
Information technology — Performance testing of biometric template protection schemes

*Technologies de l'information — Essais de performance des systèmes
de protection par modèle*

IECNORM.COM : Click to view the full PDF of ISO/IEC 30136:2018



IECNORM.COM : Click to view the full PDF of ISO/IEC 30136:2018



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2018

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Fax: +41 22 749 09 47
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Contents

Page

Foreword	iv
Introduction	v
1 Scope	1
2 Normative references	1
3 Terms and definitions	1
4 Abbreviated terms	3
5 Conformance	4
6 Methods for biometric template protection (informative)	4
6.1 General	4
6.2 Generalized architecture for biometric template protection system	5
6.3 Data separation	8
6.4 Examples of typical architectures in template protection systems	8
6.4.1 Biometric verification utilizing multiple databases	8
6.4.2 Two-factor biometric verification utilizing smart card	9
6.4.3 Two-factor biometric verification utilizing passwords	10
7 Overview of performance evaluation for biometric template protection schemes	10
7.1 Methods for attacking a biometric template protection system	10
7.2 Necessity of metrics beyond traditional recognition performance	10
7.3 Technology evaluation	11
7.4 Theoretical evaluation and empirical evaluation	11
7.5 Threat models	11
7.5.1 Naive model	12
7.5.2 Collision model	12
7.5.3 General models	12
8 Performance metrics for biometric template protection systems	13
8.1 General	13
8.2 Case of multiple biometric access control systems	13
8.3 Metrics for enrolment and verification performance	14
8.3.1 General	14
8.3.2 Accuracy degradation	14
8.3.3 Template diversity	15
8.3.4 Storage requirement per registered individual	16
8.4 Metrics for security and privacy protection performance	16
8.4.1 Irreversibility	16
8.4.2 Unlinkability	18
8.4.3 Successful Attack Rate (SAR) (optional)	19
Annex A (informative) Publication of algorithms or proofs used in performance evaluations	21
Bibliography	22

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see the following URL: www.iso.org/iso/foreword.html.

This document was prepared by ISO/IEC JTC 1, *Information technology*, SC 37, *Biometrics*.

Introduction

In conventional biometric access control systems, an adversary who compromises an enrolment database may gain access to the generative biometric data of the individuals enrolled therein. This is undesirable because, if the biometric system is vulnerable to presentation attacks or replay attacks, the adversary could impersonate an individual and gain access to the system after gaining access to the enrolment database. Furthermore, if the biometric enrolment databases contain unprotected templates and the same biometric modality is adopted in multiple applications, the adversary could link the accounts of the individual across those applications (cross-matching).

A biometric template stored in an enrolment database is a reference set of biometric features derived from the biological and behavioural characteristics of an individual. If the system implementation allows it, a biometric enrolment that is known to have been compromised may be revoked and renewed a limited number of times. However, the number of unique biometrics that can be extracted from an individual is limited and thus biometric enrolments cannot be revoked and then re-issued an unlimited number of times like new credit card numbers or passwords. The compromise of biometric enrolment records stored in an enrolment database is a serious issue. Therefore, methods and procedures to mitigate the risk of compromise are needed.

Secure biometric verification

The biometrics research community has invested significant effort in enabling biometric verification without directly needing to store an individual's biometric features in the clear at the access control device. This has led to the development of new methods referred to as "biometric template protection", "biometric information protection", or simply "secure biometrics". In this document, the term "biometric template protection" is used.

The rationale behind this strategy is that, instead of storing the biometric features directly, the access control system derives some data from the biometric features and stores this derived data on the device. During the biometric verification phase, the system receives a probe biometric sample from the individual seeking access. Then, the system combines the probe biometric sample and the derived data and generates a biometric verification decision. The main property of the derived data is that it reveals little or no information about the underlying biometric characteristic that was captured during the enrolment phase.

Thus, if the access control device is compromised by an adversary, only the derived data falls into the hands of the adversary, but this does not enable the adversary to recover the biometric characteristics of the individuals enrolled in the database. Clearly, this strategy protects the privacy of the individuals enrolled in the database.

Further, if an adversary attempts to gain access, i.e. to log in, to the system by providing a fake probe biometric sample, then in a well-designed secure biometric system, combining the fake probe biometric sample with the derived stored data results in biometric verification failure. Thus, this strategy protects the secrecy of the individuals enrolled in the database.

