On the (Non-)Reusability of Fuzzy Sketches and Extractors and Security Improvements in the Computational Setting*

Marina Blanton and Mehrdad Aliasgari
Department of Computer Science and Engineering
University of Notre Dame
{mblanton,maliasga}@cse.nd.edu

Abstract

Secure sketches and fuzzy extractors enable the use of biometric data in cryptographic applications by correcting errors in noisy biometric readings and producing cryptographic materials suitable for authentication, encryption, and other purposes. Such constructions work by producing a public sketch, which is later used to reproduce the original biometric and all derived information exactly from a noisy biometric reading. It has been previously shown that release of multiple sketches associated with a single biometric presents security problems for certain constructions. We continue the analysis to demonstrate that all other constructions in the literature are also prone to similar problems and cannot be safely reused. To mitigate the problem, we propose for each user to store one short secret string for all possible uses of her biometric, and show that simple constructions in the computational setting have numerous advantageous security and usability properties under standard hardness assumptions. Our constructions are generic in that they can be used with any existing secure sketch as a black box.

1 Introduction

The motivation for this work comes from practical use of biometric-derived data. Biometrics and derivation of cryptographic material from biometric data for authentication, encryption, or other purposes is an active research area. Secure sketches and fuzzy extractors [11, 10] were introduced as mechanisms of deriving cryptographic material from noisy biometric data for the purpose of its use in cryptographic applications. Such constructions allow one to produce a helper string (secure sketch) – which is viewed as public – from a biometric and later re-produce the cryptographic string exactly from a close noisy biometric reading using the helper string. The goal of such constructions is to keep the biometric itself hidden, which means that information leakage due to the release of the helper string must be minimized.

While this is a powerful concept that enables new applications and can also be attractive to users who no longer need to maintain secrets to participate in cryptographic protocols, it has been shown that leakage of information associated with the biometric in such constructions is unavoidable [27, 12]. Furthermore, this concept was initially proposed and primarily studied in the context when the construction is applied to a biometric only once. Consecutive publications [3, 26] explored the security guarantees of such schemes in terms of their reusability, when a single biometric or its noisy version is used to produce multiple secure sketches using the same or different algorithms. Information leakage prevents such constructions from meeting standard security requirements sought

*Portions of this work were supported by grant AFOSR-FA9550-09-1-0223 from the Air Force Office of Scientific Research.

of them in cryptographic applications such as indistinguishability (or inability to link two records to the same biometric) and irreversibility (inability to reverse the construction and directly recover information about the biometric). Some of the more popular constructions for secure sketches (namely, based on the code offset and permutation groups) has been shown to have serious security weaknesses with respect to their reusability in presence of even very weak adversaries [26]. In this work, we continue this analysis on a number of other constructions from the literature and show that they also cannot be safely reused. In particular, we show that they fail to satisfy standard security expectations with respect to reusability and therefore cannot be used in security applications.

Information leakage is generally difficult to quantify. In constructions that work with biometric data, theoretical analysis is expressed in terms of entropy loss associated with the release of the helper string and therefore is only a rough upper bound rather than a precise estimate. Also, for the current error rates in biometric data and their typical set of parameters, the information theoretic analysis provides bounds that result in leakage of most or even all entropy contained in a biometric (see [2] for an example set of parameters for iris codes and entropy loss due to the release of helper data). This presents problems even in presence of very weak adversaries.

To overcome the issues of information leakage and unsafe reuse of such constructions, we propose to use the computational setting, where a user stores a single short key (regardless of the number of applications of secure sketches or fuzzy extractors to her biometric) and the adversary is computationally bounded. Note that the key is introduced for the purpose of avoiding information leakage and improving the security of the schemes and does not change the functionality. We believe that keeping a single short key for all possible uses of biometric-based material in different security applications is a small price to pay for achieving significant security improvements (which otherwise are not possible) and the ability to safely use such constructions. We show that the use of one key and standard computational assumptions (namely, existences of pseudo-random functions and hash functions) is sufficient to achieve solutions with very attractive properties using simple schemes. Our constructions are generic in that they can be applied to any existing secure sketch scheme in a black box manner for any type of biometric (or distance metric) to produce a secure sketch or a fuzzy extractor with improved properties.

We would like to emphasize that the use of the secret in our schemes should not be confused with so-called multi-factor authentication or the use of a shared secret. There are two fundamental differences between such approaches and our work: (i) in our schemes the secret is not shared, neither the secret itself nor any function of it need to be known to any party, and (ii) a single secret is sufficient for all possible uses of fuzzy sketches and extractors including multiple biometric types, multiple applications, and multiple servers.

The security benefits of our schemes are:

- Leakage of information about a biometric is unavoidable in the information theoretic setting [27, 12], while our solution results in provably no information leakage.
- Previously, only certain restricted types of error-correcting codes could be used to ensure security of fuzzy sketches and extractors [3]. Our solution lifts such restrictions and can be used with any type of error-correcting code.
- Prior analysis of secure sketch constructions [26] showed that they fail to achieve standard security requirements for cryptographic applications. We show that other prior constructions are also susceptible to that problem, while our solution is secure in a much stronger adversarial model.
- Previously exposure of a key derived from a biometric was shown to provably reveal no information about the biometric for a specific construction in the random oracle model [3]. Our construction, on the other hand, achieves this result in the standard model using any existing secure sketch.

In our analysis of existing constructions, we use a very weak adversary. The security of our own schemes, on the other hand, is shown using a very strong adversary (the strongest in the literature for the same problem setup).

To summarize, our contributions are two-fold: (i) new analysis of fuzzy sketch schemes that shows that even a weak adversary has a significant advantage in compromising security of existing constructions, and (ii) simple schemes that use a single secret to achieve strong security under standard assumptions.

2 Model and Definitions

2.1 Fuzzy sketches and extractors

Secure (or fuzzy) sketches were introduced by Dodis at el. [11, 10] as a mechanism of correcting errors in noisy secrets (e.g., biometrics) by releasing a helper string S that does not reveal a lot of information about the secret. Let W denote a random variable and w its value.

Definition 1 *A $(\mathcal{M}, m, m', t)$-secure sketch is a pair of randomized algorithms:*

- SS *is a function that, on input w from metric space $\mathcal{M}$ with distance function* dist*, outputs a sketch S.*
- Rec *is a function that, on input $w' \in \mathcal{M}$ and $S = \mathsf{SS}(w)$, recovers and outputs the original w if $\mathsf{dist}(w, w') \leq t$.*

Secure sketches have been constructed for different types of metric spaces $\mathcal{M}$, for which the distance function $\mathsf{dist}(a, b)$ is defined for all $a, b \in \mathcal{M}$. Security of a secure sketch is evaluated in terms of entropy of W before and after releasing the string S, i.e., the entropy loss associated with making S public. The *min-entropy* (or "worst-case" entropy) of W is $H_\infty(W) = -\log \max_w \Pr[W = w]$ and the *average min-entropy* of W given S is $\bar{H}_\infty(W|S) = \log \mathbb{E}_{s \leftarrow S}[\max_w \Pr[W = w|S = s]]$. Then for any W with $H_\infty \geq m$ the probability of guessing W after observing S is at most $1/2^{-m'}$ where $m' \leq \bar{H}_\infty(W|S)$, i.e., the entropy loss due to release of S is $m - m'$ (and is unavoidable).

Fuzzy extractors allow one to extract randomness from w (to use it as cryptographic material) and later reproduce it exactly using w' close to the original w.

Definition 2 *A $(\mathcal{M}, m, m', t, \epsilon)$-fuzzy extractor is a pair of algorithms:*

- Gen *is a function that, on input $w \in \mathcal{M}$, outputs extracted random string R and a helper string P.*
- Rep *is a function that, on input w' and P reproduces and outputs R that was generated using $\mathsf{Gen}(w)$ if $\mathsf{dist}(w, w') \leq t$.*

The security requirement is such that, for any W of min-entropy m, the *statistical distance* between the distribution of R and the uniform distribution of strings of the same length is no greater than ϵ, even after observing the helper data P. Note that the statistical distance between probability distributions X and Y is defined as $\mathrm{SD}(X, Y) = \frac{1}{2} \sum_a |\Pr(X = a) - \Pr(Y = a)|$.

A fuzzy extractor can be built from a secure sketch using the following generic construction given in [11]:

$\mathsf{Gen}(w)$:

1. Execute $S \leftarrow \mathsf{SS}(w; r_1)$, where r_1 explicitly denotes random coins used by SS (if any).
2. Use a strong extractor Ext to extract a random string R from w, i.e., $R \leftarrow \mathsf{Ext}(w; r_2)$, where r_2 denotes random coins used by Ext.
3. Output public $P = (S, r_2)$ and secret R.

$\mathsf{Rep}(w', P = (S, r_2))$

1. Execute $w \leftarrow \mathsf{Rec}(w', S)$. If Rec fails (i.e., when $\mathsf{dist}(w, w') > t$ such that $S = \mathsf{SS}(w)$), stop.
2. Extract R from w using r_2 as $R \leftarrow \mathsf{Ext}(w, r_2)$ and output R.

Strong extractors [20, 19] have been well studied and can extract at most $m - 2\log(\frac{1}{\epsilon}) + O(1)$ nearly random bits (where m is min-entropy of W and ϵ is the security parameter defined above). One such construction uses universal hash function and extracts $m - 2\log(\frac{1}{\epsilon}) + 2$ random bits. Thus, entropy loss of $2\log(\frac{1}{\epsilon}) + 2$ is in addition to the entropy loss due to the release of a sketch S. If a strong extractor is modeled as a random oracle, there is no additional entropy loss.

Many constructions utilize error-correcting codes. A code C is a subset of K elements $\{w_0, \ldots, w_{K-1}\}$ of $\mathcal{M}$. The minimum distance of C is the smallest d such that $\mathsf{dist}(w_i, w_j) \geq d$ for all $i \neq j$, which implies that the code can detect up to $d-1$ errors; and the error-correcting distance is $t = \lfloor (d-1)/2 \rfloor$.

A linear error-correcting code C over field $\mathbb{F}_q$ is a k-dimensional linear subspace of the vector space $\mathbb{F}_q^n$ which uses Hamming distance as the metric, and is denoted as $(n, k, t)_{\mathbb{F}_q}$-code. For any linear code C, an $(n-k) \times n$ parity-check matrix H projects any vector $v \in \mathcal{F}_q^n$ to the space orthogonal to C. This projection is called the syndrome and denoted by $\mathsf{syn}(v) = Hv$. Then $v \in C$ iff $\mathsf{syn}(v) = 0$. The syndrome contains all information necessary for decoding. That is, when codeword c is transmitted and noisy $w = c + e$ is received, $\mathsf{syn}(w) = \mathsf{syn}(c) + \mathsf{syn}(e) = 0 + \mathsf{syn}(e)$, where $\mathsf{syn}(e)$ can be used to determine the error pattern e.

Secure sketch constructions for the Hamming distance (e.g., the code-offset construction) have been most heavily analyzed. Also, the permutation-based construction, which is applicable to any transitive metric, has been sufficiently analyzed in [26, 3]. For that reason, in this work we concentrate on constructions specific to other distance metrics, specifically the set difference and the edit distance. Recall that while the Hamming distance is used for biometric data such as iris codes, the set difference is employed for fingerprints and the edit distance is relevant to DNA comparisons.

2.2 Constructions for set difference

Throughout this work, we use notation $a \xleftarrow{R} A$ to denote that the value of a is chosen uniformly at random from the set A.

Fuzzy vault. The fuzzy vault scheme designed by Juels and Sudan [13] can be used as a fuzzy sketch when the biometric data is comprised of unordered elements $w = \{w_1, \ldots, w_s\}$ (e.g., minutiae points in fingerprints). The main idea is to disguise the points in w by adding a large number of *chaff points*. The genuine points then carry information that allows w to be reconstructed from its noisy version w'. In what follows, we assume that $t \in [1, s]$, and $r \in [s+1, n]$, where n is the set of all possible points or the universe, are system-wide parameters. Work is over field $\mathbb{F}_n$, where n is a prime power.

To compute $\mathsf{SS}(w)$:

1. Choose a random polynomial $p(\cdot)$ of degree at most $s - t - 1$ over $\mathbb{F}_n$.
2. For each $w_i \in w$, let $x_i = w_i$ and $y_i = p(x_i)$.
3. Choose $r - s$ distinct points $x_{s+1}, \ldots, x_r$ at random from $\mathbb{F}_n \setminus w$ and set $y_i \xleftarrow{R} \mathbb{F}_n \setminus \{p(x_i)\}$ for $i = s+1, \ldots, r$.
4. Output $\mathsf{SS}(w) = \{(x_1, y_1), \ldots, (x_s, y_s)\}$ sorted by the value of x_i's.

To compute $\mathsf{Rec}(w', S)$:

1. Create the set D of pairs (x_i, y_i) such that $x_i \in w'$.
2. Run Reed-Solomon decoding on D to recover the polynomial $p(\cdot)$.

3. Output s points of the form $(x_i, p(x_i))$ from S.

Privacy of the biometric depends on the number and distribution of points in S (i.e., the difficulty of identifying the original points and the number of spurious polynomials created by the chaff points). The entropy loss due to the release of S is determined in [11] to be upper bounded by $t \log n + \log \binom{n}{r} - \log \binom{n-s}{r-s} + 2$.

