

Biometric Template Protection: Bridging the Performance Gap Between Theory and Practice

Karthik Nandakumar, *Member, IEEE*, and Anil K. Jain, *Fellow, IEEE*

Abstract—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 non-invertibility, revocability, and non-linkability properties without significantly degrading the recognition performance. The objective of this work is to analyze the factors contributing to this performance gap and highlight promising research directions to bridge this gap. 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 non-invertibility and non-linkability of existing algorithms. Finally, achieving non-linkability without the use of external secrets (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.

I. 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 de-duplication), 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 Advanced Encryption Standard (AES) and RSA cryptosystem. This approach has two main drawbacks. Firstly, 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 iPhone6¹, Privaris plusID²), 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.

A. 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 (\mathbf{x}), a biometric template protection algorithm generates and stores a protected biometric reference (\mathbf{v}) derived from the original template. Note that the term “protected biometric reference” not only includes the protected biometric 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.,

K. Nandakumar is with IBM Research, Singapore, e-mail: nkarthik@sg.ibm.com

A. K. Jain is with the Department of Computer Science and Engineering, Michigan State University, East Lansing, MI 48824 USA, e-mail: jain@cse.msu.edu

¹<http://support.apple.com/en-sg/HT5949>

²<http://www.privaris.com/products/indeX.html>

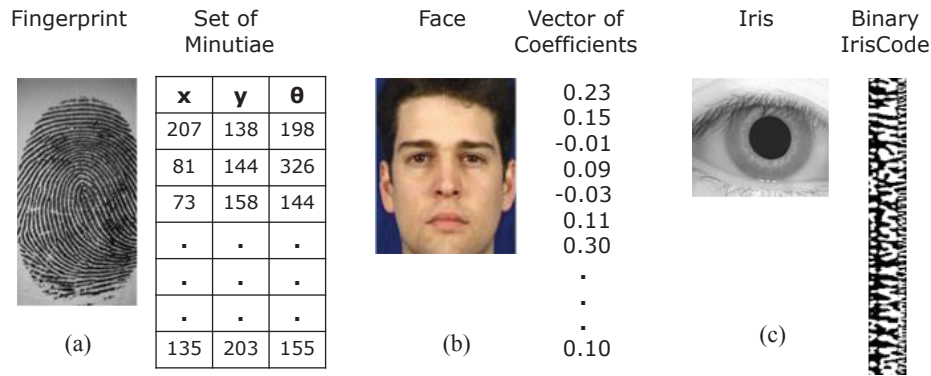


Fig. 1. Examples of biometric templates extracted from (a) fingerprint, (b) face, and (c) iris images. A fingerprint image is typically represented as an unordered set of minutiae, which encodes the location (x,y) and orientation (θ) of friction ridge discontinuities. Face images are often represented 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.

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 intra-subject 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 intra-subject 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 intra-subject variations and inter-subject variations. A highly distinctive representation should have small intra-subject variations (features extracted from multiple acquisitions of the same biometric trait of a person should be similar), but large inter-subject variations (features extracted from the same biometric trait of different individuals should be different). When minimizing intra-subject variations, care must be taken to preserve inter-subject variations. Otherwise, distinctiveness of the features may degrade, resulting in lower recognition performance.

In the context of template security, the protected biometric reference (\mathbf{v}) is typically considered as public information that is available to any adversary. Hence, \mathbf{v} should satisfy the following three properties:

- *Non-invertibility* or *Irreversibility*: It should be computationally difficult³ to obtain the original biometric template from an individual's protected biometric reference. This property 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 re-issue 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.
- *Non-linkability* 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 non-linkability property prevents cross-matching across different applications, thereby preserving the privacy of the individual.

Apart from satisfying the above 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 registry), 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.

II. BIOMETRIC TEMPLATE PROTECTION APPROACHES

Numerous template protection techniques have been proposed in the literature with the objective of ensuring non-invertibility, revocability, and non-linkability without compromising on 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, namely, *pseudonymous*

³A problem can be considered to be computationally hard or difficult if it cannot be solved using a polynomial-time algorithm.

