. **Price/Publication Information:** Individual copies: IEEE Members $20.00 (first copy only), nonmembers $309.00 per copy. (Note: Postage and handling charge not included.) Member and nonmember subscription prices available upon request. Available in microfiche and microfilm. Periodicals Postage at New York, NY and at additional mailing offices. **Postmaster:** Send address changes to IEEE SIGNAL PROCESSING LETTERS, IEEE, 445 Hoes Lane, Piscataway, NJ 08854-4141. GST Registration No. 125634188. CPC Sales Agreement #40013087. Return undeliverable Canada addresses to: Pitney Bowes IMEX, P.O. Box 4332, Stanton Rd., Toronto, ON M5W 3J4, Canada. IEEE prohibits discrimination, harassment and bullying. For more information visit http://www.ieee.org/nondiscrimination. Printed in U.S.A.


IEEE
SignalProcessing
MAGAZINE
[VOLUME 32 NUMBER 5 SEPTEMBER 2015]
BIOMETRICS SECURITY AND PRIVACY PROTECTION
RECENT ADVANCES
SIGNAL PROCESSING-DRIVEN IMAGING TECHNOLOGIES
COMMEMORATING STEREO SOUND RECORDING
CREATING ANALYTIC DSP ONLINE HOMEWORK
IEEE Signal Processing Society
IEEE

CONTENTS

[VOLUME 32 NUMBER 5]

SPECIAL SECTION—BIOMETRICS SECURITY AND PRIVACY PROTECTION

Digital Object Identifier 10.1109/MSP.2015.2444511

COLUMNS

DEPARTMENT

Call for Papers

http://ssp2016.tsc.uc3m.es

2016 IEEE Statistical Signal Processing Workshop

26-29 June 2016, Palma de Mallorca, Spain

The 2016 IEEE Workshop on Statistical Signal Processing (SSP 2016) is the 19th of a series of unique meetings that bring members of the IEEE Signal Processing Society together with researchers from allied fields such as bioinformatics, communications, machine learning, and statistics.

The scientific program of SSP 2016 will include invited plenary talks, as well as regular and special sessions with contributed research papers. All submitted papers will be reviewed by experts and only a proportion will be accepted to maintain a high quality workshop. All accepted papers will be published on IEEE Xplore. The scope of the workshop includes basic theory, methods and algorithms, and applications in the following areas:

Theoretical Topics

- Adaptive systems and signal processing
- Detection and estimation theory
- Learning theory and pattern recognition
- Multivariate statistical analysis
- System identification and calibration
- Monte Carlo methods
- Network and graph analysis
- Random matrix theory
- Time-frequency and time-scale analysis
- Compressed sensing
- Point process estimation
- Stochastic filtering

Application Areas

- Bioinformatics and genomics
- Array processing, radar and sonar
- Communication systems and networks
- Sensor networks
- Information forensics and security
- Medical imaging
- Biomedical signal processing
- Preventive, social network analysis
- Smart grids and industrial applications
- Geoscience
- Astrophysics
- New methods, directions and applications

IEEE Signal Processing Society IEEE

Venue

Es Baluard Museu d'Art Modern i Contemporani, Palma de Mallorca, Spain

Paper Submission

Prospective authors are invited to submit full-length papers, with up to four pages for technical content including figures and references, using the templates and formatting guidelines posted on the website. All accepted papers must be presented at the workshop in order to be published in the proceedings. Best student paper awards, selected by a SSP committee, will be presented at the workshop.

Special Sessions

In addition to regular sessions, the workshop will also have a number of special sessions on topics of particular relevance. Prospective organizers of special sessions are invited to submit a proposal form, available on the workshop website, by e-mail to the Special Sessions Chair.

Important Dates

Submission of proposals for special sessions	**Nov 09, 2015**
Notification of acceptance of special sessions	**Nov 30, 2015**
Full paper submission deadline	**Feb 08, 2016**
Notificacion of acceptance	**April 04, 2016**
Camera ready papers due on	**April 18, 2016**

Organization

General Chairs:

Antonio Artés-Rodríguez (Universidad Carlos III de Madrid, Spain)

Joaquín Miguez (Universidad Carlos III de Madrid, Spain)

Technical Program Chairs:

Sergios Theodoridis (National and Kapodistrian University of Athens, Greece)

Konstantinos Slavakis (University of Minnesota, USA)

Finance Chair:

Matilde Sánchez-Fernández (Universidad Carlos III de Madrid, Spain)

Special Sessions Chair:

Mónica F. Bugallo (Stony Brook University, USA)

Local Arrangements Chairs:

Guillem Femenias (Universitat de les Illes Balears, Spain)

Felip Riera-Palou (Universitat de les Illes Balears, Spain)

Publications Chair

Pau Closas (CTTC, Spain)

General Chairs
Tsuhan Chen, Cornell Univ.
Ming-Ting Sun, Univ. Washington
Cha Zhang, Microsoft Research
Program Chairs
Philip Chou, Microsoft Research
Anthony Vetro, MERL
Max Mühlhäuser, TU Darmstadt
Lap-Pui Chau, NTU
Jenq-Neng Huang, Univ. Washington
Yung-Hsiang Lu, Purdue Univ.
Finance Chairs
Ying Li, IBM Research
Yi Wu, Intel Labs
Plenary Chairs
John Apostolopoulos, Cisco
Antonio Ortega, USC
Workshop Chairs
Pascal Frossard, EPFL
Ivana Tosic, Ricoh
Tutorial Chairs
Yap-Peng Tan, NTU
Lexing Xie, Australian Natl. Univ.
Special Session Chairs
Aljoscha Smolic, Disney Research
Luigi Atzoni, Univ. of Cagliari
Panel Chairs
Fernando Pereira, IST
Gene Cheung, NII
Award Chair
Chang Wen Chen, SUNY Buffalo
Industrial Program Chairs
Onur Guleryuz, Polytechnic Univ.
Ton Kalker, Huawei
Student Program Chairs
Jane Z. Wang, UBC
Ivan Bajić, Simon Fraser Univ.
Grand Challenge Chairs
Christian Timmerer, UNIKLU
Andrew Gallagher, Google
Demo/Expo Chairs
Jacob Chakareski, Univ. of Alabama
Qiong Liu, FXPAL
Local/Events Chairs
Zicheng Liu, Microsoft Research
Jue Wang, Adobe Research
Lu Xia, Amazon
Publicity Chairs
Kiyoharu Aizawa, Univ. of Tokyo
Maria Martini, KCOL
Sponsorship Chairs
Belle Tseng, Apple Inc.
Yen-Kuang Chen, Intel Research
Publication Chairs:
Junsong Yuan, NTU
Chia-Wen Lin, NTHU
Registration Chairs:
YingLi Tian, CUNY,
Yan Tong, Univ. of South Carolina
Web Chair
Jie Liang, Simon Fraser Univ.

CALL FOR PAPERS

IEEE International Conference on Multimedia and Expo (ICME) 2016

July 11-15, 2016 · Seattle, USA

With around 1000 submissions and 500 participants each year, the IEEE International Conference on Multimedia & Expo (ICME) has been the flagship multimedia conference sponsored by four IEEE societies since 2000. It serves as a forum to promote the exchange of the latest advances in multimedia technologies, systems, and applications from both the research and development perspectives of the circuits and systems, communications, computer, and signal processing communities. In 2016, an Exposition of multimedia products, prototypes and animations will be held in conjunction with the conference.

