Secure Sketch for Biometric Templates

Qiming Li[1], Yagiz Sutcu[2], and Nasir Memon[3]

[1] Department of Computer and Information Science
[2] Department of Electrical and Computer Engineering
[3] Department of Computer and Information Science
Polytechnic University
6 Metrotech Center, Brooklyn, NY 11201
qiming.li@ieee.org ygzstc@yahoo.com memon@poly.edu

Abstract. There have been active discussions on how to derive a consistent cryptographic key from noisy data such as biometric templates, with the help of some extra information called a *sketch*. It is desirable that the sketch reveals little information about the biometric templates even in the worst case (i.e., the *entropy loss* should be low). The main difficulty is that many biometric templates are represented as points in continuous domains with unknown distributions, whereas known results either work only in discrete domains, or lack rigorous analysis on the entropy loss. A general approach to handle points in continuous domains is to quantize (discretize) the points and apply a known sketch scheme in the discrete domain. However, it can be difficult to analyze the entropy loss due to quantization and to find the "optimal" quantizer. In this paper, instead of trying to solve these problems directly, we propose to examine the *relative entropy loss* of any given scheme, which bounds the number of additional bits we could have extracted if we used the optimal parameters. We give a general scheme and show that the relative entropy loss due to suboptimal discretization is at most $(n \log 3)$, where n is the number of points, and the bound is tight. We further illustrate how our scheme can be applied to real biometric data by giving a concrete scheme for face biometrics.
Keywords: Secure sketch, biometric template, continuous domain.

1 Introduction

The main challenge in using biometric data in cryptography is that they cannot be reproduced exactly. Some noise will be inevitably introduced into biometric samples during acquisition and processing. There have been active discussions on how to extract a reliable cryptographic key from such noisy data. Some recent techniques attempt to correct the noise in the data by using some public information P derived from the original biometric template X. These techniques include fuzzy commitment [12], fuzzy vault [11], helper data [19], and secure sketch [7]. In this paper, we follow Dodis et al. [7] and call such public information P a *sketch*.

Typically, there are two main components in a secure sketch scheme. The first is the sketch generation algorithm, which we will refer to as the *encoder*. It takes the original biometric template X as the input, and outputs a sketch P. The second algorithm is the biometric template reconstruction algorithm, or the *decoder*, which takes another

biometric template Y and the sketch P as the input and outputs X'. If Y and X are sufficiently similar according to some similarity measure, we will have $X = X'$. An important requirement for such a scheme is that the sketch P should not reveal too much information about the biometric template X. Dodis et al. [7] gives a notion of *entropy loss*, which (informally speaking) measures the advantage that P gives to any adversary in guessing X, when X is discrete in nature (Section 3 provides the details). It is worth to note that the entropy loss is a worst case bound for *all* distributions of X.

There are several difficulties in applying many known secure sketch techniques to known types of biometric templates directly. Firstly, many biometric templates are represented by sequences of n points in a continuous domain (say, $\mathbb{R}$), or equivalently, points in an n-dimensional space (say, $\mathbb{R}^n$). In this case, since the entropy of the original data can be very large, and the length of the extracted key is typically quite limited, the "entropy loss" as defined in [7] can be very high for any possible scheme. For example, X is often a discrete approximation of some points in a continuous domain (e.g., decimal fractions obtained by rounding real numbers). As the precision of X gets higher, both the entropy of X and the entropy loss from P become larger, but the extracted key can become stronger. Hence, this notion of entropy loss alone is insufficient, and the seemingly high entropy loss for this type of biometric data would be misleading. We will discuss this issue in detail in Section 4, and give a complimentary definition of *relative entropy loss* for noisy data in the continuous domain. Informally speaking, the relative entropy loss of a sketch measures the imperfectness of the rounding, which is the maximum amount of additional entropy we can obtain by the "optimal" rounding. At the same time, the entropy loss from P serves as a measure of the security of the sketch in the discrete domain.

Secondly, even if the biometric templates are represented in discrete form, there are practical problems when the entropy of the original template is high. For example, the iris pattern of an eye can be represented by a 2048 bit binary string called *iris code*, and up to 20% of the bits could be changed under noise [9]. The fuzzy commitment scheme based on binary error-correcting codes [12] seems to be applicable at the first glance. However, it would be impractical to apply a binary error-correcting code on such a long string with such a large error-correcting capability. A two-level error-correcting technique is proposed in [9], which essentially changes the similarity measure. As a result, the space is no longer a metric space.

Thirdly, the similarity measures for many known biometric templates can be quite different from those considered in many theoretical works (such as Hamming distance, set difference and edit distance in [7]). This can happen as a result of technical considerations (e.g., in the case of iris codes). However, in many cases this is due to the nature of biometric templates. For instance, a fingerprint template usually consists of a set of minutiae (feature points in 2-D space), and two templates are considered as similar if more than a certain number of minutiae in one template are near distinct minutiae in the other. In this case, the similarity measure has to consider both Euclidean distance and set difference at the same time.

The secure sketch for point sets [5] is perhaps the first rigorous approach to similarity measures that do not define a metric space. A generic scheme is proposed in [5] for point

sets in bounded discrete d-dimensional space for any d, where the underlying similarity measure is motivated by the similarity measure of fingerprint templates. While such a scheme is potentially applicable to fingerprints represented as minutiae, other types of biometrics are different both in representations and similarity measures, thus require different considerations and different schemes.

In this paper, we study how to design secure sketch for biometric templates, where the worst case bound can be proved. We observe that many biometric templates can be represented in a general form: The original X can be considered as a list of n points, where each point x of X is in a bounded continuous domain. Under noise, each point can be perturbed by a distance less than δ, and on top of that, at most t points can be replaced. Similar to [5], we will refer to the first noise as the *white noise*, and the second *replacement noise*. We note that this similarity measure can be applied to handwritten online signatures [8], iris patterns [9], voice features [15], and face biometrics [17]. This formulation is different from that in [5] in two ways: (1) The points are in a continuous domain, and (2) the points are always ordered.

To handle points in continuous domain, a general two step approach is to (1) quantize (i.e., discretize) the points in X to a discrete domain with a scalar quantizer $\mathcal{Q}_\lambda$, where λ is the step size, and (2) apply secure sketch techniques on the quantized points $\widehat{X} = \mathcal{Q}_\lambda(X)$ in the quantized domain, which is discrete. For example, if points in X are real numbers between 0 and 1, assume that we have a scalar quantizer $\mathcal{Q}_\lambda$ with step size $\lambda = 0.01$, such that $\mathcal{Q}_\lambda(x) = \widehat{x}$ if and only if $\widehat{x}\lambda \leq x < (\widehat{x}+1)\lambda$, then every point in X would be mapped to an integer in $[0, 99]$. After that, we can apply a secure sketch for discrete points in the domain $[0, 99]^n$ to achieve error-tolerance.