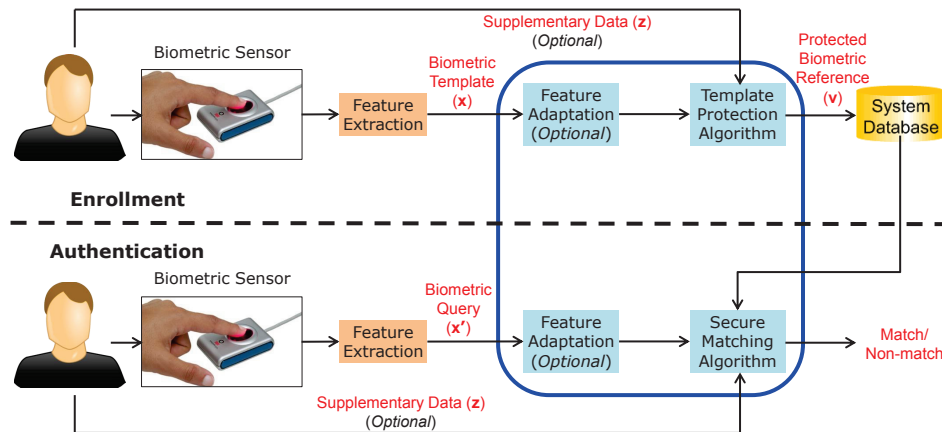


Fig. 2. General framework of a biometric system with template protection.

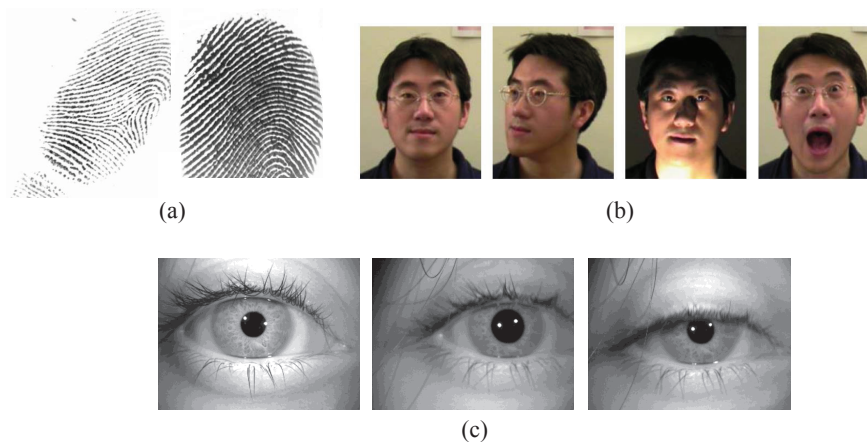


Fig. 3. Illustration of intra-subject variations observed in biometric samples. (a) Images of the same finger may exhibit variations in translation, rotation, and non-linear distortion. (b) Pose, illumination, and facial expression changes may change the appearance of face images obtained from the same person. (c) Iris images of the same eye may exhibit differences due to pupil dilation, partial closure of eyelids, and change in gaze angle.

identifier (PI) and *auxiliary data* (AD). Depending on how these two components are generated, biometric template protection schemes can be broadly categorized as: (i) feature transformation approach and (ii) biometric cryptosystems. A detailed review of biometric template protection approaches is beyond the scope of this paper and we refer the readers to [4], [5], [6] for such in-depth analysis.

In the feature transformation approach (see Figure 4(a)), a non-invertible or one-way function is applied to the biometric template (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 (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 non-invertible only when the supplementary data (e.g., key or password) is assumed to be a secret.