Authors are invited to submit a full paper (two-column format, 6 pages maximum) according to the guidelines available on the conference website at http://icme2016.org/. Only electronic submissions will be accepted. Topics of interest include, but are not limited to:

- Speech, audio, image, video, text and new sensor signal processing
- Signal processing for media integration
- 3D visualization and animation
- 3D imaging and 3DTV
- Virtual reality and augmented reality
- Multi-modal multimedia computing systems and human-machine interaction
- Multimedia communications and networking
- Media content analysis
- Multimedia quality assessment
- Multimedia security and content protection
- Multimedia databases and digital libraries
- Multimedia applications and services
- Multimedia standards and related issues

ICME 2016 aims to have high quality oral and poster presentations. Several awards sponsored by industry and institutions will be given out. Best papers will be presented in a single-track session to all participants. Accepted papers should be presented, or else they will not be included in the IEEE Xplore Library.

A number of Workshops will be organized by the sponsoring societies. To further foster new emerging topics, ICME 2016 also welcomes researchers, developers and practitioners to organize regular Workshops. Industrial exhibitions are held in parallel with the main conference. Proposals for Special Sessions, Tutorials, and Demos are also invited. Please visit the ICME 2016 website for submission details.

Special Session Proposals Due: October 1, 2015
Notification of Special Session Acceptance: October 17, 2015
Regular Paper Abstract Submission: November 30, 2015
Regular Paper Submission: December 4, 2015
Workshop Proposals Due: November 20, 2015
Notification of Workshop Proposal Acceptance: December 15, 2015
Panel/Tutorial Proposals Due: January 15, 2016
Notification of Panel/Tutorial Acceptance: February 29, 2016
Notification of Regular Paper Acceptance: March 11, 2016
Workshop & Demo Paper Submission: March 18, 2016
Notification of Workshop and Demo Paper Acceptance: April 22, 2016
Camera-Ready Papers Due: May 13, 2016
Exhibition Application: May 13, 2016
Conference Website: http://icme2016.org/

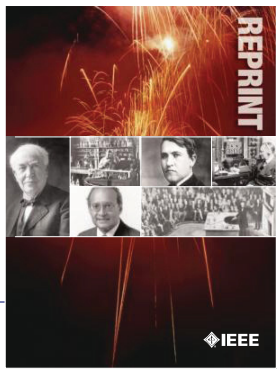

IEEE ORDER FORM FOR REPRINTS

Purchasing IEEE Papers in Print is easy, cost-effective and quick.
Complete this form, send via our secure fax (24 hours a day) to 732-981-8062 or mail it back to us.

PLEASE FILL OUT THE FOLLOWING

Author: _____

Publication Title: _____

Paper Title: _____

_____

RETURN THIS FORM TO:
IEEE Publishing Services
445 Hoes Lane
Piscataway, NJ 08855-1331

Email the Reprint Department at reprints@ieee.org for questions regarding this form

PLEASE SEND ME

■ 50 ■ 100 ■ 200 ■ 300 ■ 400 ■ 500 or _____(in multiples of 50) reprints.
■ YES ■ NO Self-covering/title page required. COVER PRICE: $74 per 100, $39 per 50.
■ $58.00 Air Freight must be added for all orders being shipped outside the U.S.
■ $21.50 must be added for all USA shipments to cover the cost of UPS shipping and handling.

PAYMENT

■ Check enclosed. Payable on a bank in the USA.
■ Charge my: ■ Visa ■ Mastercard ■ Amex ■ Diners Club

Account # _____ Exp. date _____

Cardholder's Name (please print): _____

_____

■ Bill me (you must attach a purchase order) Purchase Order Number _____

Send Reprints to:

_____
_____
_____
_____

Bill to address, if different:

_____
_____
_____
_____

Because information and papers are gathered from various sources, there may be a delay in receiving your reprint request. This is especially true with postconference publications. Please provide us with contact information if you would like notification of a delay of more than 12 weeks.

Telephone: _____ Fax: _____ Email Address: _____

2012 REPRINT PRICES (without covers)

Number of Text Pages

	1-4	5-8	9-12	13-16	17-20	21-24	25-28	29-32	33-36	37-40	41-44	45-48
50	$129	$213	$245	$248	$288	$340	$371	$408	$440	$477	$510	$543
100	$245	$425	$479	$495	$573	$680	$742	$817	$885	$953	$1021	$1088

Larger quantities can be ordered. Email reprints@ieee.org with specific details.

Tax Applies on shipments of regular reprints to CA, DC, FL, MI, NJ, NY, OH and Canada (GST Registration no. 12534188).
Prices are based on black & white printing. Please call us for full color price quote, if applicable.

2015 IEEE MEMBERSHIP APPLICATION

(students and graduate students must apply online)

IEEE Advancing Technology for Humanity

Start your membership immediately: Join online www.ieee.org/join

Please complete both sides of this form, typing or **printing in capital letters**. Use only English characters and abbreviate only if more than 40 characters and spaces per line. We regret that incomplete applications cannot be processed.

1 Name & Contact Information

Please PRINT your name as you want it to appear on your membership card and IEEE correspondence. As a key identifier for the IEEE database, circle your last/surname.

☐ Male ☐ Female Date of birth (Day/Month/Year) ____/____/____

Title First/Given Name Middle Last/Family Surname

▼ Primary Address ☐ Home ☐ Business (All IEEE mail sent here)

Street Address

City State/Province

Postal Code Country

Primary Phone

Primary E-mail

▼ Secondary Address ☐ Home ☐ Business

Company Name Department/Division

Street Address City State/Province

Postal Code Country

Secondary Phone

Secondary E-mail

To better serve our members and supplement member dues, your postal mailing address is made available to carefully selected organizations to provide you with information on technical services, continuing education, and conferences. Your e-mail address is not rented by IEEE. Please check box only if you do not want to receive these postal mailings to the selected address. ☐

2 Attestation

I have graduated from a three- to five-year academic program with a university-level degree.
☐ Yes ☐ No

This program is in one of the following fields of study:
☐ Engineering
☐ Computer Sciences and Information Technologies
☐ Physical Sciences
☐ Biological and Medical Sciences
☐ Mathematics
☐ Technical Communications, Education, Management, Law and Policy
☐ Other (please specify): _____

This academic institution or program is accredited in the country where the institution is located. ☐ Yes ☐ No ☐ Do not know

I have _____ years of professional experience in teaching, creating, developing, practicing, or managing within the following field:

☐ Engineering
☐ Computer Sciences and Information Technologies
☐ Physical Sciences
☐ Biological and Medical Sciences
☐ Mathematics
☐ Technical Communications, Education, Management, Law and Policy
☐ Other (please specify): _____

3 Please Tell Us About Yourself

Select the numbered option that best describes yourself. This information is used by IEEE magazines to verify their annual circulation. Please enter numbered selections in the boxes provided.