However, there are two difficulties when this approach is applied. Firstly, if we follow the notion of secure sketch and entropy loss as in [7], the quantization error $X - \widehat{X}$ in the first step has to be kept in the sketch, since exact reconstruction of X is required by definition. However, it can be difficult to give an upper bound on the entropy loss from the quantization errors. Even if we can, it can be very large.

Furthermore, as the quantization step λ becomes very small, the bound on the entropy loss in the quantized domain during the second step can be very high. For instance, for $x \in [0, 1)$ and $\delta = 0.01$, when $\lambda = 0.01$, the entropy loss in Step (2) will be $\log 3$, and the bound is tight. When $\lambda = 0.001$, the entropy loss will be $\log 21$. However, the big difference in entropy loss in the quantized domain can be misleading. We will revisit this example in Section 5, and will show that the second case actually results in a stronger key if X is uniformly distributed.

To address the above problems, we consider the following strategy. Instead of trying to answer the question of how much entropy is lost during quantization, we study how different quantizers affect the strength of the key that we can finally extract from the noisy data. In particular, given a secure sketch scheme in the discrete domain and a quantizer $\mathcal{Q}_1$ with step size λ_1, we consider any quantizer $\mathcal{Q}_2$ with step size λ_2. Assuming that m_1 and m_2 are the strengths of the keys under these two quantizers respectively, we found that it is possible to give an upper bound on the difference between m_1 and m_2, for any distribution of X, and any choices of λ_2 (hence $\mathcal{Q}_2$) within a certain range. This bound can be expressed as a function of λ_1. In other words, although we do not know what is the exact entropy loss due to the quantizer $\mathcal{Q}_1$, we

do know that at most how far away $\mathcal{Q}_1$ can be from the "optimal" one. Based on this, we give a notion of *relative entropy loss* for data in continuous domain. Furthermore, we show that if X is uniformly distributed, the relative entropy loss can be bounded by a constant for any choice of λ_1.

To illustrate how our general approach can be applied to practical biometric templates, we give a scheme based on the authentication scheme for face biometrics in [17]. We will also discuss some practical issues in designing secure sketch schemes for biometric templates.

We note that our proposed schemes and analysis can be applied for two parties to extract secret keys given correlated random variables (e.g., [14]), where the random variables take values in a continuous domain (e.g. $\mathbb{R}$). The entropy loss in the quantized domain measures how much information can be leaked to an eavesdropper, while the relative entropy loss measures how many additional bits that we might be able to extract.

We will give a review of related works in Section 2, followed by some preliminary formal definitions in Section 3. Our definition of secure sketch and its security will be presented in Section 4. We give a general similarity measure and our proposed schemes in Section 5, together with a security analysis and some discussions on choosing the parameters. A concrete secure sketch scheme for face biometrics will be given in 6.

2 Related Works

It is not surprising that the construction of the sketch largely depends on the representation of the biometric templates and the underlying distance function that measures the similarity. Most of the known techniques assume that the noisy data under consideration are represented as points in some metric space. The fuzzy commitment scheme [12], which is based on binary error-correcting codes, considers binary strings where the similarity is measured by Hamming distance. The fuzzy vault scheme [11] considers sets of elements in a finite field with set difference as the distance function, and corrects errors by polynomial interpolation. Dodis et al. [7] further gives the notion of *fuzzy extractors*, where a "strong extractor" (such as pair-wise independent hash functions) is applied after the original X is reconstructed to obtain an almost uniform key. Constructions and rigorous analysis of secure sketch are given in [7] for three metrics: Hamming distance, set difference and edit distance. Secure sketch schemes for point sets in [5] are motivated by the typical similarity measure used for fingerprints, where each template consists of a set of points in 2-D space, and the similarity measure does not define a metric space.

On the other hand, there have been a number of works on how to extract consistent keys from real biometric templates, which have quite different representations and similarity measures from the above theoretical works. Such biometric templates include handwritten online signatures [8], fingerprints [20], iris patterns [9], voice features [15], and face biometrics [17]. These works, however, do not have sufficiently rigorous treatment of the security, compared to well-established cryptographic techniques. Some of the works give analysis on the entropy of the biometrics, and approximated amount of efforts required by a brute-force attacker.

Boyen [2] shows that a sketch scheme that is provably secure may be insecure when multiple sketches of the same biometric data are obtained. Boyen et al. further study the security of secure sketch schemes under more general attacker models in [1], and techniques to achieve mutual authentication are proposed.

Linnartz and Tuyls [13] consider a similar problem for biometric authentication applications. They consider zero mean i.i.d. jointly Gaussian random vectors as biometric templates, and use mutual information as the measure of security against dishonest verifiers. Tuyls and Goseling [19] consider a similar notion of security, and develop some general results when the distribution of the original is known and the verifier can be trusted. Some practical results along this line also appear in [18].

3 Preliminaries

3.1 Entropy and Entropy Loss in Discrete Domain

In the case where X is discrete, we follow the definitions by Dodis et al. [7]. They consider a variant of the *average min-entropy* of X given P, which is essentially the minimum strength of the key that can be consistently extracted from X when P is made public.

In particular, the min-entropy $\mathbf{H}_\infty(A)$ of a discrete random variable A is defined as $\mathbf{H}_\infty(A) = -\log(\max_a \Pr[A = a])$. For two discrete random variables A and B, the average min-entropy of A given B is defined as $\widetilde{\mathbf{H}}_\infty(A \mid B) = -\log(\mathbb{E}_{b \leftarrow B}[2^{-\mathbf{H}_\infty(A|B=b)}])$.

For discrete X, the entropy loss of the sketch P is defined as $\mathcal{L} = \mathbf{H}_\infty(X) - \widetilde{\mathbf{H}}_\infty(X|P)$. This definition is useful in the analysis, since for any ℓ-bit string B, we have $\widetilde{\mathbf{H}}_\infty(A \mid B) \geq \mathbf{H}_\infty(A) - \ell$. For any secure sketch scheme for discrete X, let R be the randomness invested in constructing the sketch, it is not difficult to show that when R can be computed from X and P, we have

$$\mathcal{L} = \mathbf{H}_\infty(X) - \widetilde{\mathbf{H}}_\infty(X \mid P) \leq |P| - \mathbf{H}_\infty(R). \tag{1}$$

In other words, the entropy loss can be bounded from above by the difference between the size of P and the amount of randomness we invested in computing P. This allows us to conveniently find an upper bound of $\mathcal{L}$ for any distribution of X, since it is independent of X.

