

A quantitative analysis of indistinguishability for a continuous domain biometric cryptosystem

Ileana Buhan, Jeroen Breebaart, Jorge Guajardo,
Koen de Groot, Emile Kelkboom, Ton Akkermans

Philips Research Laboratories, Eindhoven, The Netherlands

Abstract. Biometric information is regarded as highly sensitive information and therefore encryption techniques for biometric information are needed to address security and privacy requirements of biometric information. Most security analyses for these encryption techniques focus on the scenario of one user enrolled in a single biometric system. In practice, biometric systems are deployed at different places and the scenario of one user enrolled in many biometric systems is closer to reality. In this scenario, cross-matching (tracking users enrolled in multiple databases) becomes an important privacy threat. To prevent such cross-matching, various methods to create renewable and indistinguishable biometric references have been published. In this paper, we investigate the indistinguishability or the protection against cross-matching of a continuous-domain biometric cryptosystem, the QIM. In particular our contributions are as follows. Firstly, we present a technique, which allows an adversary to decide whether two protected biometric reference data come from the same person or not. Secondly, we quantify the probability of success of an adversary who plays the indistinguishability game and thirdly, we compare the probability of success of an adversary to the authentication performance of the biometric system for the MCYT fingerprint database. The results indicate that although biometric cryptosystems represent a step in the direction of privacy enhancement, we are not there yet.

1 Introduction

When Alice wants to prove her identity to a biometric authentication system she provides a biometric trait and the system compares the measured biometrics to her reference biometric information. If the two match, Alice is authenticated. For the purpose of authentication, in our model Alice does not need an additional password or token and her reference biometric identity is stored by the authenticating entity. The authenticating entity, however, has to safeguard the privacy of Alice. This important responsibility can be addressed using a variety of requirements and techniques for storing and processing biometric data, for details which include template protection techniques, encryption, etc. we refer to [11].

One of the privacy threats spurred by the widespread use of biometric applications is the ability to track users across applications by comparing biometric references facilitated by the uniqueness and persistence of biometric characteristics. Several counter measures have been identified to prevent cross-matching, which include: (1) the avoidance of central databases by the application of the data separation principle, which recommends storing biometric references on an individual secure token or smartcard, (2) the provision of confidentiality of biometric references by encryption techniques such as DES or AES, and (3) the application of renewable and unlinkable biometric references by means of a diversification process. Renewable and unlinkable biometric references correspond to techniques such as discrete fuzzy extractors [4] and continuous fuzzy extractors [1]. Continuous fuzzy extractors are also referred to as biometric cryptosystem [6], while the term biometric template protection often refers to the combination of *all* the *previous* mentioned countermeasures to provide confidentiality, renewability and authenticity for biometric references [11].

In this paper, we investigate the privacy enhancement introduced by a biometric cryptosystem assuming that an attacker has access to protected biometric references in at least two databases. The biometric references are assumed to be protected only by a renewable and preferably unlinkable diversification transform, and additional methods such as data separation or data confidentiality are not used. In our

model, the biometric references are protected against abuse in two ways. Firstly, a protected biometric template reveals almost nothing about the biometric characteristics of its owner and, if a database with protected biometric templates is compromised, the attacker cannot learn much about the compromised data. Secondly, if such an intrusion is detected the protected biometric references can be revoked and renewed, since at any time the protection scheme can be reapplied on the original or newly acquired data.

There are two classes of biometric cryptosystems techniques, which are fundamentally different. The first class considers biometric information as discrete variables (a collection of points) and has been formalized by Dodis *et al.* [4] in their definitions for fuzzy extractors and fuzzy sketches. The second class, considers biometric information as continuous variables (probability distributions, which describes the behavior of a user’s biometrics over time) and has been formalized by Linnartz and Tuyls [9] and Buhan *et al.* [1]. Both methods use a random, binary string to protect the biometric information. The result of this process is known as sketch and is considered to be public.

In this paper, we investigate and quantify the indistinguishability offered by a *continuous* biometric cryptosystems scheme. The scenario is the following: the attacker, Charlie, learns that a particular *protected biometric reference* belongs to Alice. This step is not particularly challenging for Charlie since it is assumed that protected biometric references are public. Now, Charlie would like to know what other accounts Alice has and what information is associated with these accounts. Therefore, the question we ask is:

*What is the probability that given a **protected** biometric reference that belongs to Alice, Charlie can find another protected biometric reference of Alice in a target database?*

OUR CONTRIBUTIONS. Our contributions are threefold. Firstly, we present a technique, which allows an adversary to match a *protected* biometric references generated using a *continuous* method, e.g., Quantization Index Modulation, (QIM) proposed by Linnartz and Tuyls [9] and extended by Buhan *et al.* [2]. Secondly, we quantify indistinguishability by means of the indistinguishability game proposed by Simoons *et al.* [12], and the limitations of this approach are outlined. Thirdly, an alternative, practical evaluation to quantify indistinguishability is described and results for real-world biometric data are provided based on the MCYT fingerprint database.

2 Preliminaries

NOTATION. By capital letters we denote random variables while small letters are used to denote observations of random variables. A random variable is completely described by its probability density function. A random variable A is endowed with a probability distribution $f_A(a)$. With A^d we denote the random variable endowed with a discrete probability distribution $f_{A^d}(a)$ while A^c is used to denote the random variable endowed with the continuous probability distribution $f_{A^c}(a)$. We use the random variable X when referring to a biometric identifier, which is represented as an m -dimensional feature vector. We assume that elements of the feature vector are independent and identically distributed, as commonly assumed in the biometric literature, more details on transformation techniques for biometric data can be found in Duda, *et. al* [5]. Subscripts are used for referring to components of a vector, while superscripts are use for enumerating elements of the same type.

The description of the QIM-fuzzy embedder is given for one generic feature element i , which in fact completely describes the whole process due to the independence assumption. The universe of all users with a given biometric identifier is denoted by \mathcal{U} . We use variable P when referring to public, protected biometric data, also referred to as a sketch. We use K to denote the key used to protect the biometric data.




Fig. 1. By quantization, $f_A(a)$ (continuous line) is transformed into $f_{Q(A)}(a)$ (dotted line). We can write $Q(f_A(a)) = f_{Q(A)}(a)$.




Fig. 2. Quantization of X with two scalar quantizers Q_0 (the set of X points) and Q_1 (the set of o points), corresponding to key bits k_0 and k_1 respectively both with step size q , gives the result, p_0 and p_1 respectively.