Techniques that can generate non-invertible 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 auxiliary data 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

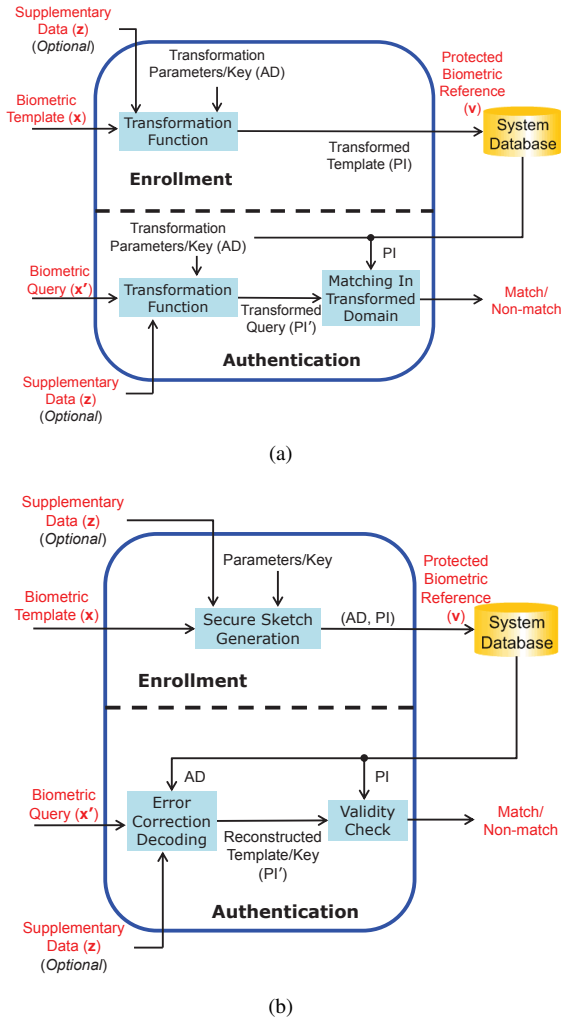


Fig. 4. There are two broad approaches for biometric template protection: (a) feature transformation and (b) biometric cryptosystem. The protected biometric reference (denoted by v) generally consists of two distinct parts, namely, *pseudonymous identifier* (PI) and *auxiliary data* (AD).

performed indirectly by attempting to recover the original template (x) using the secure sketch (AD) in conjunction with the query biometric features (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 the 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 non-invertibility, but at the same time tolerant to intra-subject 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 sets, which often leads to loss of discriminatory in-

formation and consequent degradation in recognition accuracy. Due to the properties of linear error correction codes⁴ that are commonly used in secure sketch constructions, it is difficult to achieve non-linkability in biometric cryptosystems.

One way to overcome the above 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].

A. 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 on-going Fingerprint Verification Competition (FVC-onGoing⁵). Six algorithms were able to achieve an equal error rate (EER) of less than 0.3% on the FVC-STD-1.0 benchmark dataset 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 dataset was 1.54%, which is more than 5 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 non-invertibility and non-linkability of the protected biometric reference. This is necessary to benchmark the utility of a biometric template protection scheme. In cryptography, “security strength” (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 3072-bit key can provide a security strength of approximately 128 bits⁶. However, there is no consensus within the biometrics community on analogous metrics that can be used to measure the non-invertibility, revocability, and non-linkability properties of biometric template protection algorithms as well as the methods to compute these metrics [23]. Consequently, practical template protection schemes neither have proven non-invertibility/non-linkability guarantees nor do they achieve satisfactory recognition performance. This explains why despite 20 years of research,

⁴In 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.

⁵<https://biolab.csr.unibo.it/fvcongoing/UI/Form/Home.aspx>

⁶Barker et al., “Recommendation for Key Management”, NIST 800-57, July 2012.

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 intra-subject 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 non-invertibility and recognition performance [17], [24] due to the following reason. Maximizing non-invertibility 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 be solved only 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 intra-subject variations without compromising on inter-subject 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 non-invertibility and non-linkability 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 non-invertibility. 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 non-invertibility, the problem of *ensuring non-linkability and revocability of protected biometric reference has not been adequately addressed* in the literature. While many template protection constructs claim to provide non-linkability 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 paper is to provide an in-depth analysis of the above three challenges, discuss some of the solutions that have been proposed to overcome them, and identify unresolved issues that require further research.