Improved fuzzy vault. Dodis et al. [10] observed that the polynomial in the above construction does not need to be random, which allows for a secure sketch with significantly lower entropy loss, namely $t \log n$.

To compute $\mathsf{SS}(w)$:

1. Compute unique monic polynomial $p(x) = \prod_{w_i \in w}(x - w_i)$ of degree s.
2. Output the coefficients of $p()$ of degree $s-1$ down to $s-t$, which will form $\mathsf{SS}(w) = (c_{s-1}, \ldots, c_{s-t})$.

To compute $\mathsf{Rec}(w', S = (c_{s-1}, \ldots, c_{s-t}))$:

1. Create a new polynomial p_{high} of degree s that shares the top $t+1$ coefficients with $p()$, i.e., $p_{\text{high}}(x) = x^s + \sum_{i=s-t}^{s-1} c_i x^i$.
2. Evaluate p_{high} on points of w' to obtain pairs (a_1, b_1), $\ldots$, (a_s, b_s).
3. Use Reed-Solomon decoding to find a polynomial p_{low} of degree $s-t-1$ such that $p_{\text{low}}(a_i) = b_i$ for at least $s - t/2$ values of a_i's. If none can be found, output fail.
4. Output the roots of the polynomial $p_{\text{high}} - p_{\text{low}}$.

Pinsketch. This next construction works when the universe size n is large (or could not be enumerated) and thus all computation is polynomial in $\log n$. Pinsketch also allows the biometric w to have a variable number of points, which makes the construction particularly attractive. In what follows, support $\mathsf{supp}(w)$ is used as an alternative representation of small weight w by listing the positions at which it is non-zero. This allows decoding complexity to be a function of $\log n$ instead of n.

To compute $\mathsf{SS}(w) = \mathsf{syn}(x_w)$:

1. Let $s_j = \sum_{w_i \in w} (w_i)^j$ (in $\mathbb{F}_{2^m}$ where $n = 2^m - 1$).
2. Output $\mathsf{SS}(w) = s_1, s_3, \ldots, s_{2t-1}$.

To compute $\mathsf{Rec}(w', S = (s_1, s_3, \ldots, s_{2t-1}))$:

1. Compute $(s'_1, s'_3, \ldots, s'_{2t-1}) = \mathsf{SS}(w') = \mathsf{syn}(x_{w'})$.
2. Let $\sigma_i = s'_i - s_i$ and compute $\mathsf{supp}(v)$ such that $\mathsf{syn}(v) = (\sigma_1, \sigma_3, \ldots, \sigma_{2t-1})$ and $|\mathsf{supp}(v)| \leq t$.
3. If $\mathsf{dist}(w, w') \leq t$, then $\mathsf{supp}(v) = w \Delta w'$; therefore, output $w = w' \Delta \mathsf{supp}(v)$.

Here Δ denotes symmetric difference, i.e., $\mathsf{dist}(w, w') = w \Delta w'$. This construction uses BCH codes (which are linear) and results in entropy loss of $t \log(n+1)$.

2.3 Constructions for edit distance

To the best of our knowledge, the only known constructions for the edit distance first use an embedding of the edit distance metric into a transitive metric (e.g., the Hamming distance) of larger dimension and apply a secure sketch construction to the target metric. A construction for the edit distance then can proceed as follows:

To compute $\mathsf{SS}(w)$:

1. Embed w into v in a transitive metric space (the Hamming distance or set difference) as $v = f(w)$.

2. Compute and output $\mathsf{SS}(v) = \mathsf{syn}(v)$.

To compute $\mathsf{Rec}(w', S)$:

1. Embed w' into v' in a transitive metric space.
2. Execute $v = \mathsf{Rec}(v', S)$ and output the reverse embedding of v $w = f^{-1}(v)$.

The entropy loss of this construction depends on the properties of the embedding and the secure sketch scheme in the target metric. When f^{-1} is not efficiently computable, the output of an additional function $g(w)$ can be stored in S which helps in recovering w from v.

2.4 Security notions

The original security definitions of fuzzy sketches and extractors (i.e., quantifying information leakage about the biometric due to the release of public helper data and ensuring that the output of fuzzy extractors is indistinguishable from random) were formulated for a single instance of a fuzzy sketch or extractor in isolation [11]. Consecutive literature [3, 26] considered a stronger (and more realistic) adversarial model where such constructions can be invoked multiple times and therefore the security guarantees must hold when the constructions are reused. Furthermore, the power granted to the adversary can greatly differ. In this work we use weak adversaries while analyzing existing constructions (to show that they do not provide sufficient security guarantees even in presence of weak adversaries) and strong adversaries when proving the security of our proposed solution. In a nutshell, a weak adversary is given two fuzzy sketches and tries to determine whether they were produced using the same biometric and what that biometric was, while a strong adversary can adaptively ask for fuzzy sketches and private key that fuzzy extractors output.

Let t be the maximum amount of errors that the biometric system can tolerate. We define Δ_t to be the set of all perturbation functions that represent differences in sampling biometric data; we then have $\Delta_t = \{\delta : \mathcal{M} \to \mathcal{M} \text{ such that } \mathsf{dist}(w, \delta(w)) \leq t\}$. In what follows, we first define a security game for weak adversaries with access to public sketches and then proceed with security games for strong adversaries. Two security properties for weak adversaries were defined in [26]: sketch indistinguishability and irreversibility.

2-Indistinguishability game ([26]):

1. The challenger chooses a random variable $W \in \mathcal{M}$ and samples it to obtain $w \in \mathcal{M}$. The challenger computes $S_1 = \mathsf{SS}(w)$ and gives S_1 to $\mathcal{A}$.
2. The challenger chooses a bit b at random. If $b = 1$, the challenger chooses $\delta \xleftarrow{R} \Delta_t$, and produces a related biometric $w' = \delta(w)$. Otherwise, if $b = 0$, the challenger samples W to obtain a different biometric w'. The challenger than computes $S_2 = \mathsf{SS}(w')$ and gives S_2 to $\mathcal{A}$.
3. The adversary eventually produces a bit b' and wins if $b' = b$.

The adversary $\mathcal{A}$'s advantage in the above game is defined as:

$$\mathsf{Adv}^{\mathsf{ind}}_{\mathcal{A}} = 2\left|\Pr[b' = b] - \frac{1}{2}\right| = 2\left|\Pr[b' \neq b] - \frac{1}{2}\right|$$

Definition 3 *An $(\mathcal{M}, m, m', t)$-secure fuzzy sketch (SS, Rec) is ϵ-indistinguishable in Δ_t if for any adversary $\mathcal{A}$ it holds that $\mathsf{Adv}^{\mathsf{ind}}_{\mathcal{A}} \leq \epsilon$, and the fuzzy sketch is reusable when ϵ is negligible.*

The irreversibility property of fuzzy sketch constructions means that an adversary who obtains access to multiple sketches which have been generated from the same noisy input possibly using different sketching functions is unable to recover the original input.

In the irreversibility game below, the adversary obtains access to multiple sketches which were generated from the same noisy input, but possibly using different sketching functions. The adversary's goal is then to recover the original input. For the purposes of this game, we define a family $\mathcal{F} = \{(\mathsf{SS}_i, \mathsf{Rec}_i)\}$ of $(\mathcal{M}, m, m'_i, t_i)$-secure fuzzy sketches.

Irreversibility game ([26]):

1. The challenger chooses a random variable $W \in \mathcal{M}$ and samples it to obtain $w \in \mathcal{M}$. The challenger chooses $(\mathsf{SS}_{i_1}, \mathsf{Rec}_{i_1})$ at random from $\mathcal{F}$, computes $S_1 = \mathsf{SS}_{i_1}(w)$ and gives S_1 to $\mathcal{A}$.
2. The challenger chooses $\delta \xleftarrow{R} \Delta_t$, where $t = \min\{t_i\}$, and $(\mathsf{SS}_{i_2}, \mathsf{Rec}_{i_2})$ at random from $F \setminus \{(\mathsf{SS}_{i_1}, \mathsf{Rec}_{i_1})\}$. The challenger produces a related biometric $w' = \delta(w)$, computes $S_2 = \mathsf{SS}(w')$ and gives S_2 to $\mathcal{A}$.
3. The adversary eventually produces output $\hat{w} \in \mathcal{M}$ and wins if $\hat{w} = w$.

Note that requiring the adversary to produce w is equivalent to requiring it to produce w', since knowledge of one of them is equivalent to knowledge of both in presence of sketches S_1 and S_2.

The adversary $\mathcal{A}$'s advantage in the above game is defined as:

$$\mathsf{Adv}^{\mathsf{irrev}}_{\mathcal{A}} = \frac{2^{\min(m'_{i_1}, m'_{i_2})}}{2^{\min(m'_{i_1}, m'_{i_2})} - 1} \left| \Pr[\hat{w} = w] - \frac{1}{2^{\min(m'_{i_1}, m'_{i_2})}} \right|.$$

Definition 4 *A family $\mathcal{F}$ of $(\mathcal{M}, m, m'_i, t_i)$-secure fuzzy sketches $\{(\mathsf{SS}_i, \mathsf{Rec}_i)\}$ is ϵ-irreversible in Δ_t if for any adversary $\mathcal{A}$ it holds that $\mathsf{Adv}^{\mathsf{irrev}}_{\mathcal{A}} \leq \epsilon$, and the family is just irreversible when ϵ is negligible.*

We now proceed with defining security games for more powerful adversaries. We term the next two definitions as *weak biometric privacy* and *strong biometric privacy*, respectively. In both of them the adversary will be allowed to query the scheme a large number of times. The difference between the definitions is that in the first the adversary obtains access only to the public information, while in the second it also obtains access to the key output by a fuzzy extractor. Thus, we use the first definition for secure sketches and the second one for fuzzy extractors.

The two security games below are roughly equivalent to outsider and insider chosen perturbation security, respectively, in [3], but are stronger than the respective definitions in [3]. In particular, in our definition of weak biometric security we require the adversary to only distinguish between two sketches, while the adversary was required to recover the biometric w in [3] in the corresponding definition. Furthermore, instead of allowing the adversary to query fuzzy sketches for a particular biometric w and then challenging the adversary by asking it to distinguish between a sketch for w and a sketch for a randomly chosen biometric, we setup two biometrics w_0 and w_1 and allow the adversary to query sketches for both. Then during the challenge, the adversary is asked to determine which biometric was used in producing the challenge sketch. This can potentially give the adversary advantage over the prior formulation of the game, especially in the computational setting where different users will possess different key material.

As we are now working in the computational setting, we use κ to denote the security parameter. All algorithms are assumed to be polynomial time in κ. We say that a $\epsilon(\kappa)$ function is negligible if for all positive polynomials $p(\cdot)$ and sufficiently large κ it holds that $\epsilon(\kappa) < 1/p(\kappa)$.

Weak biometric privacy:

1. (Preparation) The adversary chooses a random variable $W \in \mathcal{M}$ and sends its specification to the challenger.

2. (Sampling) The challenger randomly samples W to obtain $w_0 \in \mathcal{M}$ and $w_1 \in \mathcal{M}$ and initializes two users $\mathcal{U}_0$ and $\mathcal{U}_1$, respectively, using that information.
3. (Queries) The adversary makes up to q possibly adaptive sketching queries as follows: to form query i, the adversary chooses $\delta_i \in \Delta_t$ and sends it and a bit b_i to the challenger. The challenger computes $S_i \leftarrow \mathsf{SS}(\delta_i(w_{b_i}); r_i)$ using fresh randomness r_i and returns S_i to $\mathcal{A}$.
4. (Challenge) The challenger chooses a bit $b \xleftarrow{R} \{0,1\}$ and $\delta \xleftarrow{R} \Delta_t$, and produces a biometric $w' = \delta(w_b)$. The challenger then computes $S \leftarrow \mathsf{SS}(w'; r)$ using fresh randomness r and gives S to $\mathcal{A}$.
5. (More queries) The adversary $\mathcal{A}$ can run more queries up to the bound q as specified in step 3.
6. (Response) The adversary eventually produces a bit b' and wins if $b' = b$.

We define the adversary $\mathcal{A}$'s advantage in this game as:

$$\mathsf{Adv}^{\mathsf{wbp}}_{\mathcal{A}}(\kappa) = 2\left|\Pr[b' = b] - \frac{1}{2}\right| = 2\left|\Pr[b' \neq b] - \frac{1}{2}\right|$$

Definition 5 *We say that an $(\mathcal{M}, m, m', t)$-secure fuzzy sketch (SS, Rec) has weak biometric privacy if for any probabilistic polynomial-time (PPT) adversary $\mathcal{A}$ it holds that $\mathsf{Adv}^{\mathsf{wbp}}_{\mathcal{A}}(\kappa) \leq \epsilon(\kappa)$ for a negligibly small $\epsilon(\kappa)$.*

Note that unlike previous definitions, we explicitly specify the security parameter κ and define the adversary's advantage as a function of this parameter. As mentioned earlier, our construction considers the computational setting (and thus computationally limited adversaries), where the complexity of all algorithms will be specified as a function of κ.

In what follows, let Δ denote all perturbation functions over metric space $\mathcal{M}$, i.e., $\Delta = \{\delta : \mathcal{M} \to \mathcal{M}\}$ where $\mathsf{dist}(w, \delta(w))$ can be greater than threshold t.