A. Primary line of business → ☐

1. Computers
2. Computer peripheral equipment
3. Software
4. Office and business machines
5. Test, measurement and instrumentation equipment
6. Communications systems and equipment
7. Navigation and guidance systems and equipment
8. Consumer electronics/appliances
9. Industrial equipment, controls and systems
10. ICs and microprocessors
11. Semiconductors, components, sub-assemblies, materials and supplies
12. Aircraft, missiles, space and ground support equipment
13. Oceanography and support equipment
14. Medical electronic equipment
15. OEM incorporating electronics in their end product (not elsewhere classified)
16. Independent and university research, test and design laboratories and consultants (not connected with a mfg. co.)
17. Government agencies and armed forces
18. Companies using and/or incorporating any electronic products in their manufacturing, processing, research or development activities
19. Telecommunications services, telephone (including cellular)
20. Broadcast services (TV, cable, radio)
21. Transportation services (airline, railroad, etc.)
22. Computer and communications and data processing services
23. Power production, generation, transmission and distribution
24. Other commercial users of electrical, electronic equipment and services (not elsewhere classified)
25. Distributor (reseller, wholesaler, retailer)
26. University, college/other educational institutions, libraries
27. Retired
28. Other_____

B. Principal job function → ☐

1. General and corporate management
2. Engineering management
3. Project engineering management
4. Research and development management
5. Design engineering management—analog
6. Design engineering management—digital
7. Research and development engineering
8. Design/development engineering—analog
9. Design/development engineering—digital
10. Hardware engineering
11. Software design/development
12. Computer science
13. Science/physics/mathematics
14. Engineering (not elsewhere specified)
15. Marketing/sales/purchasing
16. Consulting
17. Education/teaching
18. Retired
19. Other_____

C. Principal responsibility → ☐

1. Engineering and scientific management
2. Management other than engineering
3. Engineering design
4. Engineering
5. Software: science/mngmnt/engineering
6. Education/teaching
7. Consulting
8. Retired
9. Other_____

D. Title → ☐

1. Chairman of the Board/President/CEO
2. Owner/Partner
3. General Manager
4. VP Operations
5. VP Engineering/Dir. Engineering
6. Chief Engineer/Chief Scientist
7. Engineering Management
8. Scientific Management
9. Member of Technical Staff
10. Design Engineering Manager
11. Design Engineer
12. Hardware Engineer
13. Software Engineer
14. Computer Scientist
15. Dean/Professor/Instructor
16. Consultant
17. Retired
18. Other_____

Are you now or were you ever a member of IEEE?
☐ Yes ☐ No If yes, provide, if known:

Membership Number Grade Year Expired

4 Please Sign Your Application

I hereby apply for IEEE membership and agree to be governed by the IEEE Constitution, Bylaws, and Code of Ethics. I understand that IEEE will communicate with me regarding my individual membership and all related benefits. **Application must be signed.**

Signature Date

Over Please

5 Add IEEE Society Memberships (Optional)

The 39 IEEE Societies support your technical and professional interests. Many society memberships include a personal subscription to the core journal, magazine, or newsletter of that society. **For a complete list of everything included with your IEEE Society membership, visit www.ieee.org/join.** All prices are quoted in US dollars.

Please check ✔ the appropriate box.		BETWEEN 16 AUG 2014-28 FEB 2015 **PAY**	BETWEEN 1 MAR 2015-15 AUG 2015 **PAY**
IEEE Aerospace and Electronic Systems ⌘ 🖂	AES010	25.00 ☐	12.50 ☐
IEEE Antennas and Propagation ⌘ 🖂	AP003	15.00 ☐	7.50 ☐
IEEE Broadcast Technology ▪ ⌘ 🖂	BT002	15.00 ☐	7.50 ☐
IEEE Circuits and Systems ▪ ⌘ 🖂	CAS004	19.00 ☐	9.50 ☐
IEEE Communications ⌘ 🖂	COM019	30.00 ☐	15.00 ☐
IEEE Components, Packaging, & Manu. Tech. ▪ ⌘ 🖂	CPMT021	15.00 ☐	7.50 ☐
IEEE Computational Intelligence ⌘ 🖂	CIS011	29.00 ☐	14.50 ☐
IEEE Computer ▪ ⌘ 🖂	C016	56.00 ☐	28.00 ☐
IEEE Consumer Electronics ⌘ 🖂	CE008	20.00 ☐	10.00 ☐
IEEE Control Systems ⌘ 🖂	CS023	25.00 ☐	12.50 ☐
IEEE Dielectrics and Electrical Insulation ⌘ 🖂	DEI032	26.00 ☐	13.00 ☐
IEEE Education ⌘ 🖂	E025	20.00 ☐	10.00 ☐
IEEE Electromagnetic Compatibility ▪ ⌘ 🖂	EMC027	31.00 ☐	15.50 ☐
IEEE Electron Devices ▪ ⌘ 🖂	ED015	18.00 ☐	9.00 ☐
IEEE Engineering in Medicine and Biology ⌘ 🖂	EMB018	40.00 ☐	20.00 ☐
IEEE Geoscience and Remote Sensing ⌘ 🖂	GRS029	19.00 ☐	9.50 ☐
IEEE Industrial Electronics ⌘ 🖂	IE013	15.00 ☐	7.50 ☐
IEEE Industry Applications ⌘ 🖂	IA034	20.00 ☐	10.00 ☐
IEEE Information Theory ▪ ⌘ 🖂	IT012	30.00 ☐	15.00 ☐
IEEE Instrumentation and Measurement ⌘ 🖂	IM009	29.00 ☐	14.50 ☐
IEEE Intelligent Transportation Systems ▪ ⌘ 🖂	ITSS038	35.00 ☐	17.50 ☐
IEEE Magnetics ▪ 🖂 ♦⌘	MAG033	26.00 ☐	13.00 ☐
IEEE Microwave Theory and Techniques ⌘ 🖂	MTT017	17.00 ☐	8.50 ☐
IEEE Nuclear and Plasma Sciences ▪ ⌘ 🖂	NPS005	35.00 ☐	17.50 ☐
IEEE Oceanic Engineering ▪ ⌘ 🖂	OE022	19.00 ☐	9.50 ☐
IEEE Photonics ▪ ⌘ 🖂	PHO036	34.00 ☐	17.00 ☐
IEEE Power Electronics ⌘ 🖂	PEL035	25.00 ☐	12.50 ☐
IEEE Power & Energy ▪ ⌘ 🖂	PE031	35.00 ☐	17.50 ☐
IEEE Product Safety Engineering ▪	PSE043	35.00 ☐	17.50 ☐
IEEE Professional Communication ▪ ⌘ 🖂	PC026	31.00 ☐	15.50 ☐
IEEE Reliability ▪ ⌘ 🖂	RL007	35.00 ☐	17.50 ☐
IEEE Robotics and Automation ⌘ 🖂	RA024	9.00 ☐	4.50 ☐
IEEE Signal Processing ▪ ⌘ 🖂	SP001	20.00 ☐	10.00 ☐
IEEE Social Implications of Technology ▪ ⌘ 🖂	SIT030	33.00 ☐	16.50 ☐
IEEE Solid-State Circuits ⌘ 🖂	SSC037	29.00 ☐	14.50 ☐
IEEE Systems, Man, & Cybernetics ⌘ 🖂	SMC028	12.00 ☐	6.00 ☐
IEEE Technology & Engineering Management ⌘ 🖂	TEM014	35.00 ☐	17.50 ☐
IEEE Ultrasonics, Ferroelectrics, & Frequency Control ▪ ⌘ 🖂	UFFC020	20.00 ☐	10.00 ☐
IEEE Vehicular Technology ▪ ⌘ 🖂	VT006	18.00 ☐	9.00 ☐

Legend—Society membership includes:

⌘ One or more Society publications
▪ Society newsletter
🖂 Online access to publication
♦ CD-ROM of selected society publications

Complete both sides of this form, sign, and return to:
IEEE MEMBERSHIP APPLICATION PROCESSING
445 HOES LN, PISCATAWAY, NJ 08854-4141 USA
or fax to +1 732 981 0225
or join online at www.ieee.org/join

Please reprint your full name here

6 2015 IEEE Membership Rates (student rates available online)

IEEE member dues and regional assessments are based on where you live and when you apply. Membership is based on the calendar year from 1 January through 31 December. All prices are quoted in US dollars.