When referring to noise we use the variable N . We write $[x] = \begin{cases} [x], \{x\} \geq \frac{1}{2} \\ [x], \{x\} < \frac{1}{2} \end{cases}$ for every real number $x \in \mathbb{R}$, whereby $\{x\}$ we denote the fractional part of number x .

QUANTIZATION. A continuous random variable A can be transformed into a discrete random variable by means of quantization, which we write $Q(A)$. Formally, a quantizer is a function $Q : \mathcal{U} \rightarrow M$ that maps each $a \in \mathcal{U}$ into the closest *reconstruction point* in the set $M = \{c_1, c_2, \dots\}$ by

$$Q(a) = \arg \min_{c_i \in M} d(a, c_i) \quad (1)$$

where d is an appropriate distance measure defined on \mathcal{U} .

The *Voronoi region* or the *decision region* of a reconstruction point c_i is the subset of all points in \mathcal{U} , which are closer to that particular reconstruction point than to any other reconstruction point, with respect to a specific distance measure. We denote with V_{c_i} the Voronoi region of the reconstruction point c_i . When A is one dimensional, Q is called a *scalar quantizer*. If all Voronoi regions of a quantizer are equal in both size and shape the quantizer is *uniform*. In the scalar case, the length of the Voronoi region is then called the *step size*. If the reconstruction points form a lattice, the Voronoi regions of all reconstruction points are congruent. By quantization the probability density function of the continuous random variable A , $f_A(a)$ (which is continuous) is transformed into the probability density function $f_{Q(A)}(a)$ (which is discrete).

HIDING CODES FOR CONTINUOUS VARIABLES. Quantization based data hiding codes as introduced by Chen *et al.* [3] (also known as quantization index modulation) can embed secret information into a real-valued quantity. A *Quantization Index Modulation*, $QIM : \mathcal{U} \times K \rightarrow M$ data hiding scheme can be seen as a set of individual quantizers $\{Q_1, Q_2, \dots, Q_{2^m}\}$, where $2^m = |K|$ and each quantizer maps $x \in \mathcal{U}$ into a reconstruction point. The quantizer is chosen by the input value $k \in K$. We write $Q_k(x)$ to denote the quantization operation $QIM(x, k)$. The set of all reconstruction points is $M = \bigcup_{k \in K} M_k$ where $M_k \subset M$ is the set of reconstruction points of the quantizer Q_k , and k is known as the *label* of the reconstruction point $Q_k(x)$.

The amount of tolerated noise or the reliability is determined by the minimum distance between two neighboring reconstruction points. The size and shape (for high dimensional quantization) of the Voronoi region determines the error tolerance. For the scalar quantizer in the previous example

$$Q_k(x) = Q_k(y) \text{ when } d(Q_k(x), y) \leq \frac{q}{2}$$

The number of quantizers in the QIM set determines the amount of information that can be embedded. By setting the number of quantizers and by choosing the shape and size of the decision region the performance properties can be finely tuned, more details can be found in Buhan, *et al.* [2].

3 QIM Biometric Cryptosystem

A biometric cryptosystem aims to protect the stored biometric identity of a user from abuse in two ways. Firstly, a protected biometric reference should reveal almost nothing about the underlying biometric and if a database with protected biometric reference is compromised, the attacker cannot learn much about the biometric. Secondly, if such an intrusion is detected the protected biometric reference can be revoked and renewed, since at any time the protection scheme can be reapplied on the original data.

The main challenge in protecting biometric references using cryptographic techniques is coping with noise, which is always introduced into biometric samples during data acquisition and processing. Biometric cryptosystems can transform a noisy, biometric measurement represented as a sequence of non-uniformly distributed real numbers into a reproducible, uniformly-distributed binary string. There are many parameters that control this transformation, for example the length of the output binary sequence, the probability that two measurements coming from the same users will be mapped to the same binary sequence, etc.

Two abstractions, secure sketches and fuzzy extractors were proposed by Dodis, *et al.* [4] to describe the process of transforming a biometric characteristic into a reproducible, uniform binary sequence. A secure sketch can correct the noise between two biometric measurements coming from the same user by using some public information called a sketch. The result of a secure sketch is a reproducible sequence, which is not, necessarily, uniformly distributed and thus not suitable to be used as cryptographic keys. Fuzzy extractors can be used to extract randomness from biometric data to make the output of a secure sketch suitable for usage as cryptographic keys. Both constructions work only on biometric data represented as discrete variables. The process of transforming a continuous variable into a discrete variable influences the performance of fuzzy extractors and secure sketches.

Fuzzy embedders were proposed by Buhan, *et al.* [2] as an extension to the fuzzy extractor idea. A fuzzy embedder can transform a noisy, non-uniform continuous variable, into a reproducible, uniformly random string, which is suitable to be used as a cryptographic key. Basically, the function of a fuzzy embedder is the same as the function of the fuzzy extractor, but its scope is extended so as to accept continuous variables as input. A fuzzy embedder is a pair of procedures. The first is the embed procedure, which is used once when the biometric data is protected and stored on an untrusted server. The second is the reproduce procedure, which is used to authenticate the user to the server.

QIM BIOMETRIC CRYPTOSYSTEM. Linnartz *et al.* [9] were the first to suggest how to use QIM for the protection of biometric data. The main advantage is that quantization (or discretization) of the biometric data is not required, since QIM works on continuously represented data. Li, *et al.* [8] argue that performance measures like min-entropy or entropy-loss are the result of the quantization parameters used. The larger the quantization step, the less entropy is left in the discrete biometric data and the easier it is to reconstruct the secret k vice-versa the smaller the quantization step, the more entropy remains in the discrete biometric and the harder it is to reconstruct k .

Definition 1 (QIM-fuzzy embedder [2]). A $(\mathcal{U}, X, K, \eta, m, q)$ - QIM-fuzzy embedder is a pair of randomized procedures $\langle \text{Embed}, \text{Reproduce} \rangle$ where

- Embed is a function used during enrollment that outputs a sketch $p \in [-\frac{q}{2}, \frac{q}{2}]^m$ on input $k \in K$ and $x \in X$;
- Reproduce is a function used during verification that given a word x' and any sketch $p = \text{Embed}(x, k)$ outputs k as long as $d(x_i, x'_i) \leq \frac{q}{2}, (\forall) i \in \{1, m\}$.

For any random variable X over \mathcal{U} the probability that an adversary who observes P guesses X is at most $\eta = I(X; P)$

For QIM the enrollment phase consists of a three step procedure that is applied on each feature vector component x_i separately as shown in Table 1.