Rationale for new metrics

There are several ways in which biometric template protection can be realized. Some of these methods are described in ISO/IEC 24745:2011. Regardless of the method employed to construct the derived data, the following questions must be asked when evaluating a biometric template protection system:

- a) What is the probability that the system rejects genuine individuals and accepts imposters? This is a natural question to ask of *any* biometric verification system. The metrics, False Non-Match Rate (FNMR) and False Match Rate (FMR) measure this performance [ISO/IEC 19795-1] for the conventional biometric system in which enrolment biometric features are matched against probe biometric features. A biometric template protection system will also inherit these metrics, though the method of measuring them may vary depending upon the particular realization of the template protection algorithm.

- b) What is the probability that an adversary enhanced with some knowledge about the database of enrolled individuals can be successfully verified as one of those enrolled?
- c) How much information can an adversary obtain by compromising an access control device and stealing the derived (stored) enrolment information? In conventional biometric systems, the adversary may obtain significant information, in the form of the stored biometric template, or the stored feature vector. The goal of biometric template protection systems is to ensure that the stored derived data does not leak much information about the enrolled individuals.
- d) What is the probability that an adversary, having successfully compromised one or more access control devices and having stolen the data stored on them, uses the information gained to be successfully verified at an access control device?

These questions form the basis for evaluating the accuracy, secrecy, and privacy of a biometric template protection system, which introduces a new set of metrics not previously associated with evaluating traditional biometric systems.

Necessity for standardization

There are several architectures under the umbrella of biometric template protection, e.g., fuzzy vault-based systems, secure sketch-based systems, cancellable biometric systems, secure multiparty computation-based systems, etc. It is necessary to define key metrics that not only answer the questions posed above, but also apply to a wide variety of biometric template protection architectures, thereby providing a common basis for comparison of systems based on different architectures. The goal of this document is to specify new metrics for evaluating template protection-based biometric verification and identification systems. Theoretical and empirical definitions are provided for each metric in [Clause 8](#).

Information technology — Performance testing of biometric template protection schemes

1 Scope

This document supports evaluation of the accuracy, secrecy, and privacy of biometric template protection schemes. It establishes definitions, terminology, and metrics for stating the performance of such schemes. Particularly, this document establishes requirements for the measurement and reporting of:

- theoretical and empirical accuracy of biometric template protection schemes,
- theoretical and empirical probability of a successful attack on biometric template protection schemes (single or multiple), and
- the information leaked about the original biometric when one or more biometric template protection schemes are compromised.

This document also gives guidance on measuring and reporting diversity and unlinkability of templates.

This document does not:

- establish template protection schemes;
- address testing of traditional encryption schemes.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 19795-1, *Information technology — Biometric performance testing and reporting — Part 1: Principles and framework*

ISO/IEC 24745:2011, *Information technology — Security techniques — Biometric information protection*

ISO/IEC 2382-37, *Information technology — Vocabulary — Part 37: Biometrics*

3 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO/IEC 2382-37 and ISO/IEC 24745 and the following apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- IEC Electropedia: available at <http://www.electropedia.org/>
- ISO Online browsing platform: available at <http://www.iso.org/obp>

3.1

accuracy degradation

difference in FNMR/FMR (or FAR/FRR) for a biometric system tested both with and without template protection schemes

3.2

adversary

one who compromises an enrolment database and may gain access to the generative biometric data of the individuals enrolled therein

3.3

biometric template protection

protection of biometric references under various requirements for secrecy, irreversibility, and renewability during storage and transfer

3.4

generative biometric data

biometric data (sample(s) or features) used as primary input to the biometric template protection scheme

3.5

irreversibility

property of a transform that creates a biometric reference from generative biometric data such that knowledge of the transformed biometric reference cannot be used to determine any information about the generative biometric data

Note 1 to entry: Metrics introduced by this document aim to measure irreversibility by the degree of difficulty faced by an adversary in recreating an original unprotected version of the biometric data.

3.6

privacy compromise

event in which an adversary discovers part of the generative biometric data of an individual enrolled in the database of a biometric verification or identification system

Note 1 to entry: Discovery of part of the generative biometric data does not mean that successful biometric recognition is achieved, i.e. biometric system secrecy is not compromised by the discovery itself.