Please check ✔ the appropriate box. RESIDENCE	BETWEEN 16 AUG 2014-28 FEB 2015 **PAY**	BETWEEN 1 MAR 2015-15 AUG 2015 **PAY**
United States	$193.00 ☐	$96.50 ☐
Canada (GST)*	$171.25 ☐	$85.63 ☐
Canada (NB, NF and ON HST)*	$182.85 ☐	$91.43 ☐
Canada (Nova Scotia HST)*	$185.75 ☐	$92.88 ☐
Canada (PEI HST)*	$184.30 ☐	$92.15 ☐
Canada (GST and QST Quebec)	$185.71 ☐	$92.86 ☐
Africa, Europe, Middle East	$158.00 ☐	$79.00 ☐
Latin America	$149.00 ☐	$74.50 ☐
Asia, Pacific	$150.00 ☐	$75.00 ☐

*IEEE Canada Business No. 125634188

Minimum Income or Unemployed Provision

Applicants who certify that their prior year income did not exceed US$14,500 (or equivalent) or were not employed are granted 50% reduction in: full-year dues, regional assessment and fees for one IEEE Membership plus one Society Membership. If applicable, please check appropriate box and adjust payment accordingly. Student members are not eligible.

☐ I certify I earned less than US$14,500 in 2014 or 2013
☐ I certify that I was unemployed in 2014 or 2013

7 More Recommended Options

Proceedings of the IEEE ..... print $45.00 ☐ **or** online $39.00 ☐
Proceedings of the IEEE (print/online combination) ..... $55.00 ☐
IEEE Standards Association (IEEE-SA) ..... $52.00 ☐
IEEE Women in Engineering (WIE) ..... $25.00 ☐

8 Payment Amount

Please total the Membership dues, Society dues, and other amounts from this page:

IEEE Membership dues ⑥ ..... $_____
IEEE Society dues (optional) ⑤ ..... $_____
IEEE-SA/WIE dues (optional) ⑦ ..... $_____
Proceedings of the IEEE (optional) ⑦ ..... $_____
Canadian residents pay 5% GST or appropriate HST (BC—12%; NB, NF, ON-13%;NS-15%) on Society payments & publications only ..... TAX $_____
AMOUNT PAID ..... **TOTAL** $_____

Payment Method

All prices are quoted in US dollars. You may pay for IEEE membership by credit card (see below), check, or money order payable to IEEE, drawn on a US bank.

☐ Check ☐ VISA ☐ MasterCard ☐ Diners Club International ☐ AMERICAN EXPRESS ☐ DISCOVER

Credit Card Number

MONTH EXPIRATION DATE | YEAR | CARDHOLDER'S 5-DIGIT ZIPCODE (BILLING STATEMENT ADDRESS) USA ONLY

Name as it appears on card

Signature

☐ Auto Renew my Memberships and Subscriptions (available when paying by credit card).
☐ I agree to the Terms and Conditions located at www.ieee.org/autorenew

9 Were You Referred to IEEE?

☐ Yes ☐ No If yes, provide the following:
Member Recruiter Name _____
IEEE Recruiter's Member Number (Required) _____

CAMPAIGN CODE	PROMO CODE

14-MEM-119 7/14

Information for Authors

(Updated/Effective January 2015)

For Transactions and Journals:

Authors are encouraged to submit manuscripts of Regular papers (papers which provide a complete disclosure of a technical premise), or Comment Correspondences (brief items that provide comment on a paper previously published in these TRANSACTIONS).

Submissions/resubmissions must be previously unpublished and may not be under consideration elsewhere.

Every manuscript must:

i. provide a clear statement of the problem and what the contribution of the work is to the relevant research community;
ii. state why this contribution is significant (what impact it will have);
iii. provide citation of the published literature most closely related to the manuscript; and
iv. state what is distinctive and new about the current manuscript relative to these previously published works.

By submission of your manuscript to these TRANSACTIONS, all listed authors have agreed to the authorship list and all the contents and confirm that the work is original and that figures, tables and other reported results accurately reflect the experimental work. In addition, the authors all acknowledge that they accept the rules established for publication of manuscripts, including agreement to pay all overlength page charges, color charges, and any other charges and fees associated with publication of the manuscript. Such charges are not negotiable and cannot be suspended. The corresponding author is responsible for obtaining consent from all co-authors and, if needed, from sponsors before submission.

In order to be considered for review, a paper must be within the scope of the journal and represent a novel contribution. A paper is a candidate for an Immediate Rejection if it is of limited novelty, e.g. a straightforward combination of theories and algorithms that are well established and are repeated on a known scenario. Experimental contributions will be rejected without review if there is insufficient experimental data. These TRANSACTIONS are published in English. Papers that have a large number of typographical and/or grammatical errors will also be rejected without review.

In addition to presenting a novel contribution, acceptable manuscripts must describe and cite related work in the field to put the contribution in context. Do not give theoretical derivations or algorithm descriptions that are easily found in the literature; merely cite the reference.

New and revised manuscripts should be prepared following the "Manuscript Submission" guidelines below, and submitted to the online manuscript system, ScholarOne Manuscripts. Do not send original submissions or revisions directly to the Editor-in-Chief or Associate Editors; they will access your manuscript electronically via the ScholarOne Manuscript system.

Manuscript Submission. Please follow the next steps.

1. *Account in ScholarOne Manuscripts.* If necessary, create an account in the on-line submission system ScholarOne Manuscripts. Please check first if you already have an existing account which is based on your e-mail address and may have been created for you when you reviewed or authored a previous paper.
2. *Electronic Manuscript.* Prepare a PDF file containing your manuscript in double-column, single-spaced format using a font size of 10 points or larger, having a margin of at least 1 inch on all sides. Upload this version of the manuscript as a PDF file "double.pdf" to the ScholarOne-Manuscripts site. Since many reviewers prefer a larger font, you are strongly encouraged to also submit a single-column, double-spaced version (11 point font or larger), which is easy to create with the templates provided **IEEE Author Digital Toolbox (http://www.ieee.org/publications_standards/publications/authors/authors_journals.html)**. Page length restrictions will be determined by the double-column version. Proofread your submission, confirming that all figures and equations are visible in your document before you "SUBMIT" your manuscript. Proofreading is critical; once you submit your manuscript, the manuscript cannot be changed in any way. You may also submit your manuscript as a .PDF or MS Word file. The system has the capability of converting your files to PDF, however it is your responsibility to confirm that the conversion is correct and there are no font or graphics issues prior to completing the submission process.
3. *EDICS (Not applicable to Journal of Selected Topics in Signal Processing).* All submissions must be classified by the author with an EDICS (Editors' Information Classification Scheme) selected from the list of EDICS published online at the at the publication's EDICS webpage (*please see the list below). Upon submission of a new manuscript, please choose the EDICS categories that best suit your manuscript. Failure to do so will likely result in a delay of the peer review process.
4. *Additional Documents for Review.* Please upload pdf versions of all items in the reference list that are not publicly available, such as unpublished (submitted) papers. Graphical abstracts and supplemental materials intended to appear with the final paper (see below) must also be uploaded for review at the time of the initial submission for consideration in the review process. Use short filenames without spaces or special characters. When the upload of each file is completed, you will be asked to provide a description of that file.
5. *Supplemental Materials.* IEEE Xplore can publish multimedia files (audio, images, video), datasets, and software (e.g. Matlab code) along with your paper. Alternatively, you can provide the links to such files in a README file that appears on Xplore along with your paper. For details, please see IEEE Author Digital Toolbox under "Multimedia." To make your work reproducible by others, these TRANSACTIONS encourages you to submit all files that can recreate the figures in your paper.
6. *Submission.* After uploading all files and proofreading them, submit your manuscript by clicking "Submit." A confirmation of the successful submission will open on screen containing the manuscript tracking number and will be followed with an e-mail confirmation to the corresponding and all contributing authors. Once you click "Submit," your manuscript cannot be changed in any way.
7. *Copyright Form and Consent Form.* By policy, IEEE owns the copyright to the technical contributions it publishes on behalf of the interests of the IEEE, its authors, and their employers; and to facilitate the appropriate reuse of this material by others. To comply with the IEEE copyright policies, authors are required to sign and submit a completed "IEEE Copyright and Consent Form" prior to publication by the IEEE. The IEEE recommends authors to use an effective electronic copyright form (eCF) tool within the ScholarOne Manuscripts system. You will be redirected to the "IEEE Electronic Copyright Form" wizard at the end of your original submission; please simply sign the eCF by typing your name at the proper location and click on the "Submit" button.