Enrollment:	
1.	Generate: $k_i \in \{0, 1\}$;
2.	Apply: $\text{Embed}(x_i, k_i) = Q_{k_i}(x_i) - x_i = p_i$;
3.	Publish: p_i ;
Verification:	
1.	Reproduce(x'_i, p_i) = k'_i , where k'_i is the label of the reconstruction point $\lceil \frac{x'_i + p_i}{q} \rceil q$
2.	If $k_i = k'_i$ accept, otherwise reject;

Table 1. Enrollment and verification algorithm for the QIM, biometric cryptosystem. We observe that the biometric keys k_i and k'_i will be exactly the same as long as $d(x_i, x'_i) \leq \frac{q}{2}$.

During authentication, a noisy biometric feature vector $x' = (x'_1, x'_2, \dots, x'_m)$ is collected. Verification of a user is performed by reproducing each bit of the biometric key, k_i from the biometric measurement x'_i and the corresponding sketch p_i . The reproduction procedure finds the closest reconstruction point for $Q(x'_i + p_i) \in M$ and returns the label, 0 or 1, associated with this point. The decision to accept or reject a user is done by comparing the obtained key, k' to the enrollment key, k .

Example 1. We want to hide one bit of information, $k \in \{0, 1\}$, into the real value x_i . For this purpose we use a scalar uniform quantizer with step size q , given by rounding x_i to the closest reconstruction point. The public sketch is computed as:

$$Q(x_i) = q \left\lceil \frac{x_i}{q} \right\rceil.$$

The quantizer Q is used to generate a set of two new quantizers $\{Q_0, Q_1\}$ defined as:

$$Q_0(x_i) = n(x_i)q \quad \text{and} \quad Q_1(x_i) = (n(x_i) + \frac{1}{2})q.$$

In Figure 2 the reconstruction points for the quantizer Q_1 are shown as circles and the reconstruction points for the quantizer Q_0 are shown as crosses. The embedding is done by mapping the point x_i to one of the reconstruction points of these two quantizers. For example, if $k = 1$, x_i is mapped to the closest \circ point. Therefore,

$$p_0 = \text{Embed}(x_i, k = 0) = Q_0(x_i) - x_i \quad \text{and} \quad p_1 = \text{Embed}(x_i, k = 1) = Q_1(x_i) - x_i$$

where $n(x_i) \in \mathbb{Z}$ is chosen such that $|Q_k(x_i) - x_i| \leq \frac{q}{2}$. The result of the embedding is the distance vector to the nearest \times or \circ as chosen by k . During the reproduction procedure x_i is perturbed by noise then quantizer will assign the received data to the closest \times or \circ point, and output 0 or 1 respectively. The set of the two quantizers $\{Q_0, Q_1\}$ is called a QIM.

Definition 2 (Related Sketches). Let $(\mathcal{U}, X, K, \eta, m, q)$ be a QIM-fuzzy embedder. We say that $p_x = \text{Embed}(x, k)$ and $p'_x = \text{Embed}(x', k')$ are related sketches as long as $d(x_i, x'_i) \leq \frac{q}{2}, (\forall) i \in \{1, m\}$ for any pair $\{k, k'\} \in K$.

4 A theoretical measure of indistinguishability for the QIM fuzzy embedder

n-INDISTINGUISHABILITY. The aim of a biometric cryptosystem, which features the *n*-indistinguishability attribute as defined by Simoens *et al.* [12] is that no adversary has a significant advantage over random guessing in determining whether *n* sketches $\{P_1, P_2, \dots, P_n\}$ are related or not.

Simoens *et al.* [12] model *n*-indistinguishability as a game where it is assumed that an adversary has obtained a database of protected biometric references and wants to find the sketches that are related to the reference he holds. As it is the customary in cryptography, the adversary is assumed to know all algorithms used to protect the biometric references.

2-INDISTINGUISHABILITY (Simoens *et al.* [12]). For completeness, we give the description of the game, for two sketches ($n = 2$) below.

1. The challenger randomly selects the variable $X \in \mathcal{U}$ and samples X to obtain $x \in X$. He also selects a secret key $k^{(1)} \in K$ and gives the output of the embed procedure, the sketch P , to the adversary.
2. The challenger flips a fair coin $\mathbf{c} \in \{0, 1\}$. If $\mathbf{c} = 1$, the challenger samples variable X again to obtain x' . If $\mathbf{c} = 0$, the challenger selects another random variable $Y \in \mathcal{U}$ and samples Y to obtain $y \in Y$. Regardless of the result of the coin flip, the challenger selects a new secret key $k^{(2)}$ and gives the output of the embed procedure P' , to the adversary.
3. The adversary's aim is to guess correctly whether P' comes from x or y . In particular, the adversary outputs a single bit $\hat{\mathbf{c}} \in \{0, 1\}$ and he wins the game if $\hat{\mathbf{c}} = \mathbf{c}$.

The advantage of the adversary in the indistinguishability game is defined as:

$$Adv_{2\text{-IND}} = 2 \left| \Pr[\hat{\mathbf{c}} = \mathbf{c}] - \frac{1}{2} \right| = 2 \left| \Pr[\hat{\mathbf{c}} \neq \mathbf{c}] - \frac{1}{2} \right|$$

Notice that we model biometrics as an m -dimensional variable and therefore an adversary who guesses $\hat{\mathbf{c}} = \mathbf{c}$ has to make m correct guesses: $(\hat{\mathbf{c}}_1 = \mathbf{c}_1) \wedge (\hat{\mathbf{c}}_2 = \mathbf{c}_2) \wedge \dots \wedge (\hat{\mathbf{c}}_m = \mathbf{c}_m)$ one for every component of the public sketch. As we made the assumption that components in the features vector are independent we can write:

$$\Pr[\hat{\mathbf{c}} = \mathbf{c}] = \prod_{i=1}^m \Pr[\hat{\mathbf{c}}_i = \mathbf{c}_i]$$

Without loss of generality, in the rest of this section we concentrate on evaluating the advantage of the adversary in the indistinguishability game for one correct guess of the form $\hat{\mathbf{c}}_i = \mathbf{c}_i$, and all definitions are given for a $(\mathcal{U}, X, K, \eta, 1, q)$ QIM-fuzzy embedder. The adversary in the above game is called $\text{Charlie}_{\text{IND}}$ and his advantage in the game is defined as:

$$Adv_{2\text{-IND}}^i = 2 \left| \Pr[\hat{\mathbf{c}}_i = \mathbf{c}_i] - \frac{1}{2} \right| = 2 \left| \Pr[\hat{\mathbf{c}}_i \neq \mathbf{c}_i] - \frac{1}{2} \right| \quad (2)$$

Definition 3 (ϵ -Indistinguishability). An $(\mathcal{U}, X, K, \eta, 1, q)$ QIM-fuzzy embedder $\langle \text{Embed}, \text{Reproduce} \rangle$ is ϵ -indistinguishable if for any adversary $\text{Charlie}_{\text{IND}}$, such that $Adv_{2\text{-IND}} = Adv_{\text{Charlie}_{\text{IND}}}$ it holds that $Adv_{2\text{-IND}}^i \leq \epsilon$.