3.7

privacy leakage

<template protection scheme> amount of information about an individual's generative biometric data which an adversary can learn from the stored reference data

3.8

pseudonymous identifier comparator

system, process or algorithm that compares the pseudonymous identifier generated during enrolment by the pseudonymous identifier encoder and the pseudonymous identifier reconstructed during verification by the pseudonymous identifier recoder, and returns a similarity score representing the similarity between the two

3.9

pseudonymous identifier recoder

system, process or algorithm that reconstructs a pseudonymous identifier based on the provided auxiliary data and the extracted features

3.10

secrecy

degree of difficulty faced by an adversary in determining input data, from a protected biometric template, that achieves biometric recognition, when impersonating an individual enrolled in the biometric enrolment database of a template protection system

3.11**secrecy compromise**

event in which an adversary achieves biometric recognition when impersonating an individual enrolled in the biometric enrolment database of a template protection system

Note 1 to entry: Secrecy compromise includes the case in which the adversary gains unlawful access without necessarily discovering the generative biometric data of the individual being impersonated, i.e. the case in which the adversary remains unable to cause a privacy compromise.

3.12**successful attack rate**

probability that an informed adversary can obtain a false accept result in a biometric system

Note 1 to entry: An informed adversary is one that has compromised (gained access to) a subset of the biometric enrolment database and the secret parameters (if any) associated with one or more biometric recognition systems (potentially including the target system under consideration) in which common individuals are enrolled.

3.13**template diversity**

expected value of the number of independent protected templates that can be generated from a given generative biometric data by a biometric template protection scheme

3.14**template size**

size of stored reference data

3.15**unlinkability**

property of two or more biometric references that they cannot be linked to each other or to the subject(s) from which they were derived

4 Abbreviated terms

For the purposes of this document, the following abbreviated terms apply.

AD	Auxiliary Data
BTP	Biometric Template Protection
ECC	Error Correcting Code, or Error Correction Coding
FAR	False Acceptance Rate
FMR	False Match Rate
FNMR	False Non-Match Rate
FRR	False Reject Rate
$H(X)$	Information-theoretic Entropy of a random variable X
$H(X Y)$	Conditional entropy of random variable X given random variable Y
$I(X; Y)$	Mutual information between random variables X and Y
PI	Pseudonymous Identifier
PIC	Pseudonymous Identifier Comparator
PIE	Pseudonymous Identifier Encoder

PIR	Pseudonymous Identifier Recoder
RBR	Renewable Biometric Reference
SAR	Successful Attack Rate

5 Conformance

To conform to this document, evaluations of the accuracy, secrecy and privacy of biometric template protection schemes shall conform to [Clause 8](#).

6 Methods for biometric template protection (informative)

6.1 General

In this document, biometric template protection refers to the category of techniques that perform biometric verification or identification without storing the enrolment template, whether “in the clear” or encrypted via traditional means. Instead, the captured biometric sample is transformed in an irreversible fashion and the transformed result is stored in the database of enrolled individuals. If an adversary gains access to the database, only the transformed data is accessible, which has two beneficial properties:

- a) The stored data for an enrolled individual may reveal partial information about the features extracted from the individual’s biometric sample, but it does not contain enough to reconstruct the individual’s biometric characteristics.
- b) It is more difficult for the adversary to achieve biometric recognition when impersonating an individual enrolled in the enrolment database compared to systems that do not employ template protection schemes. Thus, biometric template protection provides improved secrecy for individuals enrolled in the system.

The extent of irreversibility and secrecy depends on a variety of factors, such as the type of biometric characteristic used, the feature extraction algorithm employed to extract a digital representation of the biometric signal, the mechanism used to provide template protection, and the use of optional secret keys as a second factor of secrecy. Providing increased irreversibility and secrecy may come at the cost of reduced biometric verification accuracy. In order to be able to study and analyze the various performance aspects of biometric template protection systems, it is necessary to precisely characterize the various dimensions of accuracy, secrecy and irreversibility. In this clause, a generalized architecture for biometric template protection systems is presented. Using this generalized architecture, metrics that quantify the accuracy, secrecy and irreversibility are defined in [Clause 8](#).

[Figure 1](#) illustrates the information flow within a general unprotected-template biometric system, from ISO/IEC/TR 24741:2018, Clause 6. Explanations about subsystems can be found in ISO/IEC/TR 24741:2018, Clause 6.




Figure 1 — Components of a conventional biometric system

6.2 Generalized architecture for biometric template protection system

Figure 2 illustrates the information flow within a general template protection biometric system, which can be regarded as an extended system of an unprotected-template biometric system described in Figure 1.




Figure 2 — Components of a general biometric template protection system

NOTE [Figure 2](#) represents components for typical verification/identification systems and does not cover a general identification scenario that could directly execute a 1 vs N comparison (not N times 1 vs 1). In a general 1 vs N identification system, the algorithm PIR can take multiple ADs as input and output multiple PI*s, and the algorithm PIC can take multiple PIs or PI*s as input.