Comment Correspondence. Comment Correspondences provide brief comments on material previously published in these TRANSACTIONS. These items may not exceed 2 pages in double-column, single spaced format, using 9 point type, with margins of 1 inch minimum on all sides, and including: title, names and contact information for authors, abstract, text, references, and an appropriate number of illustrations and/or tables. Correspondence items are submitted in the same way as regular manuscripts (see "Manuscript Submission" above for instructions).

Authors may also submit manuscripts of overview articles, but note that these include an additional white paper approval process http://www.signalprocessingsociety.org/publications/overview-articles/. *[This does not apply to the Journal of Selected Topics in Signal Processing. Please contact the Editor-in-Chief.]*

Digital Object Identifier

Manuscript Length. For the initial submission of a regular paper, the manuscript may not exceed 13 double-column pages (10 point font), including title; names of authors and their complete contact information; abstract; text; all images, figures and tables, appendices and proofs; and all references. Supplemental materials and graphical abstracts are not included in the page count. For regular papers, the revised manuscript may not exceed 16 double-column pages (10 point font), including title; names of authors and their complete contact information; abstract; text; all images, figures and tables, appendices and proofs; and all references. For Overview Papers, the maximum length is double that for regular submissions at each stage (please reference http://www.signalprocessingsociety.org/publications/overview-articles/ for more information).

Note that any paper in excess of 10 pages will be subject to mandatory overlength page charges. Since changes recommended as a result of peer review may require additions to the manuscript, it is strongly recommended that you practice economy in preparing original submissions. Note: Papers submitted to the TRANSACTIONS ON MULTIMEDIA in excess of 8 pages will be subject to mandatory overlength page charges.

Exceptions to manuscript length requirements may, under extraordinary circumstances, be granted by the Editor-in-Chief. However, such exception does not obviate your requirement to pay any and all overlength or additional charges that attach to the manuscript.

Resubmission of Previously Rejected Manuscripts. Authors of manuscripts rejected from any journal are allowed to resubmit their manuscripts only once. At the time of submission, you will be asked whether your manuscript is a new submission or a resubmission of an earlier rejected manuscript. If it is a resubmission of a manuscript previously rejected by any journal, you are expected to submit supporting documents identifying the previous submission and detailing how your new version addresses all of the reviewers' comments. Papers that do not disclose connection to a previously rejected paper or that do not provide documentation as to changes made may be immediately rejected.

Author Misconduct. Author misconduct includes plagiarism, self-plagiarism, and research misconduct, including falsification or misrepresentation of results. All forms of misconduct are unacceptable and may result in sanctions and/or other corrective actions. Plagiarism includes copying someone else's work without appropriate credit, using someone else's work without clear delineation of citation, and the uncited reuse of an author's previously published work that also involves other authors. Self-plagiarism involves the verbatim copying or reuse of an authors own prior work without appropriate citation, including duplicate submission of a single journal manuscript to two different journals, and submission of two different journal manuscripts which overlap substantially in language or technical contribution. For more information on the definitions, investigation process, and corrective actions related to author misconduct, see the Signal Processing Society Policies and Procedures Manual, Section 6.1. http://www.signalprocessingsociety.org/about-sps/governance/policy-procedure/part-2. Author misconduct may also be actionable by the IEEE under the rules of Member Conduct.

Extensions of the Author's Prior Work. It is acceptable for conference papers to be used as the basis for a more fully developed journal submission. Still, authors are required to cite their related prior work; the papers cannot be identical; and the journal publication must include substantively novel aspects such as new experimental results and analysis or added theoretical work. The journal publication should clearly specify how the journal paper offers novel contributions when citing the prior work. Limited overlap with prior journal publications with a common author is allowed only if it is necessary for the readability of the paper, and the prior work must be cited as the primary source.

Submission Format. Authors are required to prepare manuscripts employing the on-line style files developed by IEEE, which include guidelines for abbreviations, mathematics, and graphics. All manuscripts accepted for publication will require the authors to make final submission employing these style files. The style files are available on the web at the **IEEE Author Digital Toolbox** under "Template for all TRANSACTIONS." (LaTeX and MS Word). Please note the following requirements about the abstract:

- The abstract must be a concise yet comprehensive reflection of what is in your article.
- The abstract must be self-contained, without abbreviations, footnotes, displayed equations, or references.
- The abstract must be between 150-250 words.
- The abstract should include a few keywords or phrases, as this will help readers to find it. Avoid over-repetition of such phrases as this can result in a page being rejected by search engines.

In addition to written abstracts, papers may include a graphical abstract; see http://www.ieee.org/publications_standards/publications/authors/authors_journals.html. for options and format requirements.

IEEE supports the publication of author names in the native language alongside the English versions of the names in the author list of an article. For more information, see "Author names in native languages" (http://www.ieee.org/publications_standards/publications/authors/auth_names_native_lang.pdf) on the IEEE Author Digital Toolbox page.

Open Access. The publication is a hybrid journal, allowing either Traditional manuscript submission or Open Access (author-pays OA) manuscript submission. Upon submission, if you choose to have your manuscript be an Open Access article, you commit to pay the discounted $1,750 OA fee if your manuscript is accepted for publication in order to enable unrestricted public access. Any other application charges (such as overlength page charge and/or charge for the use of color in the print format) will be billed separately once the manuscript formatting is complete but prior to the publication. If you would like your manuscript to be a Traditional submission, your article will be available to qualified subscribers and purchasers via IEEE Xplore. No OA payment is required for Traditional submission.

Page Charges.

Voluntary Page Charges. Upon acceptance of a manuscript for publication, the author(s) or his/her/their company or institution will be asked to pay a charge of $110 per page to cover part of the cost of publication of the first ten pages that comprise the standard length (two pages, in the case of Correspondences).

Mandatory Page Charges The author(s) or his/her/their company or institution will be billed $220 per each page in excess of the first ten published pages for regular papers and six published pages for correspondence items. (**NOTE: Papers accepted to IEEE TRANSACTIONS ON MULTIMEDIA in excess of 8 pages will be subject to mandatory overlength page charges.) These are mandatory page charges and the author(s) will be held responsible for them. They are not negotiable or voluntary. The author(s) signifies his willingness to pay these charges simply by submitting his/her/their manuscript to the TRANSACTIONS. The Publisher holds the right to withhold publication under any circumstance, as well as publication of the current or future submissions of authors who have outstanding mandatory page charge debt. No mandatory overlength page charges will be applied to overview articles in the Society's journals.

Color Charges. Color figures which appear in color only in the electronic (Xplore) version can be used free of charge. In this case, the figure will be printed in the hardcopy version in grayscale, and the author is responsible that the corresponding grayscale figure is intelligible. Color reproduction charges for print are the responsibility of the author. Details of the associated charges can be found on the IEEE Publications page.

Payment of fees on color reproduction is not negotiable or voluntary, and the author's agreement to publish the manuscript in these TRANSACTIONS is considered acceptance of this requirement.