The next definition corresponds to the strongest version of the insider chosen perturbation security definition in [3]. The adversary can query the challenger to obtain sketches on both related and unrelated biometrics and private key corresponding to unrelated biometrics. Note that this time we ask the adversary to distinguish between the secret key output by a fuzzy extractor on the related biometric and a randomly chosen string. The reason why we do not ask the adversary to distinguish between secret keys of two users is because the adversary has the choice of the sketch that it can use in the challenge. This means that the adversary will trivially know for which user the secret key will be produced. We, however, note that in order to distinguish secret keys corresponding to two users, the adversary need to be able to distinguish at least one of them from a random string. Thus, our definition of security will imply the security in the game with two users.

Strong biometric privacy:

1. (Preparation) The adversary chooses $W \in \mathcal{M}$ and gives its specification to the challenger.
2. (Sampling) The challenger randomly samples W to obtain $w \in \mathcal{M}$.
3. (Public queries) The adversary makes up to q possibly adaptive generation queries as follows: to form query i, the adversary chooses $\delta_i \in \Delta$ and sends it to the challenger. The challenger computes $(P_i, R_i) \leftarrow \mathsf{Gen}(\delta_i(w); r_i)$ using fresh randomness r_i and returns public P_i only to $\mathcal{A}$.
4. (Private queries) The adversary makes up to q' possibly adaptive reproduction queries that can be intersperse with public queries as follows: to form query i, the adversary chooses $\delta'_i \in \Delta$ and a public data P'_i and sends them to the challenger. The challenger computes $R'_i \leftarrow \mathsf{Rep}(\delta'_i(w); P'_i)$ and returns R'_i to $\mathcal{A}$.

5. (Challenge) The adversary chooses string $P^* \in \{P_1, \ldots, P_q\}$ from one of the strings returned by the challenger in a public query such that P^* was produced using a public query δ_i with $\mathsf{dist}(w, \delta_i(w)) \leq t$ and in any private query (δ'_i, P^*) the distance $\mathsf{dist}(w, \delta'_i(w)) > t$. $\mathcal{A}$ sends P^* to the challenger. The challenger chooses a bit $b \xleftarrow{R} \{0,1\}$. If $b = 1$, the challenger computes the corresponding private string $R \leftarrow \mathsf{Rep}(w, P^*)$ and gives it to the adversary. Otherwise, if $b = 0$, it chooses a random string of the same length and gives it to $\mathcal{A}$ instead.
6. (More queries) The adversary $\mathcal{A}$ can run additional queries as specified in steps 3 and 4 (up to q and q' queries, respectively) with the exception that any query (δ, P^*) such that $\mathsf{dist}(w, \delta(w)) \leq t$ is not allowed.
7. (Response) The adversary eventually produces a bit b' and wins if $b' = b$.

We define the adversary $\mathcal{A}$'s advantage in this game as:

$$\mathsf{Adv}^{\mathsf{sbp}}_{\mathcal{A}}(\kappa) = 2\left|\Pr[b' = b] - \frac{1}{2}\right| = 2\left|\Pr[b' \neq b] - \frac{1}{2}\right|$$

Definition 6 *We say that an $(\mathcal{M}, m, m', t, \epsilon)$-secure fuzzy extractor (Gen, Rep) has strong biometric privacy if for any PPT adversary $\mathcal{A}$ it holds that $\mathsf{Adv}^{\mathsf{sbp}}_{\mathcal{A}}(\kappa) \leq \epsilon(\kappa)$ for a negligibly small $\epsilon(\kappa)$.*

2.5 Known Privacy Weaknesses

Simoens et al. [26] show that two popular secure sketch constructions – namely, the code offset construction with a linear error-correcting code (i.e., the syndrome construction) and the construction based on permutation groups – do not withstand the requirements of indistinguishability and irreversibility, i.e., the adversary can win such experiments with overwhelming probability. The former construction is applicable to the Hamming distance metric (and is among the most popular and widely studied schemes) and the latter can be used for any transitive distance metric. For that reason, in this work we concentrate on the analysis of schemes for other distance metrics (namely, set difference and edit distance), some of which are related to the previously analyzed constructions.

3 Analysis of Existing Schemes

3.1 Constructions for set difference

3.1.1 Fuzzy vault

Attacking indistinguishability. Before proceeding with the analysis, we note that the basic idea for the strategy in attacking the fuzzy vault scheme when two or more sketches are available – computing the intersection of the points – is straightforward and is not new. This attack appeared in [23, 15, 22]. Our analysis is different from what has been done before because all previous publications assume that given sketches are related and proceed with identifying original points. Our work, however, assumes a significantly weaker (and perhaps more realistic) adversary that would like to determine if two given sketches are related or not, which is a much more difficult task. Therefore, we present a rigorous new analysis that shows weaknesses of the scheme even in the presence of the weakest adversary.

The adversary receives two secure sketches $S_1 = \{(x_1, y_1), \ldots, (x_r, y_r)\}$ and $S_2 = \{(x'_1, y'_1), \ldots, (x'_r, y'_r)\}$, and the adversary's goal is to determine the coin flip, i.e., whether the biometrics w and w' are related or not. Let $\pi_x(S_i)$ denote projection of S_i onto the x-coordinate, i.e., $\pi_x(S_1) = \{x_1, \ldots, x_r\}$ and $\pi_x(S_2) = \{x'_1, \ldots, x'_r\}$. The idea behind the attack strategy here is to compute

the intersection of $\pi_x(S_1)$ and $\pi_x(S_2)$ and use its size to make a distinction between related and unrelated biometrics. Related sketches will overlap in at least $s-t$ original biometric points, while unrelated sketches will have fewer original biometric points overlap. In addition, a number of chaff points in $\pi_x(S_1)$ can collide with chaff points in $\pi_x(S_2)$ or points in $w' \setminus (w \cap w')$ (similarly, points from $w \setminus (w \cap w')$ can collide with chaff points in $\pi_x(S_2)$). Thus, the size of $\pi_x(S_1) \cap \pi_x(S_2)$ follows a certain distribution, but the expected overlap size is larger for related sketches. Before presenting the exact attack strategy, we analyze the properties of such a distribution.

Let $\alpha = |w \cap w'|$ denote the number of biometric points in the intersection, i.e., $\alpha \geq s - t/2$ for related biometric samples and $\alpha \leq s - t/2 - 1$ otherwise. Let $a = r - \alpha$ and $b = n - \alpha$, i.e., a is the number of sketch points that do not correspond to the overlapping biometric points and b is the overall space for such points. As customary in the literature, we assume that the biometric points of w are distributed uniformly at random in the space; the chaff points are also drawn uniformly at random from the remaining space. Then to determine how many points from $S_1' = \pi_x(S_1) \setminus (w \cap w')$ will collide with points from $S_2' = \pi_x(S_2) \setminus (w \cap w')$, suppose there are $b = n - \alpha$ bins and points from S_1' occupy $a = r - \alpha$ of them, i.e., there are a random bins with a ball in them. Then we throw another a balls (i.e., points from S_2') into the bins without replacement and count the number of bins with two balls in them (i.e., if a bin has two balls, it is removed, so that no bin has more than two balls; this is dictated by the requirement that all r points in a sketch are different). The above can be modeled as hypergeometric experiment.[1] Let X be a random variable that corresponds to the number of collisions in $\pi_x(S_1)$ and $\pi_x(S_2)$ (i.e, its size is $|(\pi_x(S_1) \cap \pi_x(S_2)) \setminus (w \cap w')|$). We obtain:

$$\Pr[X = k] = \frac{\binom{a}{k}\binom{b-a}{a-k}}{\binom{b}{a}} = \binom{a}{k} \frac{\prod_{i=0}^{k-1}(a-i)\prod_{i=0}^{a-k-1}(b-a-i)}{\prod_{i=0}^{a-1}(b-i)} \tag{1}$$

where X can take values between 0 and a. The mean value of this distribution is $E[X] = a \cdot \frac{a}{b}$.

This analysis leads to the following attack strategy: given sketches S_1 and S_2, $\mathcal{A}$ computes $\pi_x(S_1)$, $\pi_x(S_2)$, and $c = |\pi_x(S_1) \cap \pi_x(S_2)|$. Let β denote the value $(r - s + t/2)^2/(n - s + t/2)$ rounded to the nearest integer. If $c \geq (s - t/2 + \beta)$, output 1, otherwise, output 0.

Let α_{auth} denote a random variable corresponding to the distribution of $|w \cap w'|$ when w and w' are related (authentic), and α_{imp} denote a random variable corresponding to the size of such overlap when w and w' are unrelated (impostor). The adversary has the smallest probability of distinguishing between authentic and impostor sketches when the values of α_{auth} and α_{imp} are as close as possible, i.e, $\alpha_{\mathsf{auth}} = s - t/2$ and $\alpha_{\mathsf{imp}} = s - t/2 - 1$. According to the indistinguishability definition, we have $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{ind}} = 2\left|\Pr[b' = b] - \frac{1}{2}\right|$. If we let X_1 denote the random variable distributed according to the hypergeometric distribution above with $\alpha_1 = s - t/2$ and X_2 denote a similar random variable with $\alpha_2 = s - t/2 - 1$, we obtain that the adversary is successful with at least the

[1] When $a \ll b$, the requirement that we sample without replacement can be dropped and the result modeled as a simpler binomial experiment. In this applications, however, a in general is not guaranteed to be much smaller than b.

following probability:

$$\begin{aligned}
\Pr[b'=b] &= \Pr[b'=1\,|\,b=1]\Pr[b=1]+\Pr[b'=0\,|\,b=0]\Pr[b=0]\geq \qquad (2)\\
&\geq \frac{1}{2}\big(\Pr[X_1\geq c-\alpha_1]+\Pr[X_2<c-\alpha_2]\big)=\frac{1}{2}\big(\Pr[X_1\geq\beta]+\Pr[X_2<\beta+1]\big)=\\
&= \frac{1}{2}\Bigg(\sum_{i=\beta}^{r-s+t/2}\frac{\binom{r-s+t/2}{i}\binom{n-s+t/2-(r-s+t/2)}{r-s+t/2-i}}{\binom{n-s+t/2}{r-s+t/2}}+\sum_{i=0}^{\beta}\frac{\binom{r-s+t/2+1}{i}\binom{n-s+t/2+1-(r-s+t/2+1)}{r-s+t/2+1-i}}{\binom{n-s+t/2+1}{r-s+t/2+1}}\Bigg)=\\
&= \frac{1}{2}\Bigg(\sum_{i=\beta}^{r-s+t/2}\frac{\binom{r-s+t/2}{i}\binom{n-r}{r-s+t/2-i}}{\binom{n-s+t/2}{r-s+t/2}}+\sum_{i=0}^{\beta}\frac{\binom{r-s+t/2+1}{i}\binom{n-r}{r-s+t/2+1-i}}{\binom{n-s+t/2+1}{r-s+t/2+1}}\Bigg).
\end{aligned}$$

This probability and therefore $\mathsf{Adv}^{\mathsf{ind}}_{\mathcal{A}}$ can be easily computed for a given set of parameters n, r, s, and t. In reality, each parameter n, s, t, and r has limitations placed on it by the behavior of the actual biometric data. In particular, Clancy et al. [6] studies applicability of the fuzzy vault construction to fingerprint data and determines optimal parameters to use in order to achieve adequate resistance of the construction against brute force search (when an adversary is given a sketch and tries to determine sensitive information by searching through polynomials). While the fuzzy vault construction was not used exactly as a secure sketch in [6] and was generalized, we nevertheless obtain information about the parameters that would be used for fingerprint data. The field $\mathbb{F}_{p^2}$, for prime p, is used for representing fingerprint features in 2-D and the value of p is set to 251 giving us $n = 251^2 = 63001$ (this value of p also provides many choices for the decoding algorithm). The number of biometric points in a fingerprint is often in the range 20–50 (this value can greatly vary based on the equipment used and quality of data) Using the analysis and empirical data from [6, 21] as guidelines for achieving good distinguishing capability, low error rate, and difficulty of brute force attack on the fuzzy vault scheme, we set $s = 30$ and $t = 16$. Finally, the value of r is constrained in that the complexity of decoding for legitimate users can grow as r increases (this is caused by spurious polynomials introduced by the chaff points). In particular, at the decoding time, when a legitimate user computes $w' \cap \pi_x(S)$, where $S = \mathsf{SS}(w)$, the decoding complexity can grow when points from $w' \setminus (w' \cap w)$ coincide with chaff points in S. Since $|w' \setminus (w' \cap w)| \leq t/2$ for legitimate users, the experiment now consists of throwing $t/2$ points in $b = n - s + t/2$ bins, where $a = r - s + t/2$ bins already have a ball in them. We then want r to be such that the expected (integer-valued) number of collisions $\frac{t}{2} \cdot \frac{a}{b}$ to be zero.

Figure 1 plots the adversary's advantage $\mathsf{Adv}^{\mathsf{ind}}_{\mathcal{A}}$ for the above parameters as a function of r near the suggested in [6] value of r of about 300. As evident from the figure, the advantage is significant even in the worst (for the adversary) case when only one extra overlapping point separates authentic data from impostor. The jumps in the plot correspond to the places where the (integer-valued) mean of the distribution, $E[X]$, increases by 1.