3.2 Secure Sketch in Discrete Domain

Our definitions of secure sketch and entropy loss in the discrete domain follow that in [7]. Let $\mathcal{M}$ be a finite set of points with a *similarity* relation $\mathsf{S} \subseteq \mathcal{M} \times \mathcal{M}$. When $(X, Y) \in \mathsf{S}$, we say the Y is similar to X, or the pair (X, Y) is similar.

Definition 1 *A sketch scheme in discrete domain is a tuple* $(\mathcal{M}, \mathsf{S}, \mathsf{Enc}, \mathsf{Dec})$*, where* $\mathsf{Enc} : \mathcal{M} \to \{0,1\}^*$ *is an encoder and* $\mathsf{Dec} : \mathcal{M} \times \{0,1\}^* \to \mathcal{M}$ *is a decoder such that for all* $X, Y \in \mathcal{M}$*,* $\mathsf{Dec}(Y, \mathsf{Enc}(X)) = X$ *if* $(X, Y) \in \mathsf{S}$*. The string* $P = \mathsf{Enc}(X)$ *is the sketch, and is to be made public. We say that the scheme is* $\mathcal{L}$*-secure if for all random variables* X *over* $\mathcal{M}$*, the entropy loss of the sketch* P *is at most* $\mathcal{L}$*. That is,* $\mathbf{H}_\infty(X) - \widetilde{\mathbf{H}}_\infty(X \mid \mathsf{Enc}(X)) \leq \mathcal{L}$.

We call $\widetilde{\mathbf{H}}_\infty(X \mid P)$ the *left-over entropy*, which in essence measures the "strength" of the key that can be extracted from X given that P is made public. Note that in most cases, the ultimate goal is to maximize the left-over entropy for some particular distribution of X. However, in the discrete case, the min-entropy of X is fixed but can be difficult to analyze. Hence, entropy loss becomes an equivalent measure which is easier to quantify.

4 Secure Sketch in Continuous Domain

In this section we propose a general approach to handle noisy data in a continuous domain. We consider points in a universe $\mathcal{U}$, which is a set that may be uncountable. Let S be a similarity relation on $\mathcal{U}$, i.e., $\mathsf{S} \subseteq \mathcal{U} \times \mathcal{U}$. Let $\mathcal{M}$ be a set of finite points, and let $\mathcal{Q} : \mathcal{U} \to \mathcal{M}$ be a function that maps points in $\mathcal{U}$ to points in $\mathcal{M}$. We will refer to such a function $\mathcal{Q}$ as a *quantizer*.

Definition 2 *A quantization-based sketch scheme is a tuple* $(\mathcal{U}, \mathsf{S}, \mathcal{Q}, \mathcal{M}, \mathsf{Enc}, \mathsf{Dec})$*, where* $\mathsf{Enc} : \mathcal{M} \to \{0,1\}^*$ *is an encoder and* $\mathsf{Dec} : \mathcal{M} \times \{0,1\}^* \to \mathcal{M}$ *is an decoder such that for all* $X, Y \in \mathcal{U}$*,* $\mathsf{Dec}(\mathcal{Q}(Y), \mathsf{Enc}(\mathcal{Q}(X))) = \mathcal{Q}(X)$ *if* $(X, Y) \in \mathsf{S}$*. The string* $P = \mathsf{Enc}(\mathcal{Q}(X))$ *is the sketch. We say that the scheme is* $\mathcal{L}$*-secure in the quantized domain if for all random variable* X *over* $\mathcal{U}$*, the entropy loss of* P *is at most* $\mathcal{L}$*, i.e.,* $\mathbf{H}_\infty(\mathcal{Q}(X)) - \widetilde{\mathbf{H}}_\infty(\mathcal{Q}(X) \mid \mathsf{Enc}(\mathcal{Q}(X))) \leq \mathcal{L}$

In other words, a quantization is applied to transform the points in the continuous domain to a discrete domain, and a sketch scheme for discrete domain is applied to obtain the sketch P. During reconstruction, we require the exact reconstruction of the quantization $\mathcal{Q}(X)$ instead of the original X in the continuous domain. When required, a strong extractor can be further applied to $\mathcal{Q}(X)$ to extract a key (as the fuzzy extractor in [7]). That is, we treat $\mathcal{Q}(X)$ as the "discrete original". Similarly, we call $\widetilde{\mathbf{H}}_\infty(\mathcal{Q}(X) \mid P)$ the left-over entropy.

When $\mathcal{Q}$ is fixed, we can use the entropy loss on $\mathcal{Q}(X)$ to analyze the security of the scheme, and bound the entropy loss of P. However, using this entropy loss alone may be misleading, since there are many ways to quantize X, and different quantizer would make a difference in both the min-entropy of $\mathcal{Q}(X)$ and the entropy loss. Since our ultimate goal is to maximize the left-over entropy (i.e., the average min-entropy $\widetilde{\mathbf{H}}_\infty(\mathcal{Q}(X) \mid P)$), the entropy loss alone is not sufficient to compare different quantization strategies.

To illustrate the subtleties, we consider the following example. Let x be a point uniformly distributed in the interval $[0, 1)$, and under noise, it can be shifted but still within the range $[x - 0.01, x + 0.01)$. We can use a scalar quantizer $\mathcal{Q}_1$ with step size 0.01, such that all points in the interval $[0, 1)$ are mapped to integers $[0, 99]$. In this case, the min-entropy $\mathbf{H}_\infty(\mathcal{Q}_1(x)) = \log 100$. As we can see later, there is an easy way to construct a secure sketch for such $\mathcal{Q}_1(x)$ with entropy loss of $\log 3$. Hence, the left-over entropy is $\log(100/3) \approx 5.06$. Now we consider another scalar quantizer $\mathcal{Q}_2$ with step size 0.001, such that the range of $\mathcal{Q}_2(x)$ is $[0, 999]$. A similar scheme on $\mathcal{Q}_2(x)$ would give entropy loss of $\log 21$, which seems much larger than the previous $\log 3$. However,

the min-entropy of $\mathcal{Q}_2(x)$ is also increased to $\log 1000$, and the left-over entropy would be $\log(1000/21) \approx 5.57$, which is slightly higher than the case where $\mathcal{Q}_1$ is used.