*EDICS Webpages:

IEEE TRANSACTIONS ON SIGNAL PROCESSING:
http://www.signalprocessingsociety.org/publications/periodicals/tsp/TSP-EDICS/
IEEE TRANSACTIONS ON IMAGE PROCESSING:
http://www.signalprocessingsociety.org/publications/periodicals/image-processing/tip-edics/
IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE / ACM:
http://www.signalprocessingsociety.org/publications/periodicals/taslp/taslp-edics/
IEEE TRANSACTIONS ON INFORMATION, FORENSICS AND SECURITY:
http://www.signalprocessingsociety.org/publications/periodicals/forensics/forensics-edics/
IEEE TRANSACTIONS ON MULTIMEDIA:
http://www.signalprocessingsociety.org/tmm/tmm-edics/
IEEE TRANSACTIONS ON COMPUTATIONAL IMAGING:
http://www.signalprocessingsociety.org/publications/periodicals/tci/tci-edics/
IEEE TRANSACTIONS ON SIGNAL AND INFORMATION PROCESSING OVER NETWORKS:
http://www.signalprocessingsociety.org/publications/periodicals/tsipn/tsipn-edics/

2015 IEEE SIGNAL PROCESSING SOCIETY MEMBERSHIP APPLICATION

Mail to: **IEEE OPERATIONS CENTER, ATTN: Louis Curcio, Member and Geographic Activities, 445 Hoes Lane, Piscataway, New Jersey 08854 USA**
or Fax to (732) 981-0225 (credit card payments only.)
For info call (732) 981-0060 or 1 (800) 678-IEEE or E-mail: new.membership@ieee.org

1. PERSONAL INFORMATION

NAME AS IT SHOULD APPEAR ON IEEE MAILINGS: SEND MAIL TO: ☐ Home Address OR ☐ Business/School Address
If not indicated, mail will be sent to home address. Note: Enter your name as you wish it to appear on membership card and all correspondence.
PLEASE PRINT Do not exceed 40 characters or spaces per line. Abbreviate as needed. Please circle your last/surname as a key identifier for the IEEE database.

TITLE | FIRST OR GIVEN NAME | MIDDLE NAME | SURNAME/LAST NAME

HOME ADDRESS

CITY | STATE/PROVINCE | POSTAL CODE | COUNTRY

2

Are you now or were you ever a member of IEEE? ☐ Yes ☐ No
If yes, please provide, if known:

MEMBERSHIP NUMBER_____

Grade_____**Year Membership Expired:**_____

3. BUSINESS/PROFESSIONAL INFORMATION

Company Name

Department/Division

Title/Position | Years in Current Position

Years in the Profession Since Graduation | ☐ PE State/Province

Street Address

City | State/Province | Postal Code | Country

4. EDUCATION

A baccalaureate degree from an IEEE recognized educational program assures assignment of "Member" grade. For other, additional information and references may be necessary for grade assignment.

A.

Baccalaureate Degree Received | Program/Course of Study

College/University | Campus

State/Province | Country | Mo./Yr. Degree Received

B.

Highest Technical Degree Received | Program/Course of Study

College/University | Campus

State/Province | Country | Mo./Yr. Degree Received

5.

Full signature of applicant

6. DEMOGRAPHIC INFORMATION – ALL APPLICANTS -

Date Of Birth _____ ☐ Male ☐ Female
Day Month Year

7. CONTACT INFORMATION

Office Phone/Office Fax | Home Phone/Home Fax

Office E-Mail | Home E-Mail

8. 2015 IEEE MEMBER RATES

IEEE DUES Residence	16 Aug 14-28 Feb 15 Pay Full Year	1 Mar -15 Aug 15 Pay Half Year**
United States	$193.00 ☐	$96.50 ☐
Canada (incl. GST)	$171.25 ☐	$85.63 ☐
Canada (incl. HST for PEI)	$184.30 ☐	$92.15 ☐
Canada (incl. HST for Nova Scotia)	$185.75 ☐	$92.88 ☐
Canada (incl. HST for NB, NF and ON)	$182.85 ☐	$91.43 ☐
Canada (incl. GST and QST Quebec)	$185.71 ☐	$92.86 ☐
Africa, Europe, Middle East	$158.00 ☐	$79.00 ☐
Latin America	$149.00 ☐	$74.50 ☐
Asia, Pacific	$150.00 ☐	$75.00 ☐

Canadian Taxes (GST/HST): All supplies, which include dues, Society membership fees, online products and publications (except CD-ROM and DVD media), shipped to locations within Canada are subject to the GST of 5% or the HST of 13%, 14% or 15%, depending on the Province to which the materials are shipped. GST and HST do not apply to Regional Assessments. (IEEE Canadian Business Number 12563 4188 RT0001)

Value Added Tax (VAT) in the European Union: In accordance with the European Union Council Directives 2002/38/EC and 77/388/EEC amended by Council Regulation (EC)792/2002, IEEE is required to charge and collect VAT on electronic/digitized products sold to private consumers that reside in the European Union. The VAT rate applied is the EU member country standard rate where the consumer is resident. (IEEE's VAT registration number is EU826000081)

U.S. Sales Taxes: Please add applicable state and local sales and use tax on orders shipped to **Alabama, Arizona, California, Colorado, District of Columbia, Florida, Georgia, Illinois, Indiana, Kentucky, Massachusetts, Maryland, Michigan, Minnesota, Missouri, New Jersey, New Mexico, New York, North Carolina, Ohio, Oklahoma, West Virginia, Wisconsin.** Customers claiming a tax exemption must include an appropriate and properly completed tax-exemption certificate with their first order.