In practice, the above model is more complex due to the need of quantizing the data and the ability to handle white noise (small differences in the positions of the feature points). As a result, this imposes a maximum packing density of points in the vault. In particular, points normally cannot be placed very close to other points, but given the acceptable distance d, they can lie anywhere in the space as long as they are at least distance d from other points. This means that the total number of elements r in a sketch cannot exceed $\frac{4\rho p^2}{d^2\pi}$ with packing density ρ (i.e., packing non-overlapping circles of radius $d/2$). The optimal density for packing circles is unachievable and instead [6] gives $\rho \approx 0.45$ in which case the packing is guaranteed to be random. When, for instance, $p = 251$, $d = 10$, and $\rho = 0.45$, $r \leq 361$.

Relating this point placement constraint to our problem of estimating the adversary's advantage in distinguishing related and unrelated sketches, we as before assume that α genuine points overlap.

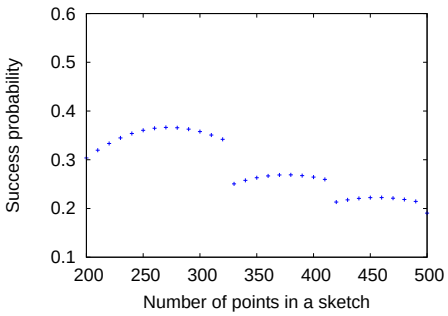

Figure 1: Adversary advantage $\mathsf{Adv}^{\mathsf{ind}}_{\mathcal{A}}$ with parameters $n = 251^2$, $s = 30$, $t = 16$, and varying r.

While a point can overlap with more than one other points, it is considered a match only with one of them (often the closest), and the remaining points can still overlap (and be considered a match) with other points. When packing non-overlapping balls in the area, the radius of each ball is $d/2$. When, however, computing the useful area, where points can be placed given a number of points already in the area, we have that a point cannot land within the distance d of another point and for that reason we model existing points as balls of radius $d' = d - \delta$ for some negligibly small δ.

The above means that two points within (Euclidean) distance of less than d from each other will be considered a match even if their positions differ. We can then take this modification into account and recompute the adversary's advantage assuming that each point occupies an area of radius $d' = d - \delta < d$. Then to determine the number of collisions caused by the chaff points (as in equation 1), we need to compute the space in which such points land. Given $\alpha = |w \cap w'|$ overlapping biometric points, they occupy an area between $\alpha \cdot \mathsf{CA}(d')$ and $\alpha \cdot 2\mathsf{CA}(d')$, where $\mathsf{CA}(r) = \pi r^2$ represents the circle area of radius r. Fortunately, we can estimate such an area more precisely as follows. Because the points are always placed at discrete locations, for any particular value of d, it is not difficult to compute the average area occupied by two overlapping balls of radius d'. Let $D_d = \{(i,j) \mid i = 0,\ldots,d-1,\ j = 1,\ldots,d-1 \text{ and } \sqrt{i^2+j^2} < d\}$. This set represents relative coordinates of a quarter of all points (excluding the perfect overlap) that would cause two points to overlap. Then the average area occupied by two overlapping balls of radius d' is:

$$\mathsf{AA}(d) = \Big(4 \sum_{(i,j)\in D_d} (2\mathsf{CA}(d') - \mathsf{IA}(d', \sqrt{i^2+j^2})) + \mathsf{CA}(d')\Big) / \left(4|D_d| + 1\right)$$

where $\mathsf{IA}(a,b)$ denotes the area of the intersection of two circles of radius a placed at distance b (the formula for which is well known). The average is computed using the distances for all overlapping points in four quarters and the perfect overlap with 0 distance.

Figure 2 plots the ratio of the average overlapping points area $\mathsf{AA}(d)$ to the area $2\mathsf{CA}(d')$ occupied by two circles of radius d' for different values of d. It is clear that the ratio remains constant, despite discretization of point locations, which could introduce an error in the computation. This means that we can approximate the area occupied by two overlapping points in our analysis by scaling $2\mathsf{CA}(r)$ by a constant factor.

Going back to modeling the distribution that would allow us to determine the adversary's advantage in distinguishing related and unrelated sketches, as before we assume that $\alpha = |w \cap w'|$ points from the original biometrics overlap. Let w_o denote the α points from w that are in the intersection, and similarly w'_o denote the corresponding α points from w'. The area occupied by the overlapping 2α points $w_o \cup w'_o$ on average is $\alpha\mathsf{AA}(d)$, while the α points from w_o and w'_o occupy $\alpha\mathsf{CA}(d')$ space each. This means that when we throw the remaining $r - \alpha$ points of w'

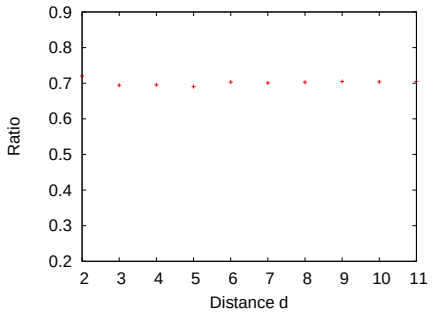

Figure 2: The ratio of the average area covered by two overlapping points to the area of two circles of diameter $d' \approx d$.

uniformly at random, the space at which they are to land is slightly different from the space where points in $\pi_x(S_1) \setminus w_o$ can reside. Let c be a constant such that $\mathsf{AA}(d) = c \cdot 2\mathsf{CA}(d')$ (i.e., $0.5 < c < 1$). Also let $\mathsf{space}(X)$ denote the space occupied by set X and $\mathsf{space}(all)$ denote the overall space of size $p \times p$. We obtain that $|\mathsf{space}(w_o)| = |\mathsf{space}(w'_o)| = \alpha\mathsf{CA}(d')$ and $|\mathsf{space}(w_o \cup w'_o)| = 2c\alpha\mathsf{CA}(d')$. Then when we throw the first point from $\pi_x(S_2) \setminus w'_o$ into $\mathsf{space}(all) \setminus \mathsf{space}(w_o)$, it has a regular probability of overlapping with a (random) point from $\pi_x(S_1) \setminus w_o$ if it lands in the space $\mathsf{space}(all) \setminus (\mathsf{space}(w_o \cup w'_o))$, and that probability is lowered if it lands in the space $\mathsf{space}(w_o \setminus w'_o)$. In the latter case, the landing point has fewer options for intersection (as a number of points in close proximity of it cannot contribute to an overlap) and the probability is lowered by a factor $2c - 1$ on average from the former case (the factor $2c - 1$ is the fraction of the thrown ball's area that can overlap with another point in w). Similar logic applies to other points that we throw with the difference that the areas occupied by existing points increase. Note that this model is still an approximation, as each ball can overlap with several balls near it. Since taking into account all possible overlaps will result in a significantly more complex model, we instead use this approximation as the lower bound on the attacker's advantage. That is, because we underestimate the remaining space at which we throw points, this results in greater probability of overlap of two balls, and the adversary's distinguishing probability is better when the number of spurious overlaps is small. Also note that the probability that a ball of radius d' overlaps with more than a single ball near it grows as the number of chaff points in a sketch increases.

To determine the number of overlapping points between $\pi_x(S_1)$ and $\pi_x(S_2)$, we might wish to rewrite equation 1 in a similar form. First observe that when we throw the first point after the α overlapping points from the original biometrics have been determined, the probability of overlap with the remaining $a = r - \alpha$ points from S_1 is

$$\frac{n - 2c\alpha\mathsf{CA}(d')}{n - \alpha\mathsf{CA}(d')} \cdot \frac{a\mathsf{CA}(d')}{n - \alpha\mathsf{CA}(d')} + \frac{(2c-1)\alpha\mathsf{CA}(d')}{n - \alpha\mathsf{CA}(d')} \cdot \frac{(2c-1)a\mathsf{CA}(d')}{n - \alpha\mathsf{CA}(d')} \tag{3}$$

Here the factors $\frac{n-2c\alpha\mathsf{CA}(d')}{n-\alpha\mathsf{CA}(d')}$ and $\frac{(2c-1)\alpha\mathsf{CA}(d')}{n-\alpha\mathsf{CA}(d')}$ correspond to the probabilities that the ball falls within the space $\mathsf{space}(all) \setminus (\mathsf{space}(w_o \cup w'_o))$ and $\mathsf{space}(w_o \setminus w'_o)$, respectively, from the total available space $n - \alpha\mathsf{CA}(d')$ (corresponding to $\mathsf{space}(all) \setminus \mathsf{space}(w_o)$). Then for the first case, the probability of overlap is $\frac{a\mathsf{CA}(d')}{n-\alpha\mathsf{CA}(d')}$, and it is lowered by a factor $2c-1$ in the second case as described above. Then when we throw the second point, the probability of overlap with the remaining points in S_1 (which do not already overlap with points from S_2) becomes $\frac{n-2c(\alpha+1)\mathsf{CA}(d')}{n-(\alpha+1)\mathsf{CA}(d')} \cdot \frac{(a-1)\mathsf{CA}(d')}{n-(\alpha+1)\mathsf{CA}(d')} + \frac{(2c-1)(\alpha+1)\mathsf{CA}(d')}{n-(\alpha+1)\mathsf{CA}(d')} \cdot \frac{(2c-1)(a-1)\mathsf{CA}(d')}{n-(\alpha+1)\mathsf{CA}(d')}$ or $\frac{n-2c\alpha\mathsf{CA}(d')-\mathsf{CA}(d')}{n-(\alpha+1)\mathsf{CA}(d')} \cdot \frac{a\mathsf{CA}(d')}{n-(\alpha+1)\mathsf{CA}(d')} + \frac{(2c-1)\alpha\mathsf{CA}(d')}{n-(\alpha+1)\mathsf{CA}(d')} \cdot \frac{(2c-1)a\mathsf{CA}(d')}{n-(\alpha+1)\mathsf{CA}(d')}$, depending on whether the previous point resulted in a fit or miss, respectively.

As it can be seen from the formulas, there are non-trivial changes to the probabilities after throwing a point depending on whether it resulted in a hit or a miss. That is, the probability of k hits after throwing a ($> k$) balls in this model is determined not only by the values of k, a, and the available space for the points, but also by the order of the hits among the a thrown points. While given a complete specification of the problem (namely, the parameters p, d, s, r, and α), we can compute the probability of the quantity $c = |\pi_x(S_1) \cap \pi_x(S_2)|$ taking on any given value, a lack of complete characterization of this distribution prevents us from determining its mean value and therefore clearly defining a strategy for the distinguishing attack.

To mitigate the issue, we propose to determine the necessary mean value by additionally simplifying the model. We proceed with the hypergeometric distribution analyzed in equation 1 for exact (rather than approximate) point overlaps. Note that when matching is approximate and each point occupies an area of size $\mathsf{CA}(d/2) = \pi(d/2)^2$, we need to use bins of size $\mathsf{CA}(d/2)$ instead of previous size 1. This means that b becomes $b = n - \alpha\mathsf{CA}(d/2)$ and each hit removes the area $\mathsf{CA}(d/2)$ from the available space. We obtain the mean value of $E[X] = \frac{a^2\mathsf{CA}(d/2)}{b}$, where $a = r - \alpha$, $b = n - \alpha\mathsf{CA}(d/2)$, and $\alpha \geq s - t/2$ for related biometric samples and $\alpha \leq s - t/2 - 1$ otherwise. Then based on this estimate of $E[X]$, the adversary's success probability can be computed using the modified equation 1 similar to equation 2 or using equation 3. We obtain that the number of spurious overlaps between chaff and other biometric points without a match would increase by approximately a multiplicative factor of $\mathsf{CA}(d/2)$ compared to the case when the points had to be matched exactly. While this lowers the adversary's distinguishing probability (compared, for instance, to the probabilities in Figure 1), it also substantially increases the cost of $\mathsf{Rec}(w', S)$ that reconstructs w from its sketch S and a legitimate related biometric w'. Therefore, to maintain usability of the secure sketch scheme, when approximate point matching is used, either the granularity of the available space n needs to be increased or the number of chaff points $r - s$ needs to be decreased. This will allow us to keep the number of spurious overlaps low, as desired. This means that we go back to the range of adversary's success probabilities depicted in Figure 1, which are clearly unacceptably high for a security construction.