Intuitively, for a given class of methods of handling noisy data in the quantized domain, it is important to examine how different precisions of the quantization process affect the strength of the extracted key. For this purpose, we propose to consider not just one, but a family of quantizers $\mathbf{Q}$, where each quantizer $\mathcal{Q}$ drawn from $\mathbf{Q}$ defines a mapping from $\mathcal{U}$ to a finite set $\mathcal{M}_{\mathcal{Q}}$. Let $\mathbf{M}$ be the set of such $\mathcal{M}_{\mathcal{Q}}$ for all $\mathcal{Q} \in \mathbf{Q}$. We also define a family of encoders $\mathbf{E}$ and decoders $\mathbf{D}$, such that for each $\mathcal{Q}$ and $\mathcal{M}_{\mathcal{Q}}$, there exist uniquely defined $\mathsf{Enc}_{\mathcal{Q}} \in \mathbf{E}$ and $\mathsf{Dec}_{\mathcal{Q}} \in \mathbf{D}$ that can handle $\mathcal{Q}(X)$ in $\mathcal{M}_{\mathcal{Q}}$.

Definition 3 *A quantization-based sketch family is a tuple* $(\mathcal{U}, \mathsf{S}, \mathbf{Q}, \mathbf{M}, \mathbf{E}, \mathbf{D})$*, such that for each quantizer* $\mathcal{Q} \in \mathbf{Q}$*, there exist* $\mathcal{M} \in \mathbf{M}$*,* $\mathsf{Enc} \in \mathbf{E}$ *and* $\mathsf{Dec} \in \mathbf{D}$*, and* $(\mathcal{U}, \mathsf{S}, \mathcal{Q}, \mathcal{M}, \mathsf{Enc}, \mathsf{Dec})$ *is a quantization-based sketch scheme. We say that such a scheme is a member of the family, and is identified by* $\mathcal{Q}$*.*

Definition 4 *A quantization-based sketch family* $(\mathcal{U}, \mathsf{S}, \mathbf{Q}, \mathbf{M}, \mathbf{E}, \mathbf{D})$ *is* $(\mathbf{L}, \mathbf{R})$*-secure for functions* $\mathbf{L}, \mathbf{R} : \mathbf{Q} \to \mathbb{R}$ *if for any member identified by* $\mathcal{Q}_1$ *(with encoder* Enc_1*) it holds that*

1. *This member is* $\mathbf{L}(\mathcal{Q}_1)$*-secure in the quantized domain; and*
2. *For any random variable* X*, and any member identified by* $\mathcal{Q}_2$ *(with encoder* Enc_2*), we have*

$$\widetilde{\mathbf{H}}_\infty(\mathcal{Q}_2(X) \mid \mathsf{Enc}_2(\mathcal{Q}_2(X))) - \widetilde{\mathbf{H}}_\infty(\mathcal{Q}_1(X) \mid \mathsf{Enc}_1(\mathcal{Q}_1(X))) \leq \mathbf{R}(\mathcal{Q}_1).$$

In other words, to measure the security of the family of schemes, we examine two aspects of the family. Firstly, we consider the entropy loss in the quantized domain for each member of the family. This is represented by the function $\mathbf{L}$, which serves as a measure of security when the quantizer is fixed. Secondly, given any quantizer in the family, we consider the question: If we use another quantizer, how many more bits can be extracted? We call this the *relative entropy loss*, which is represented by the function $\mathbf{R}$.

We observe that for some sketch families, the relative entropy loss for any given member can be conveniently bounded by the size of of the sketch generated by that member. We say that such sketch families are *well-formed*. More precisely, we have

Definition 5 *A quantization-based sketch family* $(\mathcal{U}, \mathsf{S}, \mathbf{Q}, \mathbf{M}, \mathbf{E}, \mathbf{D})$ *is well-formed if for any two members* $(\mathcal{U}, \mathsf{S}, \mathcal{Q}_1, \mathcal{M}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ *and* $(\mathcal{U}, \mathsf{S}, \mathcal{Q}_2, \mathcal{M}_2, \mathsf{Enc}_2, \mathsf{Dec}_2)$*, it holds for any random variable* X *that*

$$\widetilde{\mathbf{H}}_\infty\left(\mathcal{Q}_1(X) \mid \langle P_1, P_2 \rangle\right) = \widetilde{\mathbf{H}}_\infty\left(\mathcal{Q}_2(X) \mid \langle P_1, P_2 \rangle\right) \tag{2}$$

where $P_1 = \mathsf{Enc}_1(\mathcal{Q}_1(X))$ *and* $P_2 = \mathsf{Enc}_2(\mathcal{Q}_2(X))$*.*

Theorem 1 *For any well-formed quantization-based sketch family, given any two members* $(\mathcal{U}, \mathsf{S}, \mathcal{Q}_1, \mathcal{M}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ *and* $(\mathcal{U}, \mathsf{S}, \mathcal{Q}_2, \mathcal{M}_2, \mathsf{Enc}_2, \mathsf{Dec}_2)$*, it holds for any random variable* X *that*

$$\widetilde{\mathbf{H}}_\infty(\mathcal{Q}_2(X) \mid P_2) - \widetilde{\mathbf{H}}_\infty(\mathcal{Q}_1(X) \mid P_1) \leq |P_1|$$

where $P_1 = \mathsf{Enc}_1(\mathcal{Q}_1(X))$ *and* $P_2 = \mathsf{Enc}_2(\mathcal{Q}_2(X))$*.*

Proof: First, it is not difficult to show that for any random variables A, B and C, we have

$$\widetilde{\mathbf{H}}_\infty(A \mid B) - |C| \leq \widetilde{\mathbf{H}}_\infty(A \mid \langle B, C\rangle) \leq \widetilde{\mathbf{H}}_\infty(A \mid B). \tag{3}$$

Let $\widehat{X}_1 = \mathcal{Q}_1(X)$ and $\widehat{X}_2 = \mathcal{Q}_2(X)$. Since the sketch family is well-formed,

$$\widetilde{\mathbf{H}}_\infty\left(\widehat{X}_1 \mid \langle P_1, P_2\rangle\right) = \widetilde{\mathbf{H}}_\infty\left(\widehat{X}_2 \mid \langle P_1, P_2\rangle\right). \tag{4}$$

Substituting B by P_1, C by P_2, and A by $\widehat{X}_1$ and $\widehat{X}_2$ respectively in (3), we have

$$\begin{aligned}\widetilde{\mathbf{H}}_\infty\left(\widehat{X}_2 \mid P_2\right) - |P_1| &\leq \widetilde{\mathbf{H}}_\infty\left(\widehat{X}_2 \mid \langle P_1, P_2\rangle\right) \\ &= \widetilde{\mathbf{H}}_\infty\left(\widehat{X}_1 \mid \langle P_1, P_2\rangle\right) \leq \widetilde{\mathbf{H}}_\infty\left(\widehat{X}_1 \mid P_1\right).\end{aligned} \tag{5}$$

□

5 A General Scheme for Biometric Templates

We observe that many biometric templates can be represented as a sequence of points in some bounded continuous domain. There are two types of noise that can occur. The first noise, *white noise*, perturbs each points by a small distance, and the second noise, *replacement noise*, replaces some points by different points.