III. DESIGNING INVARIANT FEATURE REPRESENTATIONS

A traditional biometric system accounts for intra-subject variations in two ways. Firstly, the feature extraction algorithm attempts to extract an invariant representation from the noisy biometric samples. Secondly, the matching algorithm is designed to further suppress the effect of intra-subject 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 intra-subject variations like rotation, translation, and non-linear 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 non-linear 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 non-invertible and revocable template. In contrast, adapted templates need not satisfy the non-invertibility and revocability properties. Instead, feature adaptation schemes are designed to satisfy one or more of the following three objectives: (i) minimize intra-subject variations without diluting their distinctiveness, (ii) represent the original features in a simplified form, and (iii) 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




Fig. 5. 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 non-linear distortion. This explains why a simple set difference metric is unlikely to provide accurate recognition.

and orientation features and selection of good quality minutiae [13] when designing a fingerprint cryptosystem. Though the process of quantization and feature selection reduces intra-subject variations, it is also likely to decrease inter-subject variations. Thus, the challenge is to strike an optimum balance between reducing intra-subject variations and preserving inter-subject variations. Moreover, if quantization and reliable component selection is user-specific, the quantization parameters and selected components need to be stored as auxiliary data, which is likely to decrease the non-invertibility and non-linkability 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.

A. 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: (i) real vector into a binary string, and (ii) point set into a binary string.

1) *Real Vector to Binary String*: Conversion of a real vector into a binary string involves two essential steps: (i) quantization - mapping continuous values into discrete values, and (ii) 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 in order to maximize the entropy of the resulting binary string. While the DROBA approach optimizes the detection rate (genuine accept rate) at the minimum (low) false accept rate, it requires many training samples per subject in order 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.

2) *Point Set to Binary String*: The most well-known example of point set based biometric representation is a collection of fingerprint minutiae. 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 over-lapping) 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 section III-A1. 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 2-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.

B. 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 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




Fig. 6. 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].

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.

C. Open Issues in Feature Adaptation

Though a significant amount of research effort has been devoted towards feature adaptation, three main issues remain unresolved. Firstly, 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 intra-subject variations, while ignoring the need to preserve inter-subject 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 above 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 intra-subject 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 non-invertibility properties when it is eventually secured using a biometric cryptosystem (cf. Section IV). 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.

IV. SOLVING THE RECOGNITION PERFORMANCE VS. 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 towards solving this problem is to clearly define the notion of security, establish metrics to quantify security properties such as non-invertibility and non-linkability, 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 accept 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 non-invertibility 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-line⁷ zero effort impostor attack. 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 non-invertibility, one can find instances where the term privacy actually refers to non-linkability. 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 non-invertibility given only the protected biometric reference \mathbf{v} , whereas strong biometric privacy refers to non-invertibility 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 non-invertibility and non-linkability, respectively. To avoid confusion, it has been suggested that specific properties such as non-invertibility (or irreversibility) and non-linkability (or unlinkability) must be considered instead of employing generic terms like security and privacy [39].

A. Metrics for Measuring Non-invertibility

Non-invertibility 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 non-invertibility of a protected biometric reference.

A direct measure of non-invertibility is the conditional Shannon entropy of the original template \mathbf{x} given the protected

⁷Since most practical biometric systems restrict the number of failed authentication attempts, it is usually not possible to mount online zero effort impostor (FAR) attacks.

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 non-invertibility.

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 non-invertibility 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 $\tilde{H}_\infty(A|B) = -\log(\mathbb{E}_{b \rightarrow B} [2^{-H_\infty(A|B=b)}])$ 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 non-invertibility 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 non-invertibility 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 pre-image), 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 pre-image attacks is already factored into the non-invertibility analysis of such a system. Therefore, analysis of pre-image 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], [41]. However, for the sake of simplicity, we avoid a detailed discussion of these metrics in this paper.

B. Methods for Computing Non-invertibility Metrics

Since the non-invertibility metrics for feature transformation schemes are generally computed empirically, this section will focus only on methods to compute the non-invertibility metrics for biometric cryptosystems. While the metrics for measuring non-invertibility 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], [40], [37]. Typically, the entropy loss is an increasing function of the error correction capability of the system. In other words, if larger tolerance for intra-subject 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 $H(\mathbf{x}|\mathbf{v})$. This is because of the complexity in estimating the entropy of biometric features ($H(\mathbf{x})$).