Attacking irreversibility. Now the adversary is given two sketches S_1 and S_2 for related biometrics w and w' and its goal is to recover information about w beyond what can be learned due to the release of the sketch. We analyze the case when both S_1 and S_2 are produced using the fuzzy vault scheme. The attack strategy in this case consists of first computing the x-coordinates common to both S_1 and S_2, i.e., $\pi_x(S_1) \cap \pi_x(S_2)$. Recall that the resulting set contains at least $s - t/2$ genuine points from $w \cap w'$ as well as spurious points (i.e., other than points in $w \cap w'$), the expected number of which is $(r - s + t/2)^2/(n - s + t/2)$ for related biometrics with distance t. Let $c = |\pi_x(S_1) \cap \pi_x(S_2)| - (s - t/2)$. Then to attempt to recover the biometric w, the adversary will consider every subset of size $s - t/2$ from the intersection $\pi_x(S_1) \cap \pi_x(S_2)$, reconstruct the polynomials that those points form in S_1 and S_2, and use the remaining points of S_1 and S_2 to test whether the recovered polynomials could be the original polynomials that determine w and w'. In more detail, any $s - t$ points from S_1 uniquely define a polynomial of degree $s - t - 1$, which can be reconstructed from the x and y-coordinates included in the sketch. Then if this polynomial was the original polynomial chosen at the time of S_1 creation, there will be exactly t other points in S_1 lying on this polynomial, at least $t/2$ of which must be in the intersection $\pi_x(S_1) \cap \pi_x(S_2)$. The attack strategy then consists of choosing $s - t/2$ points from $\pi_x(S_1) \cap \pi_x(S_2)$, reconstructing the polynomial using the first $s - t$ points and their coordinates in S_1, and checking whether the rest of the selected points ($t/2$ of them) lie on that polynomial *and* exactly $t/2$ points from the remaining $r - s + t/2$ points of S_1 lie on that polynomial as well. The same steps are then repeated for the second sketch S_2. More formally, the attack steps are as follows:

1. Compute $\pi_x(S_1) \cap \pi_x(S_2)$.
2. For each subset C of $\pi_x(S_1) \cap \pi_x(S_2)$ of size $s - t/2$:
   (a) Reconstruct the two unique polynomials $p_1(\cdot)$ and $p_2(\cdot)$ of degree $s-t-1$ using the first $s-t$ points and their y-coordinates in S_1 and S_2, respectively.
   (b) Test whether all remaining $t/2$ points lie on p_1 and p_2 as well. If not, proceed with the next set; otherwise, continue.
   (c) Use S_1 to compute the number of points from $\pi_x(S_1) \setminus C$ lying on p_1. Similarly, use S_2 to compute the number of points from $\pi_x(S_2) \setminus C$ lying on p_2. If both numbers equal to $t/2$, store C as a potential set that represents the original polynomials.
3. Choose one of the stored sets C at random and output the s points of S_1 that lie on the corresponding polynomial $p_1(\cdot)$ as the guess for w.

To quantify the success of this attack, let for the sake of the current description assume that $\mathsf{dist}(w, w') = t$. First note that there are $s - t/2 + c$ points common to S_1 and S_2 and therefore there are $\binom{s-t/2+c}{s-t/2} = \binom{s-t/2+c}{c}$ subsets C of size $s - t/2$ that an attacker needs to try. Because the given sketches are related, there is at least one set C that passes verification in step 2 above. The attacker can only fail if there are more than one pair of candidate polynomials and the attacker chooses a wrong pair. In fact, if the number of such polynomial pairs is $k \geq 1$, then the attacker can fully recover the biometric with the probability $1/k$. This gives us:

$$\mathsf{Adv}^{\mathsf{irrev}}_{\mathcal{A}} = \frac{2^{m'}}{2^{m'} - 1} \left(\frac{1}{k} - \frac{1}{2^{m'}} \right). \tag{4}$$

We know that k is always at least 1, and we next argue that k is small. To show this, suppose that a set C is not from $w \cap w'$ (i.e., some points of C are not genuine biometric points), but passes verification in step 2 above. Both p_1 and p_2 generated by the points in C have to pass through all remaining $t/2$ points of C and exactly $t/2$ points of the remaining $r - s + t/2$ points of S_1 and S_2, respectively. The probability that each of these t points falls on a given polynomial is at most $\frac{1}{n-1}$. This is because the chaff points are chosen at random from $n-1$ options, and if the polynomial happens to fall on the remaining field value, the probability of success is 0, and it is $\frac{1}{n-1}$ otherwise. Then the probability that $t/2$ points of C happen to lie on the polynomial is at most $(\frac{1}{n-1})^{t/2}$ and the probability that any $t/2$ points from $r - s + t/2$ happen to lie on the polynomial is at most $\binom{r-s+t/2}{t/2}(\frac{1}{n-1})^{t/2}$. To pass the verification in step 2, it must also be the case that the rest of the points from the sketch (i.e., $r - s$ of them) do not fall on the reconstructed polynomials. Because the latter probability might not be significantly smaller than 1, we bound it by 1 from the above. This allows us to compute the expected value of k as follows:

$$E[k] < 1 + \left(\binom{s - t/2 + c}{c} \binom{r - s + t/2}{t/2} \left(\frac{1}{n-1} \right)^t \right)^2 < 1 + \left(\frac{(s - t/2 + c)^c (r - s + t/2)^{t/2}}{c! \cdot (t/2)! \cdot (n-1)^t} \right)^2 \tag{5}$$

where the last inequality uses the approximation $\binom{n}{k} < \frac{n^k}{k!}$. Because $s \ll r \ll n$, the expected value of k is small. For example, using the parameters $n = 63001$, $s = 30$, $t = 16$, and $r = 300$ suggested in [6] with $\mathsf{dist}(w, w') = t$ and therefore $E[c] = 1$, we obtain $E[k] < 1 + 8.9 \cdot 10^{-122} = 1$.

When, $\mathsf{dist}(w, w') < t$, there will be multiple sets C of size $s - t/2$ that pass verification in step 2 of the attack. These sets correspond to the same original polynomials $p_1(\cdot)$ and $p_2(\cdot)$ of S_1

and S_2, respectively. The number of such sets is $\binom{s-\mathsf{dist}(w_1,w_2)/2}{t/2-\mathsf{dist}(w_1,w_2)/2}$, while the number of spurious sets and the corresponding polynomials that pass the verification can be characterized using the generalization of the quantity in equation 5. We obtain:

$$E[k] < \binom{s-\mathsf{dist}(w_1,w_2)/2}{t/2-\mathsf{dist}(w_1,w_2)/2} + \left(\left(\binom{s-t/2+c}{c} - \binom{s-\mathsf{dist}(w_1,w_2)/2}{t/2-\mathsf{dist}(w_1,w_2)/2}\right)\binom{r-s+t/2}{t/2}\left(\frac{1}{n-1}\right)^t\right)^2$$

This analysis suggests a slight modification to the attack above: if multiple sets C that correspond to the same polynomials p_1 and p_2 pass the verification in step 2, instead of choosing one of the sets at random, choose a set corresponding to the repeated polynomials p_1 and p_2. In this case, the attacker's success probability approaches 1 regardless of the number of spurious sets C that pass the verification. In general, we then have that $\mathsf{Adv}^{\mathsf{irrev}}_{\mathcal{A}}$ is greater or equal to the quantity given in equation 4.

The complexity of our attack is bounded from the above by $\binom{s-t/2+c}{c}(r-s+t/2)$ operations. Because c must be low for the fuzzy vault scheme to be practical, this amount of work is not expected to result in a large computational burden.

3.1.2 Improved fuzzy vault

Attacking indistinguishability. An important observation in designing an attack strategy for this construction is that it is deterministic. This immediately implies that the same biometric will always produce the same secure sketch, giving the adversary the ability to distinguish between the sketches. Thus, as an important special case we first consider the adversary's ability to win the indistinguishability game when no noise affects multiple sketches of the same w (this arises in several applications, where multiple keys are issued using the same copy of w). Thus, when $\mathcal{A}$ obtains challenge S_2, it outputs 1 if $S_2 = S_1$ and 0 otherwise. This means that when $b = 1$, $\mathcal{A}$ will always guess the bit correctly, but when $b = 0$ it might still sometimes output 1 if the two sketches happened to be the same. The probability of the latter, however, is small and we next provide a bound on its value.

Recall that sketch S consists of t coefficients of a polynomial $p(x) = x^s + c_{s-1}x^{s-1} + \ldots + c_1x + c_0$, where for biometric $w = \{w_1, \ldots, w_s\}$ the coefficients are $c_{s-1} = \sum_i w_i$, $c_{s-2} = \sum_{i \neq j} w_i w_j$, $\ldots$, $c_{s-t} = \sum_{C \subset [1,s], |C|=t} (\prod_{i \in C} w_i)$. First, for an unrelated random biometric $\hat{w}$, the probability that $\sum_i \hat{w}_i = c_{s-1}$ is $\frac{1}{n}$. That is, without any restrictions, there are $\prod_{i=0}^{s-1}(n-i)$ choices for s elements without repetitions from the set of n elements, and when the sum of the elements is fixed (in $\mathbb{F}_n$), the number reduces to $\prod_{i=1}^{s-1}(n-i)$.

Now let us consider c_{s-2}. We start with a simpler function $x_1x_2 = b$ in $\mathbb{F}_n$ for a fixed value of b. Recall that $n = p^2$ for a prime p. We enumerate all possible solutions x_1 and x_2 for this function such that $x_1 \neq x_2$ (since all points in a biometric are different). When b is the zero element, there are $n-1$ unordered pairs (x_1, x_2) with $x_1 \neq x_2$ whose product equals to b (one value is zero and the other can take $n-1$ remaining values). All elements other than zero form a cyclic multiplicative group, and when $b \neq 0$ there are either $\frac{n-1}{2}$ or $\frac{n-1}{2} - 1$ pairs (x_1, x_2) with distinct x_1 and x_2, when b is a quadratic non-residue or quadratic residue, respectively. Therefore, the number of pairs (x_1, x_2) satisfying the congruence for any value of b is at most $n-1$ from the overall space of $\frac{n(n-1)}{2}$ such pairs, giving us the fraction $\frac{2}{n}$.

Now recall that c_{s-2} is composed of a summation of products w_iw_j for each $i \neq j$. Then when there is only one product w_1w_2 (i.e., $s = 2$), we obtain that it is equal to 0 more frequently than to other values. When, however, $s > 2$ the distribution of product values drastically changes.

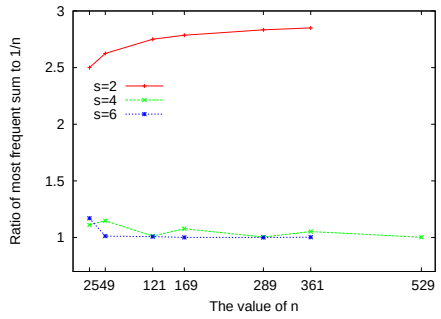

Figure 3: The ratio of the fraction of most frequent value of the sum c_{s-2} to $\frac{1}{n}$ for different values of n and s.

Because all w_i have to be unique and each w_i appears in a number of products $w_i w_j$, the value of the sum tends to be distributed more evenly as s increases. This means that the frequency of the most common value of c_{s-2} approaches $\frac{1}{n}$ when s grows. To illustrate this phenomenon, we plot empirical data for small values of $n = p^2$. In particular, for $s = 2$, 4, and 6 and all possible $w = (w_1, \ldots, w_s) \in \mathbb{F}_n^s$ we find a value of the sum which occurs the highest number of times. Let that value be denoted by $count_{\mathsf{max}}$ and the fraction of all biometrics w that results in such value by $f_{\mathsf{max}} = count_{\mathsf{max}}/\binom{n}{s}$. To evaluate how the value of f_{max} compares to $\frac{1}{n}$, we plot their ratio $f_{\mathsf{max}}/\frac{1}{n}$ in Figure 3. For $s = 2$, $f_{\mathsf{max}} = \frac{2}{n}$ is constant; for $s > 2$ it is clear that f_{max} rapidly approaches $\frac{1}{n}$ from the above even for very small values of s. This means that $\frac{2}{n}$ is a generous upper bound on the probability that c_{s-2} of a randomly chosen $\hat{w}$ will coincide with a specific value of that coefficient for an unrelated biometric w.