Without loss of generality, we assume that each biometric template X can be written as a sequence $X = \langle x_1, x_2, \cdots, x_n\rangle$, where each $x_i \in \mathbb{R}$ and $0 \leq x_i < 1$. In other words, $X \in \mathcal{U} = [0, 1)^n$. For each pair of biometric templates X and Y, we say that $(X, Y) \in \mathtt{S}$ if there exists a subset C of $\{1, \cdots, n\}$, such that $|C| \geq n - t$ for some threshold t, and for every $i \in C$, it holds that $|x_i - y_i| < \delta$, for some threshold δ.

Similar to the two-part approach in [5], we construct the sketch in two parts. The first part, the *white noise sketch*, handles the white noise in the noisy data, and the second part, the *replacement noise sketch*, corrects the replacement noise. We will concentrate on the white noise sketch in this paper, and the replacement noise sketch can be implemented using a known secure sketch scheme for set difference (e.g., that in [7, 3]).

5.1 Proposed Quantization-Based Sketch Family

Each member of the family is parameterized by a λ such that $\lambda \in \mathbb{R}$ and $0 < \lambda \leq \delta$.

Quantizer $\mathcal{Q}_\lambda$. Each quantizer $\mathcal{Q}_\lambda$ in $\mathbf{Q}$ is a scalar quantizer with step size $\lambda \in \mathbb{R}$. For each $x \in \mathcal{U}$, $\mathcal{Q}_\lambda(x) = \widehat{x}$ if and only if $\lambda\widehat{x} \leq x < \lambda(\widehat{x} + 1)$, and the quantization of X is defined as $\widehat{X} = \mathcal{Q}_\lambda(X) \triangleq \langle \mathcal{Q}_\lambda(x_1), \cdots, \mathcal{Q}_\lambda(x_n)\rangle$. The corresponding quantized domain is thus $\mathcal{M}_\lambda = [0, \lceil \frac{1}{\lambda} \rceil]^n$. The encoders and the decoders work only on the quantized domain. The white noise appeared in the quantized domain is of level $\widehat{\delta}_\lambda = \lceil \delta/\lambda \rceil$. In other words, under white noise, a point $\widehat{x}$ in the quantized domain can be shifted by a distance of at most $\widehat{\delta}_\lambda$. Let us denote $\Delta_\lambda \triangleq 2\widehat{\delta}_\lambda + 1$.

Codebook $\mathcal{C}_\lambda$. Furthermore, for each quantized domain $\mathcal{M}_\lambda$ we consider a *codebook* $\mathcal{C}_\lambda$, where every codeword $c \in \mathcal{C}_\lambda$ has the form $c = k\Delta_\lambda$ for some non-negative integer k. We use $\mathcal{C}_\lambda(\cdot)$ to denote the function such that given a quantized point $\widehat{x}$, it returns a value $c = \mathcal{C}_\lambda(\widehat{x})$ such that $|\widehat{x} - c| \leq \widehat{\delta}_\lambda$. That is, the functions finds the unique codeword c that is nearest to $\widehat{x}$ in the codebook.

Encoder Enc_λ. Given a quantized $\widehat{X} \in \mathcal{M}_\lambda$, the encoder Enc_λ does the following.

1. For each $\widehat{x}_i \in \widehat{X}$, compute $c_i = \mathcal{C}_\lambda(\widehat{x}_i)$;
2. Output $P = \mathsf{Enc}_\lambda(\widehat{X}) = \langle d_1, \cdots, d_n \rangle$, where $d_i = \widehat{x}_i - c_i$ for $1 \leq i \leq n$.

In other words, for every $\widehat{x}_i$, the encoder outputs the distance of $\widehat{x}_i$ from its nearest codeword in the codebook $\mathcal{C}_\lambda$.

Decoder Dec_λ. For a corrupted template Y, it is first quantized by $\widehat{Y} = \mathcal{Q}_\lambda(Y)$. Given $P = \langle d_1, \cdots, d_n \rangle$ and $\widehat{Y} = \langle \widehat{y}_1, \cdots, \widehat{y}_n \rangle$, and the decoder Dec_λ does the following.

1. For each $\widehat{y}_i \in \widehat{Y}$, compute $c_i = \mathcal{C}_\lambda(\widehat{y}_i - d_i)$;
2. Output $\widetilde{X} = \mathsf{Dec}_\lambda(\widehat{Y}) = \langle c_1 + d_1, \cdots, c_n + d_n \rangle$.

In other words, the decoder shifts every $\widehat{y}_i$ by d_i, maps it to the nearest codeword in $\mathcal{C}_\lambda$, and shifts it back by the same distance.

5.2 Security Analysis

For each member of the sketch family with parameter λ, the difference d_i between $\widehat{x}_i$ and p_i ranges from $-\widehat{\delta}_\lambda$ to $\widehat{\delta}_\lambda$. Intuitively, $\log \Delta_\lambda$ bits are sufficient and necessary to describe the white noise in the quantized domain (recall that $\Delta_\lambda = 2\widehat{\delta}_\lambda + 1 = 2\lceil \frac{\delta}{\lambda} \rceil + 1$). Hence, we have

LEMMA 2 *The quantization-based sketch scheme* $(\mathcal{U}, \mathsf{S}, \mathcal{Q}_\lambda, \mathcal{M}_\lambda, \mathsf{Enc}_\lambda, \mathsf{Dec}_\lambda)$ *is* $(n \log \Delta_\lambda)$*-secure in the quantized domain.*

Proof: Note that the size of each d_i generated in the second step of the encoder is $\log \Delta_\lambda$. Hence the total size of the sketch is $n \log \Delta_\lambda$. Therefore, the entropy loss of the sketch P is at most $n \log \Delta_\lambda$ by Equation (1). □

It is not difficult to see that the above bound is tight. For example, when each $\widehat{x}$ is uniformly distributed in the quantized domain, the min-entropy of each $\widehat{x}$ after quantization would be $\log \lceil \frac{1}{\lambda} \rceil$, and the average min-entropy of $\widehat{x}$ given P would be at most $\log |\mathcal{C}_\lambda| = \log \lceil \frac{1}{\lambda} \rceil - \log \Delta_\lambda$.