IEEE Signal Processing Society

2015 SPS MEMBER RATES

		16 Aug-28 Feb Pay Full Year	1 Mar-15 Aug Pay Half Year
Signal Processing Society Membership Fee*		$ 20.00 ☐	$ 10.00 ☐
Fee includes: **IEEE Signal Processing Magazine** (electronic and digital), **Inside Signal Proc. eNewsletter** (electronic) and **IEEE Signal Processing Society Content Gazette** (electronic).			
Add $15 to enhance SPS Membership and also receive:		$15.00 ☐	$ 7.50 ☐
IEEE Signal Processing Magazine (print) and **SPS Digital Library**: online access to Signal Processing Magazine, Signal Processing Letters, Journal of Selected Topics in Signal Processing, Trans. on Audio, Speech, and Language Processing, Trans. on Image Processing, Trans. on Information Forensics and Security and Trans. on Signal Processing.			
Publications available only with SPS membership:			
Signal Processing, IEEE Transactions on:	Print	$190.00 ☐	$ 95.00 ☐
Audio, Speech, and Lang. Proc., IEEE/ACM Trans. on:	Print	$145.00 ☐	$ 72.50 ☐
Image Processing, IEEE Transactions on:	Print	$188.00 ☐	$ 94.00 ☐
Information Forensics and Security, IEEE Trans. on:	Print	$163.00 ☐	$ 81.50 ☐
IEEE Journal of Selected Topics in Signal Processing:	Print	$160.00 ☐	$ 80.00 ☐
Affective Computing, IEEE Transactions on:	Electronic	$ 35.00 ☐	$ 17.50 ☐
Biomedical and Health Informatics, IEEE Journal of:	Print	$ 55.00 ☐	$ 27.50 ☐
	Electronic	$ 40.00 ☐	$ 20.00 ☐
	Print & Electronic	$ 65.00 ☐	$ 32.50 ☐
IEEE Cloud Computing	Electronic and Digital	$ 39.00 ☐	$ 19.50 ☐
***New!* IEEE Trans. on Cognitive Comm. &Networking**	Electronic	$ 26.00 ☐	$ 13.00 ☐
***New!* IEEE Trans. on Computational Imaging**	Electronic	$ 28.00 ☐	$ 14.00 ☐
***New!* IEEE Trans. on Big Data**	Electronic	$ 25.00 ☐	$ 12.50 ☐
***New!* IEEE Trans. on Molecular, Biological, & Multi-scale Communications**	Electronic	$ 24.00 ☐	$ 12.00 ☐
IEEE Internet of Things Journal	Electronic	$ 26.00 ☐	$ 13.00 ☐
IEEE Trans. on Cloud Computing	Electronic	$ 42.00 ☐	$ 21.00 ☐
IEEE Trans. on Computational Social Systems	Electronic	$ 30.00 ☐	$ 15.00 ☐
***New!* IEEE Trans. on Signal & Info Proc. Over Networks**	Electronic	$ 28.00 ☐	$ 14.00 ☐
IEEE Biometrics Compendium:	Online	$ 30.00 ☐	$ 15.00 ☐
Computing in Science & Engrg. Mag.:	Electronic and Digital	$ 39.00 ☐	$ 19.50 ☐
	Print	$ 149.00 ☐	$ 74.50 ☐
Medical Imaging, IEEE Transactions on:	Print	$ 74.00 ☐	$ 37.00 ☐
	Electronic	$ 53.00 ☐	$ 26.50 ☐
	Print & Electronic	$ 89.00 ☐	$ 44.50 ☐
Mobile Computing, IEEE Transactions on:	ELE/Print Abstract/CD-ROM	$ 40.00 ☐	$ 20.00 ☐
Multimedia, IEEE Transactions on:	Electronic	$ 42.00 ☐	$ 21.00 ☐
IEEE MultiMedia Magazine:	Electronic and Digital	$ 39.00 ☐	$ 19.50 ☐
	Print	$149.00 ☐	$ 74.50 ☐
Network Science and Engrg., IEEE Trans. on:	Electronic	$ 33.00 ☐	$ 16.50 ☐
IEEE Reviews in Biomedical Engineering:	Print	$ 25.00 ☐	$ 12.50 ☐
	Electronic	$ 25.00 ☐	$ 12.50 ☐
	Print & Electronic	$ 40.00 ☐	$ 20.00 ☐
IEEE Security and Privacy Magazine:	Electronic and Digital	$ 39.00 ☐	$ 19.50 ☐
	Print	$149.00 ☐	$ 74.50 ☐
IEEE Sensors Journal:	Print	$150.00 ☐	$ 75.00 ☐
	Electronic	$ 50.00 ☐	$ 25.00 ☐
Smart Grid, IEEE Transactions on:	Print	$100.00 ☐	$ 50.00 ☐
	Electronic	$ 40.00 ☐	$ 20.00 ☐
	Print & Electronic	$120.00 ☐	$ 60.00 ☐
Wireless Communications, IEEE Transactions on:	Print	$120.00 ☐	$ 60.00 ☐
	Electronic	$ 48.00 ☐	$ 24.00 ☐
	Print & Electronic	$120.00 ☐	$ 60.00 ☐
IEEE Wireless Communications Letters:	Print	$ 80.00 ☐	$ 40.00 ☐
	Electronic	$ 18.00 ☐	$ 9.00 ☐
	Print & Electronic	$ 95.00 ☐	$ 47.50 ☐
***New!* IEEE Life Sciences Letters** (Open Access Pub)	Electronic	No Fee	

*IEEE membership required or requested
Affiliate application to join SP Society only.

Amount Paid $_____

9.

IEEE Membership Affiliate Fee (See pricing in Section 8) $_____

Signal Processing Society Fees $_____

Canadian residents pay 5% GST or 13% HST
Reg. No. 125634188 on Society payment(s) & pubs only Tax $_____

AMOUNT PAID WITH APPLICATION TOTAL $_____
Prices subject to change without notice.

☐ **Check or money order enclosed Payable to IEEE on a U.S. Bank**
☐ **American Express** ☐ **VISA** ☐ **MasterCard**
☐ **Diners Club**

Exp. Date/ Mo./Yr.

Cardholder Zip Code Billing Statement Address/USA Only

Full signature of applicant using credit card | Date

10. WERE YOU REFERRED?

☐ Yes ☐ No If yes, please provide the follow information:
Member Recruiter Name:_____
IEEE Recruiter's Member Number (Required): _____

2015 IEEE SIGNAL PROCESSING SOCIETY STUDENT MEMBERSHIP APPLICATION

(Current and reinstating IEEE members joining SPS complete areas 1, 2, 8, 9.)

Mail to: **IEEE OPERATIONS CENTER, ATTN: Louis Curcio, Member and Geographic Activities, 445 Hoes Lane, Piscataway, New Jersey 08854 USA**
or Fax to (732) 981-0225 (credit card payments only.)
For info call (732) 981-0060 or 1 (800) 678-IEEE or E-mail: new.membership@ieee.org

1. PERSONAL INFORMATION

NAME AS IT SHOULD APPEAR ON IEEE MAILINGS: SEND MAIL TO: ☐ Home Address OR **☐ Business/School Address**
If not indicated, mail will be sent to home address. Note: Enter your name as you wish it to appear on membership card and all correspondence.
PLEASE PRINT Do not exceed 40 characters or spaces per line. Abbreviate as needed. Please circle your last/surname as a key identifier for the IEEE database.

TITLE | FIRST OR GIVEN NAME | MIDDLE NAME | SURNAME/LAST NAME

HOME ADDRESS

CITY | STATE/PROVINCE | POSTAL CODE | COUNTRY

2.

Are you now or were you ever a member of IEEE? ☐ Yes ☐ No
If yes, please provide, if known:

MEMBERSHIP NUMBER_____

Grade_____**Year Membership Expired:**_____

3. BUSINESS/PROFESSIONAL INFORMATION

Company Name

Department/Division

Title/Position | Years in Current Position

Years in the Profession Since Graduation | ☐ PE State/Province

Street Address

City | State/Province | Postal Code | Country

4. EDUCATION

A baccalaureate degree from an IEEE recognized educational program assures assignment of "Member" grade. For others, additional information and references may be necessary for grade assignment.

A.

Baccalaureate Degree Received | Program/Course of Study

College/University | Campus

State/Province | Country | Mo./Yr. Degree Received

B.

Highest Technical Degree Received | Program/Course of Study

College/University | Campus

State/Province | Country | Mo./Yr. Degree Received

5.

Full signature of applicant

6. DEMOGRAPHIC INFORMATION – ALL APPLICANTS -

Date Of Birth _____
Day Month Year

☐ Male ☐ Female

7. CONTACT INFORMATION

Office Phone/Office Fax | Home Phone/Home Fax

Office E-Mail | Home E-Mail

8.