Extending this analysis to $c_{s-3} = \sum w_i w_j w_k$, where i, j, and k are pairwise distinct, we obtain that the most frequently occurring value of c_{s-3} is 0 and when $s = 3$ (i.e., there is only one product). In that case, the number of possibilities that result in that product is $\frac{(n-1)(n-2)}{2}$ out of $\frac{n(n-1)(n-2)}{2 \cdot 3}$ total choices (and the number of possibilities when the product is non-zero is at most $\frac{n-3}{2} \cdot \frac{n-1}{2}$). This gives us that the fraction of triples that can result in any given product from the overall space is $\leq \frac{3}{n}$. For c_{s-4}, the maximum fraction is $\leq \frac{4}{n}$; for c_{s-5}, it is $\leq \frac{5}{n}$, etc. Therefore, the adversarial error is at most $\frac{t!}{n^t}$, and in practice will be close to $\frac{1}{n^t}$ because $s > t$. Both of these quantities are very low even for small values of t (e.g., 2). This gives us that the probability with which the adversary can mistakenly consider two unrelated biometrics to be related is very small. The adversary's advantage in the 2-indistinguishability game then is:

$$\begin{aligned}
\mathsf{Adv}^{\mathsf{ind}}_{\mathcal{A}} &= 2\left|\Pr[b' = b] - \frac{1}{2}\right| = 2\left|\Pr[b' = 1 | b = 1]\Pr[b = 1] + \Pr[b' = 0 | b = 0]\Pr[b = 0] - \frac{1}{2}\right| = \\
&= \left|2\Pr[b' = 1 | b = 1]\frac{1}{2} + 2\Pr[b' = 0 | b = 0]\frac{1}{2} - 1\right| = \left|\Pr[b' = 1 | b = 1] + \right. \\
&+ \left. 1 - \Pr[b' = 1 | b = 0] - 1\right| = \left|\Pr[b' = 1 | b = 1] - \Pr[b' = 1 | b = 0]\right| > 1 - \frac{t!}{n^t}.
\end{aligned}$$

To address the problem of distinguishability in the general case, assume we are given two sketches, $S_1 = (c_{s-1}, c_{s-2}, ..., c_{s-t})$ and $S_2 = (c'_{s-1}, c'_{s-2}, ..., c'_{s-t})$. From Vieta's formulas, we know the relation between the roots of a normalized polynomial (i.e., when the leading coefficient is 1) and its coefficients. For example, the summation of the roots is equal to the negative of the second leading coefficient. Suppose that the two biometrics w and w' corresponding to the sketches S_1 and S_2, respectively, have $s - k$ elements in common, where k can take any integer value from 0 up to s. That is, k elements of w are different from some other k elements of w'. In our case, k is unknown,

and our goal is to determine its value, which will allow us decide whether these two biometrics are related or not. We can easily derive equations of uncommon elements of the biometrics and the coefficients in the given sketches by subtracting coefficients of the sketches. The first two equations are:

$$\begin{aligned} c_{s-1} - c'_{s-1} &= \sum_{i=1}^{i=k} w_i - \sum_{i=1}^{i=k} w'_i \\ c_{s-2} - c'_{s-2} &= (c_{s-1} - \sum_{i=1}^{i=k} w_i) \sum_{i=1}^{i=k} w_i + \sum_{i,j=1, i\neq j}^{i,j=k} w_i w_j - (c'_{s-1} - \sum_{i=1}^{i=k} w'_i) \sum_{i=1}^{i=k} w'_i - \sum_{i,j=1, i\neq j}^{i,j=k} w'_i w'_j \end{aligned}$$

where w_i's and w'_i's are the uncommon roots of w and w', respectively. Similar to the above equations, we can construct all t equations generated by the first known t coefficients of the two sketches. To be able to do so, however, we will have to guess the value of k. Here, the adversary's strategy is to set the value of k to $t/2$. Then if these two sketches are in fact related with $\mathsf{dist}(w, w') = t$, the equations can be solved for w_i's and w'_i's. If the sketches are related and the distance between them is less than t, the adversary is overestimating their difference (i.e., $k < t/2$), yet the equations will lead to an acceptable solution. That is, in case of $\mathsf{dist}(w, w') < t$, some roots in the equations corresponding to the elements of w will be equal to roots corresponding to the elements of w'. Therefore, if the biometrics are indeed related, the adversary will find acceptable and valid solutions for the equations and with probability 1 will output a correct guess that the biometrics are related. If the sketches, however, are not related, the adversary will do her best to solve the equations. There are two possible cases: (i) the adversary does not find any valid solution in the field for this equation set and thus outputs that the biometrics are not related or (ii) the adversary in fact finds a valid and acceptable solution set to these equations and outputs that the biometrics are related. Because in the latter case the adversary makes a wrong guess, we need to find the probability that two non-related sketches lead to a set of equations that will result in a valid solution in the field.

To do so, we re-write the adversary's advantage in the 2-indistinguishability game as:

$$\mathsf{Adv}^{\mathsf{ind}}_{\mathcal{A}} = 2\left|\Pr[b' \neq b] - 1/2\right| = 2\left|\Pr[b' \neq b | b = 1]\Pr[b = 1] + \Pr[b' \neq b | b = 0]\Pr[b = 0] - 1/2\right|$$

We also know that $\Pr[b' \neq b | b = 1] = 0$ and let $\Pr[b' \neq b | b = 0] = q$. Then we obtain $\mathsf{Adv}^{\mathsf{ind}}_{\mathcal{A}} = 2\left|(1/2)q - 1/2\right| = 1 - q$. To find a bound for probability q, we use the fact that each sketch is a set of coefficients of an s-degree polynomial. This means that the total number of biometrics in this representation is $\binom{n}{s}$ and the total number of related biometrics is $R = \sum_{i=0}^{t/2} \binom{s}{i}\binom{n-s}{i}$, where i represents the number of elements in a biometric that are different from the original one. This give us the total number of non-related biometrics $NR = \binom{n}{s} - R$. We can see that R is dominated mainly by the last factor, i.e., $\binom{s}{t/2}\binom{n-s}{t/2}$.

When we reveal t coefficients, we are looking at a space of size n^t since each coefficient can take any value from the field. In total, we could have n^t possibilities for sketches. On the other hand, we can have up to R related sketches for any given biometric. To determine how closely the above observations correspond to the numbers observed in practice, we performed experiments for different values of s and t with fixed p and counted the number of unique sketches that related and non-related biometrics generate. In our experiments, we observed that each related biometric generates a different set of coefficients and thus produces a unique sketch among the sketches corresponding to related biometrics. Therefore, from the overall possible space (of size n^t), sketches for biometrics related to any given biometric w occupy a subset of size R. The next step is to determine the distribution of non-related sketches over the entire possible set of sketches, which will allows us find

Biometric size	$t = 2$		$t = 4$		$t = 6$	
	computed	observed	computed	observed	computed	observed
$s = 6$	$32,565$	$32,221$	$1,196$	$1,080$	N/A	N/A
$s = 7$	$97,652$	$97,429$	$4,081$	$4,392$	58	0
$s = 8$	$237,050$	$236,972$	$10,883$	$11,577$	179	34

Table 1: The total number of collisions between related and unrelated sketches (error count) based on theoretical analysis (denoted "computed") and experimental validation (denoted "observed").

the probability of error. That is, the distribution of the non-related sketches of a given biometric will allow us to determine the number of them colliding with the related sketches. Our experiments (detailed below) suggest that this distribution is uniform over the entire range of possibilities that non-related sketches can take. This information tells us that the expected number of non-related biometrics which generate sketches that are same as one of the related sketches will be approximately $(R/(n^t) \cdot NR)$; we call this number "the error count." The error count allows us to compute the probability that the adversary fails to determine the answer correctly, that is, given a pair of non-related sketches, it declares them as related when a related sketch with respect to one of the given sketches is the same as the other given sketch. This probability is the ratio of the error count and the total number of sketches produced in this setting, $\binom{n}{s}$, which gives us $(R/(n^t) \cdot (NR/\binom{n}{s}))$. We can see that NR is fairly close to $\binom{n}{s}$ when t is not large, giving us an approximation of the failure probability $q \approx R/(n^t)$.

To confirm this analysis, we show our experimental results. In the experiments, we counted the total number of collisions between related and non-related sketches to produce the value of the error count and compared it to the value computed according to the formula. The experiments were run for $s = 6, 7, 8$ when p was set to 5. We provide the error count values both computed according to the formula and empirically counted (observed) in Table 1. It is clear that there is a small difference between the two types of values.

To simplify the expression for q and find a lower bound on the adversary's advantage, we use an approximation of the formula $q \approx \binom{s}{t/2}\binom{n-s}{t/2}/(n^t)$ by replacing $\binom{n}{k}$ with $\leq n^k/(k!)$ and obtain $q \lesssim (s(n-s))^{t/2}/((t/2)!^2 n^t)$. The maximum of this function happens when $s = n/2$ and $t = 2$ which leads to q taking the value $1/4$. We obtain that $\mathsf{Adv}^{\mathsf{ind}}_{\mathcal{A}} \gtrsim 3/4$ regardless of the values of n, s, or t.

Note that the adversary's computation is mainly dominated by methods for solving the equations. The first equation has degree 1 and each consecutive equation's degree increases by one from the one before. Overall, we have t variables and the last equation has degree t. Solving this equation set of multivariate polynomials is generally an NP-complete problem [8]. However, in our case we know that the variables w_i are different from each others and so are w'_i. The total number of possibilities for assigning values to these variables from a field of $n = p^2$ elements is $\binom{n}{t/2}^2$, which is roughly $(\frac{1}{\pi t})(\frac{2en}{t})^t$. Therefore, for small values of n and t this attack is quite feasible for a computationally bounded adversary. On the other hand, when n and t increase, existing publications such as [8] show how to approach the general problem of solving multivariate polynomial equations over a finite field. One can apply these methods to reduce the complexity of the proposed attack. We can therefore conclude that the improved fuzzy vault scheme is computationally resistant to our attack for very large values of n and t. However, in reality, the more common case is when the difference between biometrics is fairly small. We have seen when $t = 0$ this attack becomes relatively simple to be performed. Also when t is a small number, the computation complexity of $O(n^t)$ can be affordable.

Attacking irreversibility. Now the goal is to extract the original biometric w, given its secure sketch S_1 and a sketch S_2 of a related biometric w'. Note that the strategy above recovers biometric points in order to determine whether w and w' are related or not. It is important to notice, however, that the points that the adversary recovers are those by which w and w' differ from each other. Therefore, even though the adversary can learn $t/2$ points of w and $t/2$ points of w' (including some points common to them if $\mathsf{dist}(w, w') < t$), this strategy does not lead to the recovery of sufficient information to gain a non-negligible advantage in the irreversibility game. In particular, the entropy loss of this type of secure sketch is $t \log n$, which is the same as recovering t biometric points. We thus obtain that if the adversary uses t recovered points and guesses the remaining $s - t$ points it cannot win the irreversibility game with a sufficiently large probability. The success probability, however, can substantially grow when the adversary is able to obtain more than two sketches that correspond to related biometrics. After recovering t points from each pair of sketches, the adversary will likely be able to obtain a larger number of points of w and thus gain a non-negligible advantage in recovering information about w which is not available from its single sketch S_1.

3.1.3 Pinsketch

Attacking indistinguishability. The adversary receives two secure sketches $S_1 = \mathsf{syn}(w_1) = (s_1, s_3, \ldots, s_{2t-1})$ and $S_2 = \mathsf{syn}(w_2) = (s'_1, s'_3, \ldots, s'_{2t-1})$, and its goal is to determine the coin flip, i.e., whether the biometrics w and w' are related or not. Because the reconstruction procedure requires computation of the syndrome of the (noisy) biometric, the adversary's strategy in this case is simple: compute $\sigma_i = s'_i - s_i$ for each i, and compute $|\mathsf{supp}(v)|$ such that $\mathsf{syn}(v) = (\sigma_1, \sigma_3, \ldots, \sigma_{2t-1})$. If $|\mathsf{supp}(v)| \leq t$, output 1 (related), otherwise, output 0.

To analyze the success probability of $\mathcal{A}$, we first note that the adversary will always guess correctly when w and w' are related. When w and w' are not related, the resulting $\mathsf{syn}(v)$ can either be decodable or not decodable. Because a linear code (i.e., BCH) is used, the success probability of the adversary is exactly the success probability in distinguishing sketches in the syndrome construction with linear codes analyzed in [26]. The analysis of [26] shows that the probability of decoding $\mathsf{syn}(v)$ when w and w' are unrelated is small, and the adversary wins the game with overwhelming probability. We refer the reader to [26] for additional definitions regarding linear codes and their analysis.

Attacking irreversibility. Given two biometrics w and w' produced using a (n, k_1, t_1) linear code C_1 and a (n, k_2, t_2) linear code C_2, respectively, [26] shows that when $w = w'$, the adversary's advantage in recovering w for the syndrome construction is

$$\mathsf{Adv}^{\mathsf{irrev}}_{\mathcal{A}} = \frac{1}{2^{\min(m'_1, m'_2)} - 1} \left(\frac{2^{\min(m'_1, m'_2)}}{2^{k_1 + k_2} - \mathrm{Rank}(G_{1,2})} - 1 \right)$$

where $G_{1,2}$ denotes the $(k_1 + k_2) \times n$ matrix $\begin{bmatrix} G_1 \\ G_2 \end{bmatrix}$ and G_1 (resp., G_2) is the generator matrix of C_1 (resp., C_2). When, however, $w \neq w'$, the adversary will need to iterate over all possible error patterns and verify its guess, which becomes large when t is large. We refer the reader to [26] for additional details.

3.2 Construction for edit distance

Dodis et al. [10] describe two alternative ways of realizing the secure sketch construction for the edit distance given in Section 2.3. The first consists of applying an existing low-distortion embedding

(that does not significantly change the distance between two biometrics after the mapping) of the edit distance into the Hamming distance and then using a syndrome construction for the Hamming distance to produce the public data. The second includes the application of a specially designed embedding of the edit distance into the set difference metric using so-called c-shinglings. After the embedding of a biometric is performed, the Pinsketch construction is applied to the resulting representation to compute the sketch. Note that in both of the above cases linear error-correcting codes are used, which means that the strategy and the analysis of [26] is applicable to both cases. We conclude that constructions for the edit distance metric do not achieve the indistinguishability and irreversibility properties.

4 Our Schemes

In this section we describe our simple schemes secure against strong adversaries with provably no information leakage in the computational model. In what follows, let $(\mathsf{SS}', \mathsf{Rec}')$ denote any existing fuzzy sketch scheme (for any metric). The key k denotes the long-term user's key of size κ, where κ is the security parameter. The key is not shared with any parties. Before proceeding with the description of the schemes, we provide additional definitions.