The main difference between the conventional biometric system in [Figure 1](#) and the template protection system in [Figure 2](#) is that the Pseudonymous Identifier Encoder (PIE) generates Auxiliary Data (AD) and Pseudonymous Identifier (PI) from the extracted biometric features during enrolment. A pair of AD and PI is called a Renewable Biometric Reference (RBR). A Pseudonymous Identifier Recoder (PIR) generates a different PI (labelled PI*) from an enrolled AD and the extracted biometric features during verification or identification. A Pseudonymous identifier Comparator (PIC) compares PI and PI*. The definitions of AD, PI, RBR, PIE, PIR and PIC can be found in ISO/IEC 24745:2011, Clause 2 and 5.2.3.

In the following, the algorithms executed in the PIE, PIR and PIC modules are formally redefined in compliance with ISO/IEC 24745:2011[18]. Let U be a set consisting of all individuals' biometric characteristics. Assuming that biometric features are extracted from an acquired sample during enrolment, verification/identification can be represented as an element x of a space M and PI, AD and PI* are represented elements of spaces M_{PI} , M_{AD} and M_{PI}^* .

A tuple of the three algorithms, PIE, PIR and PIC, described below, is then called a biometric template protection (BTP) algorithm.

- A PIE that takes generative biometric data $x \in M$ and a randomness as input and returns a pair (α, π) of an AD $\alpha \in M_{AD}$ and a PI $\pi \in M_{PI}$. In [Figure 2](#), the PIE can be regarded as a creation module for the Renewable biometric references (AD, PI).
- A PIR that takes as input an auxiliary data α and generative biometric data $x' \in M$ and returns a PI $\pi' \in M_{PI}^*$ for verification or identification. In [Figure 2](#), the PIR can be regarded as a creation module for the authentication data (PI*) creation module.

- A PIC that takes as input two pseudonymous identifiers π and π' , returns either a match or non-match decision. In [Figure 2](#), the PIC can be regarded as a comparison module.

In the enrolment phase, a biometric characteristic u is presented to the system, generative biometric data $x \in M$ are extracted from the characteristic u . The PIE takes as an input x and outputs a pair $(\alpha, \pi) \in M_{AD} \times M_{PI}$ of an AD and a PI and α and π are stored in storage modules labeled as the AD database and the PI database respectively. Note that α and π may be stored in separate storage modules (see [6.3](#)).

In the verification/identification phase, a biometric characteristic u' is freshly presented to the system, generative biometric data $x' \in M$ are extracted from the characteristic u' . The PIR takes as an input x' and outputs π' , the PIC compares π and π' , and outputs a comparison result that is sent to the decision subsystem.

The above concepts can be alternatively represented in [Figure 3](#) below, which is reproduced from ISO/IEC 24745:2011, Figure 5.




Figure 3 — Architecture for renewable biometric references (ISO/IEC 24745:2011)

[Figure 2](#) shows a generalized architecture of a biometric template protection system. This architecture is called “generalized” because it captures various ways in which biometric recognition schemes can be realized, including, but not necessarily limited to:

- Biometric fuzzy vault**, a template protection scheme in which the individual data stored in the biometric enrolment database is a function of a randomly generated quantity and the biometric sample [\[6\]](#)[\[7\]](#).
- Biometric secure sketch**, a template protection scheme in which the individual data stored in the biometric enrolment database is a “sketch” that, by itself, provides – at most – partial information about the underlying biometric sample [\[8\]](#)[\[9\]](#).
- Cancelable biometrics**, a template protection scheme in which the individual data stored in the biometric enrolment database is one of several non-invertible transformation or warping of the biometric sample [\[10\]](#).
- Biohashing**, a biometric feature vector is projected onto a random vector space using a secret random projection matrix, and biometric recognition is done in the space of the random projections [\[11\]](#).

The examples of the AD and PI for each explicit protection scheme are given in ISO/IEC 24745:2011, Table D.1.

In conventional biometric recognition without template protection, the AD or PI is the biometric sample itself, or an unprotected template derived from the biometric sample. In biometric template protection systems, the generative biometric data is only an intermediate step towards biometric recognition. A second signal is derived from the generative biometric data, which is stored in the biometric enrolment database.

Furthermore, biometric template protection systems can be implemented by cascading conventional biometric recognition systems with a template protection primitive such as a public-key cryptosystem^{[19][21]}, a secret transformation^{[10][11]}, or an ECC^[20]. Indeed, many implementations of the biometric fuzzy vault, biometric secure sketch, cancellable biometrics, and bihashing are realized in this way.