Definition 4 (QIM-Distinguisher). For any two sketches p_{x_i} and p_{y_i} generated by an $(\mathcal{U}, X, K, \eta, 1, q)$ QIM-fuzzy embedder $\langle \text{Embed}, \text{Reproduce} \rangle$ the function \mathcal{H}^δ , defined as:

$$\mathcal{H}^\delta(p_{x_i}, p_{y_i}) = \begin{cases} 1, & \text{if } |p_{x_i} - p_{y_i}| \leq \delta, \text{ or } |p_{x_i} - p_{y_i} - \frac{q}{2}| \leq \delta, \text{ or } |p_{x_i} - p_{y_i} + \frac{q}{2}| \leq \delta; \\ 0, & \text{otherwise.} \end{cases}$$

is a QIM-distinguisher.

A few explanations are in order to motivate the introduction of the parameter δ in the definition of the distinguisher. For an average user Alice, the distance between two random samplings x and x' of variable X is at most $\frac{q}{2}$, $d(x_{\text{Alice}}, x'_{\text{Alice}}) \leq \frac{q}{2}$. However, if Charlie, knows that the biometric data of the user he is targeting for cross-matching (Dave) is better (less noise between different samplings) compared to that of the average user (Alice), Charlie has an additional advantage. We model this advantage by the introduction of the parameter δ . By choosing a value δ Charlie has control over the distance between two biometric measurements of Dave, $d(x_{\text{Dave}}, x'_{\text{Dave}}) \leq \delta \ll \frac{q}{2}$, for example.

Lemma 1 (Distinguishing related sketches.). *Let x_i and x'_i be two samples of random variable X , furthermore let $x'_i = x_i + \delta_i$, with $|\delta_i| \leq \frac{q}{2}$. For any, two related sketches p_{x_i} and $p_{x'_i}$ generated by an $(\mathcal{U}, X, K, \eta, 1, q)$ QIM-fuzzy embedder $\langle \text{Embed}, \text{Reproduce} \rangle$ the QIM-distinguisher always outputs the value 1.*

Proof. To make the proof more readable, we firstly analyze the simple case when the sampling of variable X , yields $x_i = x'_i$. This case corresponds to the scenario when there is no noise between different enrollment samples of user X . Secondly we extend the simple case to the scenario when different enrollment samples, x and x' of user X are subjected to noise, $d(x_i, x'_i) \leq \delta$. For both cases we derive the value of the difference $p_{x_i} - p_{y_i}$ when the two sketches are related and we show that $\mathcal{H}^{\delta=0}(p_{x_i}, p_{y_i})$ and $\mathcal{H}^{\delta}(p_{x_i}, p_{y_i})$ are equal to 1 in both cases.

SIMPLE CASE ($x_i = x'_i$). Although the different keys $k^{(1)}$ and $k^{(2)}$ are used to generate sketches for x and x' respectively, we discovered there is a simple test to verify whether the resulting sketches $p^{(x)} = (p_{x_1}, p_{x_2}, \dots, p_{x_m})$ and $p^{(x')} = (p_{x'_1}, p_{x'_2}, \dots, p_{x'_m})$ are related. Each element p_{x_i} and $p_{x'_i}$ of the public sketches is computed as: $p_{x_i} = \text{Embed}(x_i, k_i^{(1)}) = Q_{k_i^{(1)}}(x_i) - x_i$ and $p_{x'_i} = \text{Embed}(x'_i, k_i^{(2)}) = Q_{k_i^{(2)}}(x'_i) - x'_i$, where quantization is defined in equation (1). In deriving a distinguishing function, the adversary can distinguish three cases:

Case I: The result of the coin flip is $c_i = 1$, ($x_i = x'_i$) and the key bits are equal ($k_i^{(1)} = k_i^{(2)}$). By subtracting the two sketches the adversary obtains:

$$p_{x_i} - p_{x'_i} = (Q_{k_i^{(1)}}(x_i) - x_i) - (Q_{k_i^{(1)}}(x_i) - x_i) = 0;$$

Case II: The result of the coin flip is $c_i = 1$, ($x_i = x'_i$) however the key bits are different ($k_i^{(1)} \neq k_i^{(2)}$). By subtracting the two sketches the adversary obtains:

$$\begin{aligned} p_{x_i} - p_{x'_i} &= Q_{k_i^{(1)}}(x_i) - x_i - (Q_{k_i^{(2)}}(x_i) - x_i) \\ &= Q_{k_i^{(1)}}(x_i) - Q_{k_i^{(2)}}(x_i) = \pm \frac{q}{2} \end{aligned}$$

Figure 2 shows that embedding two different bits in the same value will lead always to sketches that are complementary, $p_0 - p_1 = \pm \frac{q}{2}$.

Case III: The result of the coin flip is $c_i = 0$ and $x_i \neq x'_i$ when subtracting two sketches the result is different from 0 or $\pm \frac{q}{2}$.

To summarize, by subtracting p_{x_i} and p_{y_i} we obtain

$$p_{x_i} - p_{y_i} = \begin{cases} 0, & \text{if } (c_i = 1) \wedge (k_i^{(1)} = k_i^{(2)}) \\ \pm \frac{q}{2}, & \text{if } (c_i = 1) \wedge (k_i^{(1)} \neq k_i^{(2)}) \\ p_{x_i} - p_{y_i} & \text{if } (c_i = 0); \end{cases}$$

Therefore when p_{x_i} and p_{y_i} are related $|p_{x_i} - p_{y_i}| \in \{0, \frac{q}{2}\}$ and $\mathcal{H}^{\delta=0}(p_{x_i}, p_{y_i}) = 1$.