Definition 7 *Let $F : \{0,1\}^\kappa \times \{0,1\}^{\ell_1(\kappa)} \to \{0,1\}^{\ell_1(\kappa)}$ be a family of functions. For $k \in \{0,1\}^\kappa$, the function $F_k : \{0,1\}^{\ell_1(\kappa)} \to \{0,1\}^{\ell_1(\kappa)}$ is defined as $F_k(x) = F(k,x)$. F is said to be a family of pseudo-random functions (PRF) if for every PPT adversary $\mathcal{A}$ with oracle access to a function and all sufficiently large κ $|\Pr[\mathcal{A}^{F_k}(1^\kappa) - \Pr[\mathcal{A}^f(1^\kappa)]|$ is negligible in κ, where $k \stackrel{R}{\leftarrow} \{0,1\}^\kappa$ and f is a function chosen at random from all possible functions mapping $\ell_1(\kappa)$-bit inputs to $\ell_1(\kappa)$-bit outputs.*

Definition 8 *A family of functions $h : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^{\ell_2(\kappa)}$ is pairwise independent universal hash function if for all $x, x' \in \{0,1\}^n$, where $x \neq x$, $\Pr[h_y(x) = h_y(x')] = 1/2^{\ell_2(\kappa)}$ for $y \in \{0,1\}^\kappa$.*

In the following secure sketch construction, it is required that $\ell_1(\kappa) \geq |\mathsf{SS}'(w)|$, where $|a|$ denotes the length of string a. We discuss the choice of parameters in more detail later in this section.

To compute $\mathsf{SS}(w,k)$:

1. Choose $r_1 \in \{0,1\}^{\ell_1(\kappa)}$ at random.
2. Output $S = (S_1, S_2) = (r_1, F_k(r_1) \oplus \mathsf{SS}'(w))$.

To compute $\mathsf{Rec}(w', k, S = (S_1, S_2))$:

1. Compute $u \leftarrow F_k(S_1)$.
2. Output what $\mathsf{Rec}'(w', S_2 \oplus u)$ outputs.

Theorem 1 *Assuming that F is a family of PRFs, the above fuzzy sketch scheme achieves weak biometric privacy.*

Proof Let the adversary attacking the scheme be denoted by $\mathcal{A}$. We relate $\mathcal{A}$'s advantage to the advantage of adversary $\mathcal{A}_F$ attacking the security of PRF family and show that $\mathcal{A}$'s advantage in the weak biometric privacy game is negligible. In the description below, $q = poly(\kappa)$ is the number of sketching queries that $\mathcal{A}$ makes, and $\epsilon_{\mathsf{PRF}}(\kappa)$ denotes $\mathcal{A}_F$'s advantage in breaking the security of PRF family.

Let us define games $\mathbf{G}_0$, $\mathbf{G}_1$, and $\mathbf{G}_2$. The game $\mathbf{G}_0$ is the same as the weak biometric privacy game. To form the game $\mathbf{G}_1$, the challenger ensures that the values r_1 chosen in step 1 of SS

algorithm when forming a sketch to answer a query or create the challenge are unique. Let $r_1^{(i)}$ denote the string chosen in the first step of SS during query i. To form the response to the first query, $\mathcal{A}$ chooses $r_1^{(i)}$ and stores it in the database. To form the response to query i, the challenger chooses $r_1^{(i)}$ and searches its database of $i-1$ entries. If $r_1^{(i)}$ is found, it chooses a new value until it does not appear in the database. The new $r_1^{(i)}$ is used to answer the query and is added to the database. Thus, in game $\mathbf{G}_1$ all $q+1$ values $r_1^{(i)}$ used to answer the queries and form the challenge are different.

The only difference between games $\mathbf{G}_0$ and $\mathbf{G}_1$ is that no collisions happen in $\mathbf{G}_1$. Because the adversary is likely to learn some information about $\mathsf{SS}'(w)$ if a collision occurs, we obtain: $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{wbp1}}(\kappa) \geq \mathsf{Adv}_{\mathcal{A}}^{\mathsf{wbp}}(\kappa) - \binom{q+1}{2}/2^{\ell_1(\kappa)}$, where $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{wbp1}}(\kappa)$ denotes $\mathcal{A}$'s advantage in winning game $\mathbf{G}_1$ and $\binom{q+1}{2}/2^{\ell_1(\kappa)}$ is the upper bound on the probability that at least any two $r_1^{(i)}$ values from the challenger's $q+1$ responses coincide.

To construct game $\mathbf{G}_2$, we use a series of sub-games $G_0, \ldots, G_{q+1}$, where G_0 is the same as $\mathbf{G}_1$ and G_{q+1} is the same as $\mathbf{G}_2$. To form game G_1, the challenger changes its response to the first sketching query by replacing $F_k(r_1^{(1)})$ with a string $r_2^{(1)}$ chosen uniformly at random from $\{0,1\}^{\ell_1(\kappa)}$ and returns the sketch $S_1 = (r_1^{(1)}, r_2^{(1)} \oplus \mathsf{SS}''(\delta_1(w_{b_1})))$ instead. The rest of the game proceeds unmodified as in $\mathbf{G}_1$. To form game G_i, the challenger similarly modifies its response to the first i queries: instead of using F's pseudorandom output in forming the response to an $\mathcal{A}$'s sketching query, the challenger replaces it with a random string $r_2^{(j)}$. Finally, in the game $G_{q+1} = \mathbf{G}_2$ all instances of F's output in the q queries and the challenge are replaced with uniformly random strings.

Now observe that the adversary $\mathcal{A}$ cannot learn any information about sketches in game $\mathbf{G}_2$ because all biometric-related information is perfectly protected. This implies that the only way for $\mathcal{A}$ to obtain any advantage in breaking game $\mathbf{G}_1$ is by using $\mathcal{A}_F$'s advantage. In other words, $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{wbp1}}(\kappa) \leq \epsilon_{\mathsf{PRF}}(\kappa)$. Putting everything together, we obtain $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{wbp}}(\kappa) \leq \epsilon_{\mathsf{PRF}}(\kappa) + \binom{q+1}{2}/2^{\ell_1(\kappa)}$. This means that the only way for $\mathcal{A}$ to have a non-negligible advantage in breaking the security of the scheme is by having non-negligible advantage $\epsilon_{\mathsf{PRF}}(\kappa)$, which violates the assumption that the PRF family is secure. □

Note that in our fuzzy sketch construction deterministic constructions for the underlying sketch SS' are preferred because they normally produce most concise sketches. This means that the pseudorandom output does not have to be needlessly increased to hide the entire underlying sketch. In the above description, we assumed that the output length of F $\ell_1(\kappa)$ is at least as large as the output length of the given secure sketch $|\mathsf{SS}'(w)|$. While this will hold for many types of biometrics and a reasonable choice of security parameter κ, in some cases the representation of $\mathsf{SS}'(w)$ can be longer. Instead of increasing the security parameter (which will increase the complexity of the scheme), we suggest modifying the algorithm to use more than one application of F to produce a longer pseudo-random sequence. For instance, if $\ell_1(\kappa) < |\mathsf{SS}'(w)| \leq 2\ell_1(\kappa)$, the sketch will be produced as $(r_1, (F_k(r_1)||F_k((r_1+1) \bmod 2^\kappa)) \oplus \mathsf{SS}'(w))$, where $||$ denotes string concatenation. This increases the number of random values on which F is evaluated and thus the probability of their collision. However, as long as the quantity $|\mathsf{SS}'(w)|/\ell_1(\kappa)$ is constant or polynomial in κ, the security guarantees still hold.

In the fuzzy extractor construction below we split the key k into two keys k_1 and k_2. This is done to simplify the analysis. In practice, the sub-keys k_1 and k_2 can be computed by applying a PRF keyed with k to two different inputs.

To compute $\mathsf{Gen}(w, k_1, k_2)$:

1. Compute $S = \mathsf{SS}(w, k_1)$ using the fuzzy sketch scheme above.

2. Choose $r_2 \in \{0,1\}^\kappa$ at random and compute $s \leftarrow h_{r_2}(w)$.
3. Output $P = (S, r_2)$ and $R \leftarrow F_{k_2}(s)$.

To compute $\mathsf{Rep}(w', k_1, k_2, P = (P_1, P_2))$

1. Run $\mathsf{Rec}(w', k_2, P_1)$ to recover w. If it fails, output $\perp$.
2. Otherwise, reproduce the key R as $F_{k_2}(s')$, where $s' \leftarrow h_{P_2}(w)$, and output R.

The above algorithm reports an error if the private string R cannot be reproduced. When it is desirable that failures are not reported explicitly, the Rep procedure can output a (wrong) private string, e.g., computed as $R = F_{k_2}(h_{P_2}(w'))$.

Before proceeding with showing the security of the scheme, we explain the design choices made in this construction. Because a pseudo-random function is a powerful primitive, it by itself is sufficient to produce the private string R which is indistinguishable from random. For example, setting $R \leftarrow F_{k_2}(w||r)$, where r is a randomly chosen string, would satisfy the requirements of the security game. The reason for including the hash function h in the construction is to compress the biometric w without loosing the amount of its unpredictability. That is, the n-bit representation of biometric is normally substantially longer than the m bits of entropy it contains. For example, for iris the standard values of these parameters are $n = 2048$ and $m = 256$. Because $m \sim \kappa$, we can use a hash function $h : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^m$ to reduce the size of w from n to m bits without loosing its entropy. In cases when the value of m exceeds the desired length to be used as the input to a PRF, the hash function output length can be further reduced, i.e., in general $\ell_2(\kappa) \leq m$.

We note that the existing generic conversion of a secure sketch to a fuzzy extractor (given in Section 2.1) uses a strong extractor, which can be built using a universal hash function alone. The use of the hash function in a strong extractor is, however, constrained in that the output length of the extractor must necessarily be smaller than m to be able to meet the requirement of the output being close to the uniform distribution. In particular, at least $2\log(\frac{1}{\epsilon}) - 2$ bits of entropy are lost, where the parameter ϵ determines the statistical distance between the distribution of the output and the uniform distribution. In our case, no requirements on the uniformity of the output must be met, and therefore no reduction of the output length or entropy loss has to take place.

Theorem 2 *Assuming that F is a family of PRFs and h is a universal hash function, the above fuzzy extractor scheme achieves strong biometric privacy.*

Proof To show strong biometric security, we consider two games $\mathbf{G}_0$ and $\mathbf{G}_1$. The game $\mathbf{G}_0$ is as in the security definition, and there is a series of games $G_0, \ldots, G_{2q'+q}$ that lead from $\mathbf{G}_0$ to $\mathbf{G}_1$. Instead of modifying the game to first avoid collisions between randomly chosen values, this time we first replace all pseudo-random strings with truly random and then take collisions into account during the analysis of the modified game. Because of the possibility of collisions, we need to ensure that all queries that happen to use the same randomness are answered consistently. In particular, in game G_1, the challenger chooses $r_1^{(1)}$ to answer the first public query and replaces the output of $F_{k_1}(r_1^{(1)})$ with a truly random string $\hat{r}^{(1)}$. The challenger stores the pair $(r_1^{(1)}, \hat{r}^{(1)})$. Before any other public query is answered, the challenger checks whether the string $r_1^{(i)}$ equals to $r_1^{(1)}$. If it does, it uses $\hat{r}^{(1)}$ to create the sketch, otherwise, it proceeds as specified by using $F_{k_1}(r_1^{(i)})$. In game G_2, the use of $F_{k_1}(r_1^{(2)})$ in the second query is also replaced with a random string. The pairs $(r_1^{(j)}, \hat{r}^{(j)})$ for $j = 1, 2$ are stored in a database, and the rest of the queries are answered by first checking the database to ensure that all queries are answered consistently. Similarly, in game $G_{q'}$, invocations of $F_{k_1}(r_1^{(i)})$ in all public queries are replaced with truly random strings.

Now to form game $G_{q'+1}$, the challenger chooses a hash function in the first sketching query by selecting $r_2^{(1)}$, but replaces the invocation of $F_{k_2}(h_{r_2^{(1)}}(\delta_i(w)))$ with a random string $\tilde{r}^{(1)}$. The

challenger stores the pair $(h_{r_2^{(1)}}(\delta_1(w), \tilde{r}^{(1)})$ in the second database. Then in all consecutive public or *private* queries, if the evaluation of $h_{r_2^{(i)}}$ on $\delta_i(w)$ equals to the first element stored in the database, the output of the PRF F_{k_2} on that input is replaced with $\tilde{r}^{(1)}$. In the consecutive games $G_{q'+2}, \ldots, G_{2q'+q}$ invocations of the PRF $F_{k_2}(\cdot)$ are replaced with truly random string one at a time in all public and private queries. The databases are maintained to answer all queries consistently.

Let $\mathsf{Adv}^{\mathsf{wbp1}}_{\mathcal{A}}(\kappa)$ denote $\mathcal{A}$'s advantage in winning game $\mathbf{G}_1$. We obtain that $\mathsf{Adv}^{\mathsf{wbp}}_{\mathcal{A}}(\kappa) \leq \mathsf{Adv}^{\mathsf{wbp1}}_{\mathcal{A}}(\kappa) + \epsilon_{\mathsf{PRF}}(\kappa)$. We next analyze the success probability of $\mathcal{A}$ winning game $\mathbf{G}_1$.