Now we consider the relative entropy loss. First of all, we observe that the proposed sketch family is well-formed according to Definition 5.

LEMMA 3 *The quantization-based sketch family defined in Section 5.1 is well-formed.*

Proof: We consider any two members in the sketch family. The first is identified by $\mathcal{Q}_{\lambda_1}$ with step size λ_1, and the second is identified by $\mathcal{Q}_{\lambda_2}$ with step size λ_2.

For any point $x \in X$, let $\widehat{x}_1 = \mathcal{Q}_{\lambda_1}(x)$. Recall that during encoding, a codeword is computed as $c_1 = \mathcal{C}_{\lambda_1}(\widehat{x}_1)$, and the difference $d_1 = \widehat{x}_1 - c_1$ is put into the sketch. Similarly, let $\widehat{x}_2 = \mathcal{Q}_{\lambda_2}(x)$, $c_2 = \mathcal{C}_{\lambda_2}(\widehat{x}_2)$ and $d_2 = \widehat{x}_2 - c_2$.

Since $\lambda_1 \leq \delta$ and $\lambda_2 \leq \delta$, it is easy to see that if d_1, d_2 and $\widehat{x}_1$ is known, we can compute $\widehat{x}_2$ deterministically. Similarly, given d_1, d_2 and $\widehat{x}_2$, $\widehat{x}_1$ can also be determined. Thus, we have

$$\widetilde{\mathbf{H}}_\infty\left(\widehat{x}_1 \mid \langle d_1, d_2\rangle\right) = \widetilde{\mathbf{H}}_\infty\left(\langle \widehat{x}_1, \widehat{x}_2\rangle \mid \langle d_1, d_2\rangle\right) = \widetilde{\mathbf{H}}_\infty\left(\widehat{x}_2 \mid \langle d_1, d_2\rangle\right). \tag{6}$$

The same arguments can be applied to all the points in X. Hence, let $P_1 = \mathsf{Enc}_{\lambda_1}(X)$ and $P_2 = \mathsf{Enc}_{\lambda_2}(X)$, we have

$$\widetilde{\mathbf{H}}_\infty\left(\widehat{X}_1 \mid \langle P_1, P_2\rangle\right) = \widetilde{\mathbf{H}}_\infty\left(\left\langle \widehat{X}_1, \widehat{X}_2\right\rangle \mid \langle P_1, P_2\rangle\right) = \widetilde{\mathbf{H}}_\infty\left(\widehat{X}_2 \mid \langle P_1, P_2\rangle\right). \tag{7}$$

That is, the proposed sketch family is well-formed. □

By combining Theorem 1 and Lemma 3, and considering that for the member of the sketch family identified by $\mathcal{Q}_{\lambda_1}$ with step size λ_1, the size of the sketch $|P_1| = n(\log \Delta_{\lambda_1})$, we have the following lemma.

LEMMA 4 *For the quantization-based sketch family defined in Section 5.1, given any member identified by $\mathcal{Q}_{\lambda_1}$ with step size λ_1 and encoder Enc_{λ_1} it holds that, for every random variable $X \in \mathcal{U}$ and any member identified by $\mathcal{Q}_{\lambda_2}$ with step size λ_2 and encoder Enc_{λ_2}, we have*

$$\widetilde{\mathbf{H}}_\infty(\mathcal{Q}_{\lambda_2}(X) \mid \mathsf{Enc}_{\lambda_2}(\mathcal{Q}_{\lambda_2}(X))) - \widetilde{\mathbf{H}}_\infty(\mathcal{Q}_{\lambda_1}(X) \mid \mathsf{Enc}_{\lambda_1}(\mathcal{Q}_{\lambda_1}(X))) \leq n(\log \Delta_{\lambda_1}).$$

In other words, the relative entropy loss is at most $n(\log \Delta_{\lambda_1})$ for $\mathcal{Q}_{\lambda_1}$.

Not only the above is a worst case bound, we can show that the worst case can indeed happen.

LEMMA 5 *The relative entropy loss in Lemma 4 is tight for sufficiently small δ.*

Proof: For any given λ_1, we find a λ_2 such that it is possible to find $\Delta_{\lambda_1} \triangleq (2\lceil \delta/\lambda_1 \rceil + 1)$ points $W = \{w_0, \cdots, w_{\Delta_{\lambda_1}-1}\}$ such that $\mathcal{Q}_{\lambda_1}(w_i) - \mathcal{C}_{\lambda_1}(\mathcal{Q}_{\lambda_1}(w_1)) = i - \lceil \delta/\lambda_1 \rceil$, and $\mathcal{C}_{\lambda_2}(w_i) = c_i$ for some codeword $c_i \in \mathcal{C}_{\lambda_2}$. In other words, we want to find points such that each of them would generate a different d_i in the final sketch with $\mathcal{Q}_{\lambda_1}$, but would generate exactly the same number (i.e., 0) in the sketch when $\mathcal{Q}_{\lambda_2}$ is used. Note that when δ is sufficiently small, there would be sufficiently many codewords in $\mathcal{C}_{\lambda_1}$, and it is always possible to find such λ_2 (e.g., $\lambda_2 = \lambda_1/2$).

When each $x \in X$ is uniformly distributed over W, we can see that the sketch from the scheme identified by $\mathcal{Q}_{\lambda_1}$ would reveal all information about X, but in the case of $\mathcal{Q}_{\lambda_2}$, the left-over entropy would be exactly $\log \Delta_{\lambda_1}$. □

Therefore, combining lemmas 2, 4 and 5 we have

THEOREM 6 *The quantization-based sketch family defined in Section 5.1 is* $(\mathbf{L}, \mathbf{R})$*-secure where for each member in the family identified by* $\mathcal{Q}_\lambda$ *with step size* λ*, where* $\mathbf{L}(\mathcal{Q}_\lambda) = \mathbf{R}(\mathcal{Q}_\lambda) = n \log \Delta_\lambda$*. Furthermore, the bounds are tight.*

For example, if $\lambda = \delta$, we would have $\mathbf{L}(\mathcal{Q}_\lambda) = \mathbf{R}(\mathcal{Q}_\lambda) = n(\log 3)$. Note that although decreasing λ *might* give a larger left-over entropy, this is not guaranteed. In fact, if we use a $\lambda' < \lambda$, by applying the above theorem on $\mathcal{Q}_{\lambda'}$, we can see that it may result in a smaller left-over entropy than using $\mathcal{Q}_\lambda$ (e.g., consider the example in the proof of Lemma 5).