GENERAL CASE ($d(x_i, x'_i) \leq \delta$). During enrollment, multiple measurements for the same individual are taken. The average of these measurements is computed and stored as reference information. Due to the unpredictable amount of noise existent in each measurement the reference information changes as well. The extended case, models this scenario by assuming that the biometric reference information of person X gives two different reference values x_i and x'_i that are within distance δ and therefore we set $d(x_i, x'_i) \leq \delta$. Derivation of function $|p_{x_i} - p_{y_i}|$ is straightforward, by replacing $x' = x + \delta$ in the three cases derived in the previous paragraph. As a result we obtain:

$$|p_{x_i} - p_{y_i}| \in \begin{cases} (-\delta, \delta), & \text{if } (\mathbf{c}_i = 1) \wedge (k_i^{(1)} = k_i^{(2)}) \\ (-\frac{q}{2} - \delta, -\frac{q}{2} + \delta) \cup (\frac{q}{2} - \delta, \frac{q}{2} + \delta) & \text{if } (\mathbf{c}_i = 1) \wedge (k_i^{(1)} \neq k_i^{(2)}) \\ p_{x_i} - p_{y_i} & \text{if } (\mathbf{c}_i = 0); \end{cases} \quad (3)$$

Therefore when p_{x_i} and p_{y_i} are related $\mathcal{H}^{\delta}(p_{x_i}, p_{y_i}) = 1$. □

In fact, x_i and x'_i can be samples from different distributions, as long as the conditions of lemma 1 are satisfied, the distinguisher returns 1.

Lemma 2 (ϵ -Indistinguishability for QIM- fuzzy embedder). An $(\mathcal{U}, X, K, \eta, 1, q)$ QIM-fuzzy embedder $\langle \text{Embed}, \text{Reproduce} \rangle$ is ϵ -indistinguishable for any adversary $\text{Charlie}_{\text{IND}}$, and it holds that:

$$\left| \int_{\Delta} f_{D_{P_i}}(p(t))dt - 1 \right| \geq \epsilon$$

where $\Delta = (-\frac{q}{2} - \delta, -\frac{q}{2} + \delta) \cup (-\delta, \delta) \cup (\frac{q}{2} - \delta, \frac{q}{2} + \delta)$ and $f_{D_{P_i}}$ is the probability distribution of the difference between $P_{x_i} - P_{y_i}$ and P_{x_i}, P_{y_i} are the random variables from which p_{x_i} and p_{y_i} are sampled.

Proof. If the adversary uses \mathcal{H}^{δ} for guessing the challengers coin flip he will always guess that p_{x_i}, p_{y_i} are not related when $\mathcal{H}^{\delta}(p_{x_i}, p_{y_i}) = 0$, regardless of the coin flip. The adversary also guesses that p_{x_i}, p_{y_i} are related when $\mathcal{H}^{\delta}(p_{x_i}, p_{y_i}) = 1$ and $\mathbf{c}_i = 1$. The adversary makes an incorrect guess when the coin flip is $\mathbf{c}_i = 0$, (the sketches are not related) but $\mathcal{H}^{\delta}(p_{x_i}, p_{y_i}) = 1$. It follows that the probability of an incorrect guess can be derived from :

$$\Pr[\hat{\mathbf{c}}_i \neq \mathbf{c}_i] = \Pr[\hat{\mathbf{c}}_i = 0 | \mathbf{c}_i = 1] \Pr[\mathbf{c}_i = 1] + \Pr[\hat{\mathbf{c}}_i = 1 | \mathbf{c}_i = 0] \Pr[\mathbf{c}_i = 0]$$

According to lemma 1, the distinguisher always returns 1 when the sketches are related (FRR=0). Therefore, the probability that the adversary guesses $\hat{\mathbf{c}}_i = 0$, when $\mathbf{c}_i = 1$, is 0. Therefore $\Pr[\hat{\mathbf{c}}_i = 0 | \mathbf{c}_i = 1] = 0$. The probability that the adversary guesses $\hat{\mathbf{c}}_i = 1$ when $\mathbf{c}_i = 0$ is:

$$\Pr[\hat{\mathbf{c}}_i = 1 | \mathbf{c}_i = 0] = \Pr[\mathcal{H}^{\delta}(p_{x_i}, p_{y_i}) = 1 | \mathbf{c}_i = 0]$$

When knowing the probability distribution of the sketch P_i , denoted by $f_{P_i}(p)$, we can compute probability distribution for the difference between variables $D_{P_i} = P_{x_i} - P_{y_i}$ (as an exercise, in Appendix A we compute $f_{D_{P_i}}$ when $f_{\mathcal{U}}(u)$ is modeled as a normal distribution with mean $\mu = 0$ and variance




Fig. 3. $\Pr[\hat{c}_i \neq c_i]$ for $q \in \{0.5, 1, 1.5, 2, 2.5\}$ when variable X is sampled from $\mathcal{U} = N(0, 1)$.




Fig. 4. ϵ -Indistinguishability for an $\mathcal{U}, X, \eta, 1, q$ -fuzzy embedder for $q \in \{0.5, 1, 1.5, 2, 2.5\}$ and $\mathcal{U} = N(0, 1)$.

$\sigma^2 = 1$), we can write:¹

$$\Pr[\mathcal{H}^\delta(p_{x_i}, p_{y_i}) = 1 | \mathbf{c}_i = 0] = \int_{\Delta} f_{D_{P_i}}(p(t)) dt$$

The probability that the adversary makes an incorrect guess is therefore:

$$\Pr[\hat{c}_i \neq c_i] = \frac{1}{2} \int_{\Delta} f_{D_{P_i}}(p(t)) dt. \quad (4)$$

Substitutions equation (4) in equation (2), proves the values is $\epsilon = \left| \int_{\Delta} f_{D_{P_i}}(p(t)) dt - 1 \right|$. An adversary with an improved strategy or a superior distinguisher has an advantage that is larger compared to ϵ which completes our proof. \square

EVALUATION OF THE ADVERSARY ADVANTAGE. Figure 4 shows the bounds on ϵ -Indistinguishability for an $(\mathcal{U}, X, \eta, 1, q)$ fuzzy embedder for several quantization steps q , when the probability distribution of \mathcal{U} is modeled as a normal distribution with $\mu = 0$ and $\sigma = 1$. We conclude that the advantage of Charlie in the 2-Indistinguishability game is significant and it increases with the quantization step (which is preferred in practice as it increases the classification performance of the biometric authentication). Another interesting observation is that, Charlie has a larger advantage when targeting a person which has very stable biometric identifiers. This means that an adversary can easily identify protected biometric references which come from the same person and the more stable the biometric the more accurate is his identification (Dave vs. Alice).

In the next section we use the distinguisher \mathcal{H}^δ , as defined in this section to execute a cross-matching attack on real biometric data.