First of all, the adversary may have non-negligible advantage in recovering biometric information when collisions in $r_1^{(i)}$'s used by the challenger occur. The probability of such collisions in public queries is bounded by $\binom{q'}{2}/2^{\ell_1(\kappa)}$ from the above. Collisions in private queries, where the adversary can create public P_i''s with the same $r_1^{(i)}$'s, however, do not provide any advantage to $\mathcal{A}$. Furthermore, if two values $h_{r_2^{(i)}}(\delta_i(w))$ collide, the adversary might also obtain non-trivial advantage in winning the game. The probability that in public queries $h_{r_2^{(i)}}(\delta_i(w))$ and $h_{r_2^{(j)}}(\delta_j(w))$ collide when $\delta_i(w) = \delta_j(w)$ is $1/2^{\kappa}$, i.e., $r_2^{(i)}$ must equal to $r_2^{(j)}$. When $\delta_i(w) \neq \delta_j(w)$, the probability of collision of the hash function outputs $h_{r_2^{(i)}}(\delta_i(w))$ and $h_{r_2^{(j)}}(\delta_j(w))$ is at most $1/2^{\ell_2(\kappa)}$. As before, collisions in private queries do not give the adversary additional advantage because only P_i's returned from the public queries can be used in the challenge. Such collisions can give the adversary advantage if it first obtains P_i's from the challenger and then uses at most $q' - 1$ of them in private queries to obtain the corresponding private keys. If the key associated with the remaining P_i collides with one of the previously queried keys, the adversary will be able to recognize it. Assuming that $\ell_2(\kappa) \geq \kappa$, we obtain that the probability of a collision of the output of the hash function in one of the queried values and the challenge is bounded by $\min(q' - 1, q)/2^{\kappa}$ from the above.

The adversary might attempt to use the private queries to guess the target biometric directly. In particular, $\mathcal{A}$ can try to guess w (or a sufficiently close function of it), compute $\mathsf{SS}'(w)$, and replace it in P_i from one of the public queries with $\mathsf{SS}'(w')$ for sufficiently close w'. If the private query executed on such modified query comes back as not failed, the adversary's guess was correct. The success of such trials is limited by the number private queries and is at most $q/2^m$.

Putting everything together, we have $\mathsf{Adv}^{\mathsf{wbp1}}_{\mathcal{A}}(\kappa) \leq \binom{q'}{2}/2^{\ell_1(\kappa)} + \min(q' - 1, q)/2^{\kappa} + q/2^m$. This gives us $\mathsf{Adv}^{\mathsf{wbp}}_{\mathcal{A}}(\kappa) \leq \epsilon_{\mathsf{PRF}}(\kappa) + \binom{q'}{2}/2^{\ell_1(\kappa)} + \min(q' - 1, q)/2^{\kappa} + q/2^m$. Because each term in this summation is negligible, we conclude that the adversary has only negligible advantage in winning the strong biometric privacy game. □

We would like to note that certain constructions of pseudo-random functions are known to produce uniformly distributed sequences. For example, Shparlinski [24] shows that Naor-Reingold PRF [18] has this property for almost all values of parameters. For our schemes this means that the adversary does not obtain advantage in distinguishing pseudo-random strings from random.

We also note that results similar to ours can be achieved by using encryption instead of PRF, and such schemes might be known or used in industry.

5 Related Work

The overall literature on fuzzy sketches and extractors is very extensive, especially in biometric-related venues, and its overview is beyond the scope of this work. We therefore highlight the most fundamental results and the analysis related to this work. Davida et al. [7] first proposed an off-line biometric authentication scheme, where a user authenticates by presenting a signed output of a hash function over her biometric and other attributes tied to her. Error-correcting codes are used

to reconstruct the original biometric from its noisy consecutive readings. Juels and Wattenberg [14] developed a so-called fuzzy commitment scheme, which became the basis of the code-offset secure sketch construction for the Hamming distance. Juels and Sudan proposed a fuzzy vault scheme in [13], which is a secure sketch construction for the set difference metric. Dodis et al. formalized the notion of secure sketches and fuzzy extractors in their seminal work [11, 10]. That work proposed a generic conversion from a secure sketch to a fuzzy extractor, and developed a number of new or improved schemes for three distance metrics (the Hamming distance, set difference, and edit distance), most of which are outlined earlier in this work.

Boyen et al. [4] introduced the notion of robust fuzzy extractors secure against active adversaries. In that work, the reconstruction process detects tampering with the helper data and fails if the sketch has been modified; the approach relies on random oracles. Dodis et al. [9] continue that line of research and design solutions for certain distance metrics in the standard model. In addition, the keyed setting in the bounded storage model (BSM) is investigated. The use of the key in that work is fundamentally different from our keyed setting: in [9] two parties share a long-term secret key and use it to generate a secret session key with the help of close, but different strings w and w', and the key is used for data authentication. We note that the constructions presented in our work can potentially be applied to a robust fuzzy extractor to improve their properties with respect to reusability.

There are also publications that combine fuzzy extractors with passwords to improve their security properties. One such work [1] develops a biometric key generation algorithm where biometric information is combined with feature selection, after which the derived sketch is secured by using a password and random oracles. This work offers a simpler and more flexible construction.

Security requirements for adequate use of fuzzy sketches and extractors in cryptographic applications have been developing over time. Boyen [3] showed that a number of the original constructions cannot be safely applied multiple times to the same biometric, significantly limiting their usability in practice. That work developed improved constructions using a certain type of error-correcting codes and permutation groups that satisfy the reusability requirements. Our security definitions for the strong adversary were influenced by that work. Later Scheirer and Boult [23] proposed three classes of attacks on secure sketches and fuzzy vault in particular: (1) the record multiplicity attack which takes advantage of a link between related helper data (similar to multiple uses above), (2) the surreptitious key-inversion attack, where the adversary tries to recover the biometric based on any revealed key and corresponding helper data, and (3) the blended substitution attack which considers the problem of injecting false data into the stored records of helper data. The above record multiplicity (or correlation) attack has been empirically evaluated by Kholmatov and Yanikoglu [15] on the fuzzy vault scheme using a database of 400 fuzzy vault sketches (200 matching pairs). The authors were able to unlock (i.e., reconstruct the polynomial) 118 out of 200 pairs within a short period of time using two related vaults. The fuzzy vaults were constructed using polynomials of degree 8 and 200 chaff points. We note that this evaluation was performed on a specific set of parameters already knowing that two stored sketches are related. Our analysis, on the other hand, is more general and can be applied to a wide variety of parameters. Furthermore, it does not assume prior knowledge of related sketches, but rather helps to identify those records. Poon and Miri [22] also describe collusion attacks on the fuzzy vault scheme assuming that the sketches are related. Finally, Simoens et al. [26] introduced the notions of indistinguishability and irreversibility for reusable sketches and showed weaknesses of code-offset and permutation groups constructions. Here we analyze other existing constructions with respect to the indistinguishability and irreversibility properties. The follow-up work [5] investigates similar issues in the continuous domain.

More recently, Simoens et al. [25] provide three different attack strategies for an internal adver-

sary. The approach is general, although it require a considerable number of queries and thus can be prevented by limiting such queries. In addition, our use of key in the proposed scheme makes such attacks irrelevant. In a two-part series of papers [16, 17] for any key-binding or key-generating biometric cryptosystems, Lai et al. study the fundamental trade-offs among security (the length of the generated key), privacy (conditional entropy of the biometric measurements given the helper data), and key protection (conditional entropy of the key given the helper data) in two different cases. The first case [16] is when the biometric sample is used in only one system and the second study [17] considers the case when the same biometric information is used in multiple systems and the attacker will try to combine the data stored in different databases to gain information about either the biometric measurements or the generated keys. In both studies, the authors propose schemes that achieve any point on the chosen trade-off curve. Lastly, Wang et al. provide in [28] an information-theoretic analysis of information leakage and revocability for error-correcting code based implementation of fuzzy commitments and secure sketches. They, too, show that if the stored data is padded with a one-time key, then the system is resistant against linkage attacks across multiple enrollments.

6 Conclusions

This work investigates the security properties of a number of constructions for secure sketches and corresponding fuzzy extractors. We show that, in addition to the constructions that have been previously shown to have security weaknesses, other existing constructions do not meet our security expectations when they are reused on related biometrics. In particular, we analyze a number of secure sketch constructions from the literature for the set difference and edit distance metrics with respect to their indistinguishability and irreversibility in presence of very weak adversaries. Our analysis indicates that none of the schemes can be safely reused.

To mitigate the problem, we propose to use the computational setting, where a user stores a short key for all possible uses in such schemes. This change results in simple solutions with remarkable security and usability improvements which work with any existing secure sketch, mitigate information leakage associated with biometrics, and rely on generic hardness assumptions.

References

[1] L. Ballard, S. Kamara, F. Monrose, and M. Reiter. Towards practical biometric key generation with randomized biometric templates. In *ACM Conference on Computer and Communications Security (CCS)*, pages 235–244, 2008.

[2] M. Blanton and W. Hudelson. Biometric-based non-transferable anonymous credentials. In *International Conference on Information and Communications Security (ICICS)*, pages 165–180, 2009.

[3] X. Boyen. Reusable cryptographic fuzzy extractors. In *ACM Conference on Computer and Communications Security (CCS)*, pages 82–91, 2004.

[4] X. Boyen, Y. Dodis, J. Katz, R. Ostrovsky, and A. Smith. Secure remote authentication using biometric data. In *Advances in Cryptology – EUROCRYPT*, pages 147–163, 2005.

[5] I. Buhan, J. Breebaart, J. Guajardo, K. de Groot, E. Kelkboom, and T. Akkermans. A quantitative analysis of indistinguishability for a continuous domain biometric cryptosystem.

In *International Workshop on Data Privacy Management (DPM)*, volume 5939 of *LNCS*, pages 78–92, 2010.

[6] T. Clancy, N. Kiyavash, and D. Lin. Secure smartcard-based fingerprint authentication. In *ACM SIGMM Workshop on Biometrics Methods and Applications*, pages 45–52, 2003.

[7] G. Davida, Y. Frankel, and B. Matt. On enabling secure applications through off-line biometric identification. In *IEEE Symposium on Security and Privacy*, pages 148–157, 1998.

[8] J. Ding and D. Schmidt. Mutant Zhuang-Zi algorithm. In *Post-Quantum Cryptography*, volume 6061 of *LNCS*, pages 28–40, 2010.

[9] Y. Dodis, J. Katz, L. Reyzin, and A. Smith. Robust fuzzy extractors and authenticated key agreement from close secrets. In *Advances in Cryptology – CRYPTO*, pages 232–250, 2006.

[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 of Computing*, 38(1):97–139, 2008.

[11] Y. Dodis, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *Advances in Cryptology – EUROCRYPT*, pages 523–540, 2004.

[12] Y. Dodis and A. Smith. Correcting errors without leaking partial information. In *ACM Symposium on Theory of Computing (STOC)*, pages 654–663, 2005.

[13] A. Juels and M. Sudan. A fuzzy vault scheme. In *International Symposium on Information Theory*, 2002.

[14] A. Juels and M. Wattenberg. A fuzzy commitment scheme. In *ACM Conference on Computer and Communications Security (CCS)*, pages 28–36, 1999.

[15] A. Kholmatov and B. Yanikoglu. Realization of correlation attack against the fuzzy vault scheme. In *Security, Forensics, Steganography, and Watermarking of Multimedia Contents*, 2008.

[16] L. Lai, S.-W. Ho, and H. Poor. Privacy-security trade-offs in biometric security systems — Part I: Single use case. *IEEE Transactions on Information Forensics and Security*, 6(1):122–139, 2011.

[17] L. Lai, S.-W. Ho, and H. Poor. Privacy-security trade-offs in biometric security systems — Part II: Multiple use case. *IEEE Transactions on Information Forensics and Security*, 6(1):140–151, 2011.

[18] M. Naor and O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 458–467, 1997.

[19] N. Nisan and A. Ta-Shma. Extracting randomness: A survey and new constructions. *Journal of Computer and System Sciences*, 58:148–173, 1999.

[20] N. Nisan and D. Zuckerman. Randomness is linear in space. *Journal of Computer and System Sciences*, 52(1):43–53, 1996.

[21] S. Pankanti, S. Prabhakar, and A. Jain. On the individuality of fingerprints. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(8):1010–1025, August 2002.

[22] H. Poon and A. Miri. A collusion attack on the fuzzy vault scheme. *ISC International Journal of Information Security*, 1(1):27–34, 2009.

[23] W. Scheirer and T. Boult. Cracking fuzzy vaults and biometric encryption. In *IEEE Biometrics Symposium*, pages 1–6, 2007.

[24] I. Shparlinski. On the uniformity of distribution of the Naor-Reingold pseudo-random function. *Finite Fields and Their Applications*, 7(2):318–326, 2001.

[25] K. Simoens, J. Bringer, H. Chabanne, and S. Seys. A framework for analyzing template security and privacy in biometric authentication systems. *IEEE Transactions on Information Forensics and Security*, 7(3):833–841, 2012.

[26] K. Simoens, P. Tuyls, and B. Preneel. Privacy weaknesses of biometric sketches. In *IEEE Symposium on Security and Privacy*, pages 188–203, 2009.

[27] A. Smith. *Maintaining secrecy when information leakage is unavoidable*. PhD dissertation, MIT, August 2004.

[28] Y. Wang, S. Rane, S. Draper, and P. Ishwar. An information-theoretic analysis of revocability and reusability in secure biometrics. In *Information Theory and Applications Workshop (ITA)*, 2011.