2015 IEEE STUDENT MEMBER RATES

IEEE DUES Residence	16 Aug 14-28 Feb 15 Pay Full Year	1 Mar -15 Aug 15 Pay Half Year**
United States	$32.00 ☐	$16.00 ☐
Canada (incl. GST)	$33.60 ☐	$16.80 ☐
Canada (incl. HST for NB, NF, and ON)	$36.16 ☐	$18.08 ☐
Canada (incl. HST for Nova Scotia)	$36.80 ☐	$18.40 ☐
Canada (incl. HST for PEI)	$36.48 ☐	$18.24 ☐
Canada (incl. GST and QST Quebec)	$36.79 ☐	$18.40 ☐
Africa, Europe, Middle East, Latin America, Asia, Pacific	$27.00 ☐	$13.50 ☐

Canadian Taxes (GST/HST): All supplies, which include dues, Society membership fees, online products and publications (except CD-ROM and DVD media), shipped to locations within Canada are subject to the GST of 5% or the HST of 13%,14% or 15%, depending on the Province to which the materials are shipped. GST and HST do not apply to Regional Assessments. (IEEE Canadian Business Number 12563 4188 RT0001)

Value Added Tax (VAT) in the European Union: In accordance with the European Union Council Directives 2002/38/EC and 77/388/EEC amended by Council Regulation (EC)792/2002, IEEE is required to charge and collect VAT on electronic/digitized products sold to private consumers that reside in the European Union. The VAT rate applied is the EU member country standard rate where the consumer is resident. (IEEE's VAT registration number is EU826000081)

U.S. Sales Taxes: Please add applicable state and local sales and use tax on orders shipped to **Alabama, Arizona, California, Colorado, District of Columbia, Florida, Georgia, Illinois, Indiana, Kentucky, Massachusetts, Maryland, Michigan, Minnesota, Missouri, New Jersey, New Mexico, New York, North Carolina, Ohio, Oklahoma, West Virginia, Wisconsin.** Customers claiming a tax exemption must include an appropriate and properly completed tax-exemption certificate with their first order.

2015 SPS STUDENT MEMBER RATES

	16 Aug-28 Feb Pay Full Year	1 Mar-15 Aug Pay Half Year
Signal Processing Society Membership Fee*	$10.00 ☐	$ 5.00 ☐

Fee includes: **IEEE Signal Processing Magazine** (electronic and digital), **Inside Signal Processing eNewsletter** (electronic) and **IEEE Signal Processing Society Content Gazette** (electronic).

Add $8 to enhance SPS Membership and also receive: $ 8.00 ☐ $ 4.00 ☐
IEEE Signal Processing Society Magazine (print) and **SPS Digital Library**: online access to Signal Processing Magazine, Signal Processing Letters, Journal of Selected Topics in Signal Processing, Trans. on Audio, Speech, and Language Processing, Trans. on Image Processing, Trans. on Information Forensics and Security and Trans. on Signal Processing.

Publications available only with SPS membership:

Signal Processing, IEEE Transactions on:	Print	$ 95.00 ☐	$ 47.50 ☐
Audio, Speech, and Lang. Proc., IEEE/ACM Trans. on:	Print	$ 73.00 ☐	$ 36.50 ☐
Image Processing, IEEE Transactions on:	Print	$ 94.00 ☐	$ 47.00 ☐
Information Forensics and Security, IEEE Trans. on:	Print	$ 82.00 ☐	$ 41.00 ☐
IEEE Journal of Selected Topics in Signal Processing:	Print	$ 80.00 ☐	$ 40.00 ☐
Affective Computing, IEEE Transactions on:	Electronic	$ 18.00 ☐	$ 9.00 ☐
Biomedical and Health Informatics, IEEE Journal of:	Print	$ 28.00 ☐	$ 14.00 ☐
	Electronic	$ 20.00 ☐	$ 10.00 ☐
	Print & Electronic	$ 65.00 ☐	$ 32.50 ☐
IEEE Cloud Computing	Electronic and Digital	$ 20.00 ☐	$ 10.00 ☐
***New!* IEEE Trans. on Cognitive Comm. &Networking**	Electronic	$ 13.00 ☐	$ 6.50 ☐
***New!* IEEE Trans. on Computational Imaging**	Electronic	$ 14.00 ☐	$ 7.00 ☐
***New!* IEEE Trans. on Big Data**	Electronic	$ 13.00 ☐	$ 6.50 ☐
***New!* IEEE Trans. on Molecular, Biological, & Multi-Scale Communications**	Electronic	$ 12.00 ☐	$ 6.00 ☐
IEEE Internet of Things Journal	Electronic	$ 13.00 ☐	$ 6.50 ☐
IEEE Trans. on Cloud Computing	Electronic	$ 21.00 ☐	$ 10.50 ☐
IEEE Trans. on Computational Social Systems	Electronic	$ 15.00 ☐	$ 7.50 ☐
***New!* IEEE Trans. on Signal & Info Proc. Over Networks**	Electronic	$ 14.00 ☐	$ 7.00 ☐
IEEE Biometrics Compendium:	Online	$ 15.00 ☐	$ 7.50 ☐
Computing in Science & Engrg. Mag.:	Electronic and Digital	$ 20.00 ☐	$ 10.00 ☐
	Print	$ 75.00 ☐	$ 37.50 ☐
Medical Imaging, IEEE Transactions on:	Print	$ 37.00 ☐	$ 18.50 ☐
	Electronic	$ 27.00 ☐	$ 13.50 ☐
	Print & Electronic	$ 45.00 ☐	$ 22.50 ☐
Mobile Computing, IEEE Transactions on:	ELE/Print Abstract/CD-ROM	$ 20.00 ☐	$ 10.00 ☐
Multimedia, IEEE Transactions on:	Electronic	$ 21.00 ☐	$ 10.50 ☐
IEEE MultiMedia Magazine:	Electronic and Digital	$ 20.00 ☐	$ 10.00 ☐
	Print	$ 75.00 ☐	$ 37.50 ☐
Network Science and Engrg., IEEE Trans. on:	Electronic	$ 17.00 ☐	$ 8.50 ☐
IEEE Reviews in Biomedical Engineering:	Print	$ 13.00 ☐	$ 6.50 ☐
	Electronic	$ 13.00 ☐	$ 6.50 ☐
	Print & Electronic	$ 20.00 ☐	$ 10.00 ☐
IEEE Security and Privacy Magazine:	Electronic and Digital	$ 20.00 ☐	$ 10.00 ☐
	Print	$ 75.00 ☐	$ 37.50 ☐
IEEE Sensors Journal:	Print	$150.00 ☐	$ 75.00 ☐
	Electronic	$ 28.00 ☐	$ 14.00 ☐
Smart Grid, IEEE Transactions on:	Print	$ 50.00 ☐	$ 25.00 ☐
	Electronic	$ 20.00 ☐	$ 10.00 ☐
	Print & Electronic	$ 60.00 ☐	$ 30.00 ☐
Wireless Communications, IEEE Transactions on:	Print	$ 60.00 ☐	$ 30.00 ☐
	Electronic	$ 24.00 ☐	$ 12.00 ☐
	Print & Electronic	$ 60.00 ☐	$ 30.00 ☐
IEEE Wireless Communications Letters:	Print	$ 40.00 ☐	$ 20.00 ☐
	Electronic	$ 9.00 ☐	$ 4.50 ☐
	Print & Electronic	$ 48.00 ☐	$ 24.00 ☐
***New!* IEEE Life Sciences Letters** (Open Access Pub)	Electronic	No Fee	

*IEEE membership required or requested
Affiliate application to join SP Society only.

9.

Amount Paid $_____

IEEE Membership Fee (See pricing in Section 8) $_____
Signal Processing Society Fees $_____

Canadian residents pay 5% GST or 13% HST
Reg. No. 125634188 on Society payment(s) & pubs only Tax $_____

AMOUNT PAID WITH APPLICATION TOTAL $_____
Prices subject to change without notice.

☐ Check or money order enclosed Payable to IEEE on a U.S. Bank
☐ American Express ☐ VISA ☐ MasterCard ☐ Diners Club

Exp. Date/ Mo./Yr.

Cardholder Zip Code Billing Statement Address/USA Only

Full signature of applicant using credit card Date

10. WERE YOU REFERRED?

☐ Yes ☐ No If yes, please provide the follow information:
Member Recruiter Name:_____
IEEE Recruiter's Member Number (Required): _____

IEEE Signal Processing Society