5.3 A Special Case

We further study a special case when each point $x \in X$ is independently and uniformly distributed over $[0, 1)$. We further assume that $1/\delta$ is an integer, and the family of schemes only consists of members with step size λ such that $1/\lambda$ is an integer that is a multiple of Δ_λ. This additional assumption is only for the convenience of the analysis, and would not make too much difference in practice.

In this case, the entropy loss in the quantized domain for the member identified by $\mathcal{Q}_\lambda$ with step size λ would be exactly $n(\log \Delta_\lambda)$, which shows that Lemma 2 is tight. Moreover, it is interesting that the relative entropy loss in this case can be bounded by a constant.

COROLLARY 7 *When each* $x \in X$ *is independently and uniformly distributed, the quantization-based sketch family defined in Section 5.1 is* $(\mathbf{L}, \mathbf{R})$*-secure where for each member in the family identified by* $\mathcal{Q}_\lambda$ *with step size* λ*, where* $\mathbf{L}(\mathcal{Q}_\lambda) = n(\log \Delta_\lambda)$*, and* $\mathbf{R}(\mathcal{Q}_\lambda) = n \log(1 + \frac{\lambda}{2\delta}) \leq n \log(3/2)$*.*

Proof: The claim $\mathbf{L}(\mathcal{Q}_\lambda) = n(\log \Delta_\lambda)$ follows directly from Lemma 2, so we only focus on $\mathbf{R}$. Consider two members of the family identified by $\mathcal{Q}_{\lambda_1}$ and $\mathcal{Q}_{\lambda_2}$ respectively. Without loss of generality, we assume $\lambda_1 > \lambda_2$. Consider any $x \in X$, let $\widehat{x}_1 = \mathcal{Q}_{\lambda_1}(x)$, $c_1 = \mathcal{C}_{\lambda_1}(\widehat{x}_1)$. Similarly we define $\widehat{x}_2 = \mathcal{Q}_{\lambda_2}(x)$ and $c_2 = \mathcal{C}_{\lambda_2}(\widehat{x}_2)$. Hence, the min-entropy in the quantized domain would be $\log(1/\lambda_1)$ and $\log(1/\lambda_2)$ respectively.

Clearly, c_1 and c_2 are also uniformly distributed over $\mathcal{C}_{\lambda_1}$ and $\mathcal{C}_{\lambda_2}$ respectively, and do not depend on d_1 and d_2. Hence, the left-over entropy for these two members would be $\log(|\mathcal{C}_{\lambda_1}|) = \log \frac{1}{\lambda_1 + 2\delta}$ and $\log(|\mathcal{C}_{\lambda_2}|) = \log \frac{1}{\lambda_2 + 2\delta}$ respectively. Furthermore, recall that $0 < \lambda_2 < \lambda_1 \leq \delta$, and the difference between these two quantities can be bounded as

$$\log(|\mathcal{C}_{\lambda_2}|) - \log(|\mathcal{C}_{\lambda_1}|) = \log \frac{\lambda_1 + 2\delta}{\lambda_2 + 2\delta} < \log(1 + \frac{\lambda_1}{2\delta}) \leq \log \frac{3}{2}.$$

Therefore, the relative entropy loss is bounded by $n \log(3/2)$ as claimed. □

5.4 Remarks

Choosing the step size λ*.* We can view the step size λ as a measure of the precision of $\widehat{X}$. Since the white noise in the continuous domain is fixed at δ, when λ becomes

smaller, the corresponding white noise in the quantized domain would increase, and vice versa. That is intuitively why it is not possible to obtain much more left-over entropy by simply having X represented in a higher precision. In fact, it is not difficult to show that there are certain distributions of X such that a smaller step size would reveal more information. Furthermore, the scheme can be more efficient if we use a relatively larger step size, since we would need fewer bits to represent both X and the white noise in the quantized domain. If we use the same quantizer for both encoding and decoding, the simplest form of white noise in the quantized domain can be achieved when $\lambda = \delta$, where a quantized $\widehat{x}$ can be either left unchanged, or shifted by 1. In this case, from Theorem 6, we can get at most $n \log 3$ additional bits if we choose other $\lambda' < \delta$. If X is uniformly distributed, the increment is at most $n \log(3/2)$ by Corollary 7.

When $\lambda > \delta$, the form of white noise in the quantized domain would remain unchanged, but we may lose too much information about X due to the large quantization step, which may result in a much lower left-over entropy. Therefore, it is not desirable to have a step size larger than δ in general. If different quantizers are used during encoding and decoding, with large step size (e.g., 2δ), it is possible to reduce the white noise in the quantized domain to a special `0-1` noise, under which an $\widehat{x}$ is either left unchanged or shifted to $\widehat{x} + 1$, as observed in [4]. Nevertheless, this strategy may give lower left-over entropy.

Handling replacement noise. After the white noise has been corrected, an existing scheme for set difference can be applied in the quantized domain to correct the replacement noise. There are known schemes that can achieve entropy loss of $O(t \log \lceil \frac{1}{\lambda} \rceil)$ with small leading constant, such as those in [7, 3]. Although the replacement noise is not considered for the face biometrics that we study in Section 6, it may need to be addressed for other biometric templates (e.g., iris patterns [9]).

Extension to higher dimensions. It is straightforward to extend our scheme to higher dimensions, where each $x \in X$ is a point in some d-dimensional space. For example, we can apply a scalar quantizer on each coordinate of every point, and let the distance of two points in d-dimensional space be measured by max-norm (i.e., the maximum distance in all dimensions). The entropy loss of the resulting scheme would be d times that in the current construction for 1-D points. If there is no replacement noise, we could also expand the n points in d-dimensional space into nd points in 1-D and apply the proposed scheme.

The choice of the sketch family. It is important to note that even if a quantization-based sketch family is well-formed, it does not guarantee the existence of a "good" quantizer in that family. Nevertheless, it does allow us to evaluate any given member in the family with respect to the "optimal" member in the family. We consider it a challenging open problem to find a general algorithm to find the optimal quantizer among all possible quantizers, given certain practical constraints (e.g., the smallest possible quantization step and the distribution of X).