¹ We note that when evaluating $\Pr[\hat{c}_i = 0 | \mathbf{c}_i = 1]$ one should condition the output on the values of the key bits, therefore have $\Pr[\hat{c}_i = 0 | \mathbf{c}_i = 1] = \Pr[\hat{c}_i = 0 | \mathbf{c}_i = 1, k_i^{(1)} = k_i^{(2)}] \Pr[k_i^{(1)} = k_i^{(2)}] + \Pr[\hat{c}_i = 0 | \mathbf{c}_i = 1, k_i^{(1)} \neq k_i^{(2)}] \Pr[k_i^{(1)} \neq k_i^{(2)}]$ However, when the sketches are not related the probability that the function \mathcal{H}^δ gives a 0 or a $\lfloor \frac{q}{2} \rfloor$ is not influenced by the key bits.

5 A practical measure of indistinguishability for the QIM fuzzy embedder

Although the concept of n -indistinguishability is very suitable to describe the theoretical advantage of an attacker, it has its limitations. The concept describes the advantage of an attacker with respect to a perfectly indistinguishable system. On the one hand, perfect indistinguishability is hard to achieve when employing only biometric cryptosystems due to the inherent correlation of the data used to generate related biometric references. On the other hand, indistinguishability is *not* hard to achieve when biometric template protection techniques are employed. The protected biometric reference is encrypted and the encryption key is stored outside the database, for example on a smartcard. A user who wishes to verify his identity uses the key stored on the smartcard to decrypt the sketch and then proceeds to perform biometric authentication. The ciphertext indistinguishability of most encryption schemes guarantees that the biometric sketches achieve indistinguishability. Therefore, in this section we investigate the other side of the indistinguishability game presented in the previous section, namely *is there a gain in privacy with respect to cross-matching, when using a biometric cryptosystem?*

To answer this question we compare the classification performance of the MCYT fingerprint database in two distinct scenarios. *Scenario I* models the classification performance of the QIM-fuzzy embedder during normal operation. User x is first enrolled in the biometric system and the public sketch p_x is computed as $p_x = \text{Encode}(x, k)$. During verification the user presents his biometric x' and the server computes $\text{Reproduce}(x', p_x) = k'$. Authentication is considered successful if the Hamming distance between k and k' is zero. For a QIM-fuzzy embedder the percentage of successful authentication depends on the distance tolerated between x and x' , which is a function of the quantization step, q , given by $d(x, x') \leq \frac{q}{2}$. *Scenario II* corresponds to the scenario when the adversary has access to the protected biometric references, $p_x = \text{Embed}(x, k_x)$ and $p_y = \text{Embed}(y, k_y)$. In this case, classification performance is evaluated using the distinguisher function, $\mathcal{H}^\delta(p_x, p_y)$ constructed in Section 4.

We propose to use cross-matching performance differences between unprotected and protected biometric references as relevant measure for indistinguishability. These properties can be described by a receiver operating characteristic curve (ROC) indicating false match and false non-match rates. In this section, for the evaluation we use the MCYT database [10], which is known in the literature as a good quality data set and has good classification performance. In the current context of testing for indistinguishability a good quality database is rather a pessimistic choice. We expect that the sketch classification performance improves with good quality data (less noise expected between biometric references collected from the same person).

DATA SET DESCRIPTION. The MCYT database consists of fingerprints collected from 323 individuals. For each individual, 12 fingerprint images have been captured under the supervision of an operator. Fingerprint images were collected with an optical sensor (Digital Persona), which gives as output images having resolution of 256×400 pixels and 8 bit gray-scale levels. From the total of 323 individuals, 80% are used for training the algorithm and the rest of 20% (approximately 66 individuals) are used for testing the performance of the algorithm. During testing, the data is split into two sets. The first set consisting of 4 fingerprints are used as enrollment data. The second set consisting of 8 fingerprints is used as verification set. For each round of experiments 5 random splits are performed on the testing data and the results are averaged.

Each fingerprint in the database is processed and represented as a fixed length vector. To describe the shape of the fingerprint, two types of features are extracted. The first feature vector is the squared directional field and the second feature vector is the Gabor response of the fingerprint, details can be found in Tuyls *et al.* [13]. The resulting feature vector is a concatenation of the squared directional field and the Gabor response and describes the global shape of the fingerprint in 1536 elements. Prior to applying




Fig. 5. (Scenario I) Classification performance when the QIM biometric cryptosystem is applied on the MCYT fingerprint database. Testing data is split into two sets. The first set (4 measurements) is used during enrolment while the second set (8 measurements) is used during verification. Several ROC curves are obtained by varying the length of the feature vector m , and the quantization step q . The classifier used is the Hamming distance between k and $k' = \text{Reproduce}(x, p_x)$, where $p_x = \text{Embed}(x, k)$. The EER value, expressed in percent, for each of the experiments is plotted, as a function of the amount of noise tolerated between biometric samples of the same individual, $d(x, x') \leq \frac{q}{2}$.

QIM, Principal Component Analysis (PCA) and Linear Discriminant Analysis (LDA) transformations are applied on each continuous-domain feature vector to reduce its dimensionality to the desired length while maintaining maximum discrimination. The PCA and LDA parameters are obtained from the training set. The enrollment feature vector is constructed by averaging the set of enrollment feature vectors of 4 measurements.

ERROR RATES. There are two dual measures in the biometric literature to measure resilience to noise. The first is *False Rejection Rate* (FRR), which estimates the probability that the public sketch of person A and a measurement of person A produce a faulty secret key. The second measure is *False Acceptance Rate* (FAR), which estimates the probability that a public sketch of person A and a measurement of person B produce the correct secret key of person A . For QIM the factors that influences the FAR and FRR values (besides the quality of the data, which is determined by the impostor versus genuine standard deviation) are: (a) the quantization step q , which determines the amount of noise tolerated between biometric measurements of the same individual, $d(x, x') \leq \frac{q}{2}$ and (b) the number of features m that are used. Different values for the FAR and FRR are obtained by varying the maximum accepted Hamming distance between measurements coming from the same person. This curve is called *Receiver Operating Characteristic* curve (ROC). The point where the FAR and the FRR are equal is known as the *Equal Error Rate* (EER) and is used as reference point.