1) *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 inter-subject 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 individuality⁸. Moreover, they rely on some simplifying assumptions in order 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

⁸More precisely, the goal in [26], [42] is to estimate the probability of a false correspondence/match between minutiae templates from two arbitrary fingerprints belonging to different fingers.

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, intra-subject 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.

C. Open Issues in Non-invertibility Analysis

Despite significant progress in analyzing the non-invertibility of template protection schemes, there is no consensus yet on the standard metrics to be used for measuring non-invertibility and well-defined methodologies to compute these metrics. Efforts to standardize these metrics are still under progress [23]. Once such metrics are standardized, the focus should shift towards the development of a suitable framework that allows joint optimization of recognition performance and non-invertibility for both feature transformation schemes and biometric cryptosystems.

One way to overcome the inherent trade-off between non-invertibility and recognition performance is to develop techniques for multibiometric⁹ template protection. 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 be higher, thereby leading to stronger non-invertibility. 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 above challenge.

V. ACHIEVING REVOCABILITY AND NON-LINKABILITY

While revocability and non-linkability are also core requirements of a template protection scheme, the analysis of these

⁹Multibiometric systems accumulate evidence from more than one biometric identifier (multiple traits like fingerprint and iris or multiple fingers/irides) in order to recognize a person.

two properties has received considerably less attention in the literature compared to non-invertibility. Recently, it has been demonstrated that many well-known biometric cryptosystems do not generate revocable or non-linkable templates [24], [27], [47], [48]. Though feature transformation schemes are widely proclaimed as “cancelable biometrics” in acknowledgement of their strengths in achieving revocability and non-linkability, 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 non-linkability of feature transformation schemes appear to depend on the difficulty in obtaining a pre-image of the transformed template. When the pre-image is easy to compute given the transformation parameters and the transformed template, it may be possible to correlate the pre-images 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 pre-image.

One possible way to achieve revocability and non-linkability is to use hybrid biometric cryptosystems [19], [20]. While a combination of secure sketch and feature transformation enhances the non-invertibility of the protected biometric reference, the feature transformation step ensures the revocability and non-linkability properties. However, this may come at the cost of a degradation in the recognition performance.

Another practical solution for achieving revocability and non-linkability 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 auxiliary data (AD) and pseudonymous identifier (PI) are not stored together in order 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 PIN 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 non-linkability, 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 & password or biometrics & 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 multi-factor authentication protocols vanish, and their properties are no better than those of the underlying template protection scheme.

VI. 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 paper, 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 non-invertibility of the protected template. This is because the information leaked by a protected template is often proportional to the tolerance allowed to account for intra-subject variations. Furthermore, standardized metrics are required for measuring the security properties of a template protection scheme, especially non-invertibility. 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 non-linkability of protected templates.

Apart from the open research issues identified earlier in the context of feature adaptation (cf. Section III-C) and non-invertibility analysis (cf. Section IV-C), 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 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 re-issue a protected biometric reference without re-enrolling the user, which is often impractical.
- 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 integration of biometric technologies. For example, the Touch-ID fingerprint recognition system in 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.