6 A Concrete Construction for Face Biometrics

Face images, especially those taken from a controlled environment, can be used as the basis of identity verification, Here we follow the techniques employed in [17] and make use of the *singular value decomposition* (SVD) of the face images for verification, which is a well-known strategy in the face recognition literature (such as [10, 6]). Given a face image A of size $M \times N$, we can always find matrices U, Σ and V such that $A = U\Sigma V^T$, where Σ is an $M \times N$ matrix with $\min(M, N)$ non-zero elements ordered according to their significance. As noted in [17], some (say, n) most significant coefficients of Σ contain significant identity information of the individual. Typically n is chosen such that the sum of these n coefficients is more than, say, 98% of the sum of all the coefficients.

In [17], the biometric template of an individual is obtained as follows. First, we take a few face images, compute the SVD, and obtain the minimum min_i and maximum max_i of the i-th significant coefficient, for $1 \le i \le n$, where n is chosen to be 20. The mean value $a_i = (max_i + min_i)/2$ is then taken as a point in the template. When a new face image is presented for verification, its SVD is computed, and if for $1 \le i \le n$, the i-th significant coefficient is sufficiently close to a_i, it is considered as authenticated. The scheme in [17] is applied to face images from the Essex Faces94 Database [16], which contains 152 faces with 20 images for each face (24bit color JPEG). Twelve images per face are randomly chosen to compute the templates, and the rest 8 are used for testing. The experiments show that when the false accept rate is 0.005, the false reject rate is less than 0.045.

To apply our sketch scheme, for each coefficient, we further compute the minimum `min` and the maximum `max` of all the templates in the database (assuming that the number of templates is large). Hence, we can compute our biometric template X as a sequence of n points, where the i-th point $x_i = \frac{a_i - \mathtt{min}}{\mathtt{max} - \mathtt{min}}$. We set the noise level $\delta_i = \frac{k(max_i - a_i)}{\mathtt{max} - \mathtt{min}}$ for some constant $k \ge 1$. In this way, each point x_i will be between 0 and 1 so that our scheme can be applied. There is a difference, however, that we have a different δ_i for each point, which we have to put as part of the sketch. Nevertheless, our analysis on the entropy loss can be easily adapted to this case, and the difference here will not affect the security of the scheme. Here we choose $\lambda_i = \delta_i$ for all $1 \le i \le n$.

In this way, the sketch produced by our proposed scheme, would be the tuple

$$P = (\mathtt{min}, \mathtt{max}, \lambda_1, \cdots, \lambda_n, \widehat{x}_1 - \mathcal{C}_{\lambda_1}(\widehat{x}_1), \cdots, \widehat{x}_n - \mathcal{C}_{\lambda_n}(\widehat{x}_n))$$

where $\widehat{x}_i = \mathcal{Q}_{\lambda_i}(x_i)$ for $1 \le i \le n$. By applying the arguments in Theorem 6 and Corollary 7 to each point in X, we have

Corollary 8 *The entropy loss in the quantized domain for the aforementioned scheme is at most $n \log 3$. Let m be the left-over entropy. When $\lambda_i < \delta_i$ for any i, $1 \le i \le n$, let the left-over entropy be m'. We have $m' - m \le n \log 3$. If all points are uniformly distributed, we have $m' - m \le n \log(3/2)$.*

When $n = 20$, the above bounds are approximately 31.7 and 11.7 respectively.

References

1. X. Boyen, Y. Dodis, J. Katz, R. Ostrovsky, and A. Smith. Secure remote authentication using biometric data. In *Eurocrypt*, 2005.
2. Xavier Boyen. Reusable cryptographic fuzzy extractors. In *ACM CCS*, pages 82–91, Washington DC, USA, 2004. ACM Press.
3. Ee-Chien Chang, Vadym Fedyukovych, and Qiming Li. Secure sketch for multi-set difference. Cryptology ePrint Archive, Report 2006/090, 2006. `http://eprint.iacr.org/`.
4. Ee-Chien Chang and Qiming Li. Small secure sketch for point-set difference. Cryptology ePrint Archive, Report 2005/145, 2005. `http://eprint.iacr.org/`.
5. Ee-Chien Chang and Qiming Li. Hiding secret points amidst chaff. In *Eurocrypt*, volume 4004 of *LNCS*, pages 59–72, 2006.
6. Yong-Qing Cheng. Human face recognition method based on the statistical model of small sample size. In *SPIE Proc. Intell. Robot and Compu. Vision*, pages 85–95, 1991.
7. Yevgeniy Dodis, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *Eurocrypt*, volume 3027 of *LNCS*, pages 523–540. Springer-Verlag, 2004.
8. F. Hao and C.W. Chan. Private key generation from on-line handwritten signatures. *Information Management and Computer Security*, 10(2), 2002.
9. Feng Hao, Ross Anderson, and John Daugman. Combining cryptography with biometrics effectively. Technical Report UCAM-CL-TR-640, University of Cambridge, 2005.
10. Z. Hong. Algebraic feature extraction of image for recognition. *Pattern Recognition*, 24:211–219, 1991.
11. Ari Juels and Madhu Sudan. A fuzzy vault scheme. In *IEEE Intl. Symp. on Information Theory*, 2002.
12. Ari Juels and Martin Wattenberg. A fuzzy commitment scheme. In *ACM CCS*, pages 28–36, 1999.
13. J.-P. Linnartz and P. Tuyls. New shielding functions to enhance privacy and prevent misuse of biometric templates. In *AVBPA*, pages 393–402, 2003.
14. Ueli Maurer and Stefan Wolf. Information-theoretic key agreement: From weak to strong secrecy for free. In *Eurocrypt*, 2000.
15. F. Monrose, M.K. Reiter, Q. Li, and S. Wetzel. Cryptographic key generation from voice. In *IEEE Symp. on Security and Privacy*, 2001.
16. Libor Spacek. The essex faces94 database. http://cswww.essex.ac.uk/mv/allfaces/.
17. Y. Sutcu, T. Sencar, and N. Memon. A secure biometric authentication scheme based on robust hashing. In *ACM MM-SEC Workshop*, 2005.
18. P. Tuyls, A.H.M. Akkermans, T.A.M. Kevenaar, G.J. Schrijen, A.M. Bazen, and R.N.J. Veldhuis. Practical biometric authentication with template protection. In *AVBPA*, pages 436–446, 2005.
19. P. Tuyls and J. Goseling. Capacity and examples of template-protecting biometric authentication systems. In *ECCV Workshop BioAW*, pages 158–170, 2004.
20. Shenglin Yang and Ingrid Verbauwhede. Automatic secure fingerprint verification system based on fuzzy vault scheme. In *IEEE Intl. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 609–612, 2005.