SCENARIO I. The first set of experiments, corresponds to *Scenario I* and measures the performance of the biometric recognition algorithm in a classical use scenario: user x is first enrolled in the biometric system and the public sketch p_x is computed as $p_x = \text{Encode}(x, k)$. During verification the user presents his biometric x' and the server computes $\text{Reproduce}(x', p_x) = k'$. Figure 5 shows the EER for different quantization steps (q) and different number of features (m), when the Hamming distance is used to compute the distance between $\text{Reproduce}(x, p_x)$ and $\text{Reproduce}(x', p_x)$. As expected, the smaller the




Fig. 6. (Scenario II) Classification performance for the biometric sketches produced by the QIM-fuzzy embedder applied on the MCYT fingerprint database. Testing data is split into two databases (4 measurements each) and sketches are generated for each biometric reference. The classifier used in this case is the QIM-distinguisher $\mathcal{H}^{\frac{q}{4}}(p_x, p_y)$ given in definition 4. The EER values are obtained by varying the length of the feature vector m , and the size of quantization step q . The EER value, expressed in percentage, for each of the experiments is plotted, as a function of the amount of noise tolerated between biometric samples of the same individual, $d(x, x') \leq \frac{q}{2}$.

quantization step, the less noise is tolerated and the higher the EER. For example the EER goes up from 5.38% for $q = 3.5$ to 8.03% for $q = 1$ (for $m = 150$ features). Also, the more features, the less accurate the classification performance becomes. For example, for $m = 200$ features the EER is approximately 5.5% while for $m = 50$ the error rates are significantly lower, approximately 2.73% (for $q=3.5$). The reason for plotting the curves in *Figure 5* is to have a reference for the classification performance of the public sketches.

SCENARIO II. The second set of experiments, correspond to *Scenario II* and measure the performance of the cross-matching algorithm. The classifier used in this case is the distinguishing function $\mathcal{H}^{\delta}(p_x, p_y)$, where p_x is a sketch found in the first database and p_y is a sketch found in the second database. The two databases are obtained by randomly splitting the MCYT database. The result of $\mathcal{H}^{\delta}(p_x, p_y)$, is a binary string of length m (the number of components in p_x and p_y , respectively) where each bit is obtained from equation (3). As shown in *Figure 6* the EER goes down from 28.4% for $q = 1$ to only 1.86% for $q = 3.5$ ($m = 200$ features). Also, the more features are used, the less accurate the classification performance becomes. For example the 1.15% EER obtained for $m = 50$ features increases to 5.03% EER for $m = 200$ features (for $q = 2$).² The results were obtained for setting $\delta = \frac{q}{4}$, see *Section 4*. As expected it seems that the same settings that improve the classification performance of the QIM biometric cryptosystem improve the sketch classification performance. In other words linking users across databases becomes easier when the biometric classification performance improves. Comparing the classification results obtained in *Scenario I* and *Scenario II* we conclude that employing biometric cryptosystems improves, the biometric sketch indistinguishability in most cases. For example, for $m=200$ features, the classification performance for $q = 1$ is 9% in *Scenario I* decreases to 28.40%, for the same quantization step, and $m=100$ features, the classification performance is 7.05% in *Scenario I* and 22.56% in *Scenario II*.

² We note that a classification performance of 1.15% means that an attacker can guess with 98.85% probability whether two sketches are related or not.

The surprising result of these experiments is that, not only that indistinguishability is not achieved for all quantization steps but in some cases the sketch classification performance (*Scenario 2*), see *Figure 6* offers better performance compared to the biometric classification (*Scenario 1*), see *Figure 5*. We explain this phenomenon by the fact that in the classical biometric classification a 4:1 matching (4 measurement used during enrollment and 1 measurement used for authentication) is employed. On the other hand matching sketches represent a 4:4 comparison (4 measurements are used when computing the sketches during enrollment) and thus sketch classification is less corrupted by noise. Kelkboom, *et al.* [7] show that a 4:4 matching (4 enrollment measurements:4 verification measurements) has superior classification results compared to a 4:1 matching (4 enrollment measurements:1 verification measurement). This supports, from a theoretical perspective the results we obtained in practice. We consider these settings to be realistic as current practice in the field is to collect multiple samples during enrollment and less samples during verification.

The main conclusion is positive in the sense that biometric cryptosystem have a positive effect on privacy, with respect to cross-matching, however we seem to have a lot to improve in this sense.

6 Conclusions

Privacy compliant databases should ensure that users are indistinguishable. In this paper, we show how an adversary can distinguish between protected biometric references generated with the QIM-fuzzy embedder. In this context we show that the advantage of an adversary who plays the indistinguishability game is non-negligible. Secondly, we look at the indistinguishability property from a practical perspective. We first randomly split the MCYT fingerprint database into two databases such that each user in the MCYT database can be found both databases. We apply the QIM fuzzy embedder for each user using different random sequences to protect the reference biometric samples in the two different databases. On the protected references we apply the distinguishing function, to determine whether they belong to the same user or not. As the performance of the cross-matching attack depends on the amount of noise tolerated between different samplings of a users biometric we compare the results to the error rates of the normal operation point. The results indicated that the QIM method does provide a certain amount of cross-matching resilience, but at the same time does not meet the desired requirement of complete unlinkability (and hence indistinguishability) when using the MCYT fingerprint database.

7 Acknowledgements

We would like to thank Willem Jonker and Raymond Veldhuis for their support in presenting this work.

References

1. I.R. Buhan, J. Doumen, P.H. Hartel, and R.N.J. Veldhuis. Fuzzy extractors for continuous distributions. In R. Deng and P. Samarati, editors, *Proceedings of the 2nd ACM Symposium on Information, Computer and Communications Security (ASIACCS), Singapore*, pages 353–355, New York, March 2007. ACM.
2. I.R. Buhan, J. Doumen, P.H. Hartel, and R.N.J. Veldhuis. Embedding renewable cryptographic keys into continuous noisy data. In *(to appear) 10th International Conference on Information and Communications Security (ICICS)*, Lecture Notes in Computer Science, Birmingham, UK, October 2008. Springer-Verlag.
3. B. Chen and G.W. Wornell. Quantization index modulation methods for digital watermarking and information embedding of multimedia. *The Journal of VLSI Signal Processing, Springer Netherlands*, 27(1-2):7–33, February 2001.
4. Y. Dodis, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology ,Proceedings of International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT 2004), Interlaken, Switzerland*, volume 3027 of *Lecture Notes in Computer Science*, pages 523–540. Springer, May 2004.