REFERENCES

- [1] A. K. Jain, A. Ross, and K. Nandakumar, *Introduction to Biometrics*. Springer, 2011.
- [2] P. Campisi, Ed., *Security and Privacy in Biometrics*. Springer, 2013.
- [3] R. Cappelli, A. Lumini, D. Maio, and D. Maltoni, "Fingerprint Image Reconstruction From Standard Templates," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 9, pp. 1489–1503, 2007.
- [4] A. K. Jain, K. Nandakumar, and A. Nagar, "Biometric Template Security," *EURASIP Journal on Advances in Signal Processing, Special Issue on Advanced Signal Processing and Pattern Recognition Methods for Biometrics*, January 2008.
- [5] C. Rathgeb and A. Uhl, "A Survey on Biometric Cryptosystems and Cancelable Biometrics," *EURASIP Journal on Information 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 Magazine*, vol. 30, no. 5, pp. 51–64, September 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 Transactions on Pattern Analysis and Machine Intelligence*, 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 Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 4, pp. 561–572, April 2007.
- [9] S. Tulyakov, F. Farooq, P. Mansukhani, and V. Govindaraju, "Symmetric hash functions for secure fingerprint biometric systems," *Pattern Recognition Letters*, 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 Journal on Computing*, vol. 38, no. 1, pp. 97–139, 2008.
- [11] A. Juels and M. Sudan, "A Fuzzy Vault Scheme," in *Proceedings of IEEE International Symposium on Information Theory, Lausanne, Switzerland, 2002*, p. 408.
- [12] A. Juels and M. Wattenberg, "A Fuzzy Commitment Scheme," in *Proceedings of Sixth ACM Conference on Computer and Communications Security*, Singapore, November 1999, pp. 28–36.
- [13] K. Nandakumar, A. K. Jain, and S. Pankanti, "Fingerprint-based Fuzzy Vault: Implementation and Performance," *IEEE Transactions on Information Forensics and Security*, vol. 2, no. 4, pp. 744–757, December 2007.
- [14] Y. Sutcu, Q. Li, and N. Memon, "Protecting Biometric Templates with Sketch: Theory and Practice," *IEEE Transactions on Information Forensics and Security*, vol. 2, no. 3, pp. 503–512, September 2007.
- [15] F. Hao, R. Anderson, and J. Daugman, "Combining Crypto with Biometrics Effectively," *IEEE Transactions on Computers*, vol. 55, no. 9, pp. 1081–1088, September 2006.
- [16] J. Bringer, H. Chabanne, G. Cohen, B. Kindarji, and G. Zemor, "Theoretical and Practical Boundaries of Binary Secure Sketches," *IEEE Transactions on Information Forensics and 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 *Proceedings of SPIE, Electronic Imaging, Media Forensics and Security XII*, San Jose, January 2010.
- [18] T. Ignatenko and F. M. J. Willems, "Biometric Systems: Privacy and Secrecy Aspects," *IEEE Transactions on Information Forensics and Security*, vol. 4, no. 4, pp. 956–973, 2009.
- [19] T. E. Boulton, W. J. Scheirer, and R. Woodworth, "Fingerprint Revocable Biotokens: Accuracy and Security Analysis," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Minneapolis, USA, 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 Transactions on Information Forensics and Security*, vol. 5, no. 1, pp. 103–117, March 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 Magazine*, vol. 30, no. 2, pp. 42–52, March 2013.