5. R.O. Duda, P.E. Hart, and D.G. Stork. *Pattern Classification (2nd Edition)*. Wiley-Interscience, October 2000.
6. A.K. Jain, K. Nandakumar, and A. Nagar. Biometric template security. *Journal on Advances in Signal Processing (EURASIP)*, 2008:17, 2008.
7. E. Kelkboom, G. Garcia Molina, J. Breebaart, T.A.M. Kevenaar, R.N.J. Veldhuis, and W. Jonker. Binary biometrics: An analytic framework to estimate the performance curves under gaussian assumptions. In *IEEE Transactions on Systems, Man and Cybernetics (to appear)*, 2009.
8. Q. Li, Y. Sutcu, and N. Memon. Secure sketch for biometric templates. In Xuejia Lai and Kefei Chen, editors, *Advances in Cryptology 12th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT 2006)*, Shanghai, China, volume 4284 of *Lecture Notes in Computer Science*, pages 99–113. Springer, December 2006.
9. J.P. Linnartz and P. Tuyls. New shielding functions to enhance privacy and prevent misuse of biometric templates. In Josef Kittler and Mark S. Nixon, editors, *4th International Conference on Audio-and Video-Based Biometric Person Authentication (AVBPA 2003)*, Guildford, UK, volume 2688 of *Lecture Notes in Computer Science*, pages 393–402. Springer, June 2003.
10. J. Ortega-Garcia, J. Fierrez-Aguillar, D. Simon, J. Gonzalez, M. Faundez-Zanuy, V. Espinosa, A. Satue, I. Hernaez, J.-J. Igarza, C. Vivaracho, D. Escudero, and Q.-I. Moro. Myct baseline corpus: a bimodal biometric database. In *IEEE Proceedings on Vision, Image and Signal Processing, Special Issue on Biometrics on the Internet*, volume 150, pages 395–401. IEEE Computer Society, 2003.
11. ISO/IEC JTC1 SC27. CD 24745 - information security - biometric template protection.
12. K. Simoens, P. Tuyls, and B. Preneel. Privacy weakness in biometric sketches. In *IEEE Symposium on Security and Privacy, Oakland, California, USA*, May 2009.
13. P. Tuyls, A. Akkermans, T. Kevenaar, G. Schrijen, A. Bazen, and R. Veldhuis. Practical biometric authentication with template protection. In Takeo Kanade, Anil K. Jain, and Nalini K. Ratha, editors, *Proceedings of the 5th International Conference on Audio- and Video-Based Biometric Person Authentication (AVBPA 2005)*, Hilton Rye Town, NY, USA, volume 3546 of *Lecture Notes in Computer Science*, pages 436–446. Springer, July 2005.
14. B.K. Widrow, I. Kollar, and M.C. Liu. Statistical theory of quantization. *IEEE Transactions on Instrumentation and Measurement*, 45(2):353–361, December 1996.

Appendix A

Computing $f_{D_{P_i}}(p)$, where $D_{P_i} = P_i^1 - P_i^2$

We compute $f_{D_{P_i}}$ when $f_U(u)$ is modeled as a normal distribution with mean $\mu = 0$ and variance $\sigma^2 = 1$. For simplicity, in the rest of this paragraph we omit the subscript i .

Firstly, we concentrate on estimating the probability distribution of the public sketch P_i , as a function of the quantization step. Widrow, *et al.* [14] show how the probability density of $f_P(p)$ can be constructed: the value of the sketch results from the quantization of x_i falling at just the right places within all the quantization boxes. Thus when Q is a scalar uniform quantizer with step size q and reconstruction points given by $Q(n \cdot q), \forall n \in \mathbb{Z}$, we can cut $f_{X_i}(x)$ into strips of length q , stacking the strips and then adding we arrive at:

$$f_P(p) = \begin{cases} \sum_n f_X(Q(nq) + p) & \text{if } |p| \leq \frac{q}{2} \\ 0, & \text{elsewhere.} \end{cases}$$

Figure 7 shows the probability distribution function of the sketch P_i for the variable X_i quantized using different quantization steps $q \in \{0.5, 1, 2, 3\}$. Widrow, *et al.* [14] observe that quantization is a deterministic process as long as $q \leq \sigma$, $f_P(p)$ is uniform.

Secondly we compute the probability distribution of variable D_P , which represents the difference of two random sketches $P^{(1)}$ and $P^{(2)}$. The joint probability is denoted with $f_{P^1 P^2}(p^1, p^2)$ and is taken over all the pairs of (p^1, p^2) where $p = p^1 - p^2$. We note that for P^1 and P^2 generated from a quantization step q we have $-\frac{q}{2} \leq p^1, p^2 \leq \frac{q}{2}$ and $-q \leq p \leq q$ where $p = p^1 - p^2$.




Fig. 7. Probability density function $f_{P_i}(p)$ for a random variable X modeled as a normal variable $N(\mu, \sigma)$ with $\mu = 0$ and $\sigma = 1$ and for different quantization steps $q \in \{0.5, 1, 2, 3\}$ and $-\frac{q}{2} \leq p \leq \frac{q}{2}$.

If we replace $p^1 = t$ and $p^2 = t - p$ Therefore:

$$f_P(p) = \int_{-q}^q f_{P^1 P^2}(t, t - p) dt$$

To compute the joint probability of P^1 and P^2 we observe that they are independent when the coin flip is $\mathbf{c} = 0$, (the sketches are not related) and they are completely determined (sampled from the same random variable) when the sketches are related, $\mathbf{c} = 1$. The joint probability of P^1 and P^2 is then:

$$\begin{aligned} f_{P^1 P^2}(t, t - p) &= f_{P^1 P^2}(t - p | t) f_{P^1}(t) \\ &= f_{P^1 P^2}(t - p | t, \mathbf{c}) f_{P^1}(t) f_{\mathbf{C}}(\mathbf{c}) \\ &= \frac{1}{2} f_{P^1 P^2}(t - p | t, \mathbf{c} = 0) f_{P^1}(t) + \frac{1}{2} f_{P^1 P^2}(t - p | t, \mathbf{c} = 1) f_{P^1}(t) \\ &= \frac{1}{2} f_{P^1}(t) (f_{P^1 P^2}(t - p | t, \mathbf{c} = 0) + f_{P^1 P^2}(t - p | t, \mathbf{c} = 1)) \\ &= \frac{1}{2} f_{P^1}(t) (f_{P^2}(t - p) + 1) \end{aligned}$$

Since variables P^1 and P^2 are identically distributed the joint probability can be written as:

$$f_{P^1 P^2}(t, t - p) = \frac{1}{2} f_{P^1}(t) (1 + f_{P^1}(t - p)).$$

Therefore, the probability density function of the difference of P^1 and P^1 is:

$$f_P(p) = \int_{-q}^q \frac{1}{2} f_{P^1}(t) (1 + f_{P^1}(t - p)) dt$$