- [22] D. Gafurov, B. Yang, P. Bours, and C. Busch, "Independent Performance Evaluation of Pseudonymous Identifier Fingerprint Verification Algorithms," in *Proceedings of Intl. Conf. on Image Analysis and Recognition*, June 2013.
- [23] S. Rane, "Standardization of Biometric Template Protection," *IEEE MultiMedia*, vol. 21, no. 4, pp. 94–99, October 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 Transactions on Information Forensics and Security*, vol. 7, no. 6, pp. 1825–1840, December 2012.
- [25] A. W. K. Kong, "A Statistical Analysis of IrisCode and Its Security Implications," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 3, pp. 513–528, March 2015.
- [26] C. Su and S. Srihari, "Evaluation of Rarity of Fingerprints in Forensics," in *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 Transactions on Information Forensics and Security*, vol. 8, no. 9, pp. 1433–1445, September 2013.
- [28] S. Yang and I. Verbauwhede, "Secure Iris Verification," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 2, 2007.
- [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 *Proceedings of IEEE 3rd International Conference on 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 Journal on Advances in Signal Processing*, May 2009.
- [31] M.-H. Lim and A. B. J. Teoh, "A Novel Encoding Scheme for Effective Biometric Discretization: Linearly Separable Subcode," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 2, pp. 300–313, February 2013.
- [32] A. Nagar, S. Rane, and A. Vetro, "Privacy and Security of Features Extracted from Minutiae Aggregates," in *Proceedings IEEE International Conference on Acoustics, Speech and Signal Processing*, Dallas, TX, March 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 Transactions on Information Forensics and Security*, vol. 4, no. 3, pp. 397–409, September 2009.
- [34] K. Nandakumar, "A Fingerprint Cryptosystem based on Minutiae Phase Spectrum," in *Proceedings of Second IEEE Workshop on Information Forensics and Security*, Seattle, USA, December 2010.
- [35] R. Cappelli, M. Ferrara, and D. Maltoni, "Minutia Cylinder-Code: A New Representation and Matching Technique for Fingerprint Recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 12, pp. 2128–2141, December 2010.
- [36] I. R. Buhan, J. M. Doumen, P. H. Hartel, and R. N. J. Veldhuis, "Fuzzy Extractors for Continuous Distributions," in *Proceedings of ACM Symp. on Information, Computer and Comm. Security*, Singapore, March 2007, pp. 353–355.
- [37] L. Lai, S.-W. Ho, and H. V. Poor, "Privacy-Security Trade-Offs in Biometric Security Systems," *IEEE Transactions on Information Forensics and Security*, vol. 6, no. 1, pp. 122–151, March 2011.
- [38] L. Ballard, S. Kamara, and M. K. Reiter, "The Practical Subtleties of Biometric Key Generation," in *Proceedings of the 17th Conference on Security Symposium*, 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 *Proceedings of 5th IAPR International Conference on Biometrics*, March 2012, pp. 498–505.
- [40] T. Ignatenko and F. M. J. Willems, "Information Leakage in Fuzzy Commitment Schemes," *IEEE Transactions on Information Forensics and 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 *Proceedings of IEEE Sixth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, September 2013, pp. 1–8.
- [42] Y. Zhu, S. C. Dass, and A. K. Jain, "Statistical Models for Assessing the Individuality of Fingerprints," *IEEE Transactions on Information Forensics and Security*, vol. 2, no. 3, pp. 391–401, September 2007.
- [43] A. Adler, R. Youmaran, and S. Loyka, "Towards a Measure of Biometric Feature Information," *Pattern Analysis and Applications*, 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," *Proceedings of the IEEE*, vol. 94, no. 11, pp. 1927–1935, 2006.
- [45] K. Takahashi and T. Murakami, "A Measure of Information Gained Through Biometric Systems," *Image and Vision Computing*, vol. 32, no. 12, pp. 1194–1203, December 2014.
- [46] B. Fu, S. X. Yang, J. Li, and D. Hu, "Multibiometric Cryptosystem: Model Structure and Performance Analysis," *IEEE Transactions on Information Forensics and Security*, vol. 4, no. 4, pp. 867–882, December 2009.
- [47] X. Boyen, "Reusable Cryptographic Fuzzy Extractors," in *ACM Conference on Computer and Communications Security*, Washington D.C, USA, October 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 Transactions on Information Forensics and Security*, vol. 6, no. 1, pp. 107–121, March 2011.
- [49] X. Boyen, Y. Dodis, J. Katz, R. Ostrovsky, and A. Smith, "Secure Remote Authentication Using Biometric Data," in *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 *The IEEE International Workshop on Information Forensics and Security*, December 2010.

Karthik Nandakumar Karthik Nandakumar is a Research Staff Member at IBM Research, Singapore. Prior to joining IBM Research, he was a Scientist at 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 co-authored two books and received a number of awards including the 2010 IEEE Signal Processing Society Young Author Best Paper Award.

Anil K. Jain Anil K. Jain is a University Distinguished Professor in the Department of Computer Science & Engineering at Michigan State University. His research interests include pattern recognition, computer vision and biometric recognition. He has received Guggenheim fellowship, Humboldt Research award, Fulbright fellowship, IEEE Computer Society Technical Achievement award, IEEE W. Wallace McDowell award, IAPR King-Sun Fu Prize, and the IEEE ICDM Research Contribution Award for contributions to pattern recognition and biometrics. He is a Fellow of the ACM, IEEE, AAAS, IAPR and SPIE. He is the author of several books and ISI has designated him as a highly cited researcher.