6.3 Data separation

In the real world, there are many large public databases of biometric samples for some modalities. When an individual's enrolled AD and PI are simultaneously leaked to the adversary, they can find a sample which matches the leaked AD and PI with complexity of approximately $1/\text{FAR}$ trials by testing each sample in such a public database. For almost every existing modality, since the biometric sample entropy $\log(1/\text{FAR})$ is not quite large enough, BTP algorithms, with openly accessible parameters, cannot achieve sufficiently strong secrecy when both AD and PI are simultaneously compromised. Therefore, in order to prevent simultaneous leakage of both AD and PI, each individual's AD and PI are recommended to be stored in separate storage modules. In 6.4, some examples of typical applications which store AD and PI separately are described. The reference document^[15] discusses the secrecy for some data-separation models.

When individuals' AD and PI are separately stored in different storage modules, the following cases of data leaks can be considered:

- a) either AD or PI is leaked to the adversary;
- b) both AD and PI are simultaneously leaked to the adversary.

Although case b) may occur with a low probability, this document specifies performance metrics which covers all of the above cases in [Clause 8](#).

6.4 Examples of typical architectures in template protection systems

6.4.1 Biometric verification utilizing multiple databases

[Figure 4](#) illustrates the case where each individual's enrolled AD and PI are separately stored along with the individual's ID or a common linking data in different databases. During verification, the algorithm generates a new PI (described as PI^*) from a submitted biometric sample and an AD (linked to the claimed ID), compares the PI^* with the enrolled PI corresponding to the AD, returns a comparison result, and outputs a verification result.




Figure 4 — Application model of biometric template protection system with separate AD and PI databases

6.4.2 Two-factor biometric verification utilizing smart card

Figure 5 illustrates the case where each individual's enrolled AD is stored in the individual's smart card, while the PI is stored in a separate PI database situated, either at the template protection system, or in distributed storage. During verification, the algorithm generates a new PI from submitted biometric sample and AD, compares the PI* with the enrolled PI (corresponding to the claimed ID), returns a comparison result, and outputs a verification result.




Figure 5 — Application model of two-factor biometric template protection system based on smart card

6.4.3 Two-factor biometric verification utilizing passwords

Figure 6 illustrates the case where each individual's enrolled PI is stored along with the individual's ID in the database and the individual memorises his/her own AD as a password. During verification, the algorithm generates a new PI from submitted biometric sample and AD (password), compares the PI* with the enrolled PI (corresponding to the claimed ID), returns a comparison result, and outputs a verification result.




Figure 6 — Application model of password two-factor biometric template protection system

7 Overview of performance evaluation for biometric template protection schemes

7.1 Methods for attacking a biometric template protection system

Clauses 7.1 to 7.5 describe important information intended to assist in the understanding and/or use of metrics and requirements provided in Clause 8.

Based on Figures 2 and 3, there are several ways in which an adversary can compromise a biometric template protection system. The following is a non-exhaustive list of attacks that can be carried out individually or in combination, at one or more biometric system in which an individual is enrolled:

- 1) **Attack on the biometric database:** In this attack, an adversary gains access to the stored AD and/or PI corresponding to a subset of the enrolled identities.
- 2) **Attack on the generative biometric data:** In this attack, an adversary gains access to the generative biometric data used to register one or more individuals in the database of biometric identities.
- 3) **Attack on the optional key:** This attack applies only to a two-factor biometric verification or identification system. In this attack, an adversary gains access to the auxiliary data, namely the optional key returned to one or more enrolled individuals during enrolment.

7.2 Necessity of metrics beyond traditional recognition performance

Clause 8 includes metrics to measure the recognition performance of a BTP algorithm. These include the traditional metrics, such as False Match Rate (FMR) and False Non-Match Rate (FNMR), in conformity

with ISO/IEC 19795-1 and ISO/IEC 19795-2:2007. However, the existing International Standards do not cover all the metrics required to evaluate BTP algorithms. The goal of this document is to establish these new performance metrics and specify methods for evaluating them.

The purpose of a BTP algorithm is to provide countermeasures to eliminate privacy and security risks in biometric systems. The potential risks and the properties of privacy and security are elaborated in ISO/IEC 24745:2011. Depending on these properties, biometric information protection should achieve the following important targets: template diversity, irreversibility and unlinkability.

7.3 Technology evaluation

This document addresses technology evaluation of BTP algorithms. Technology evaluation is the offline evaluation of BTP algorithms (not including the sample acquisition algorithm and the feature extraction algorithm) for a given biometric modality using a pre-existing or specially collected corpus of samples.

7.4 Theoretical evaluation and empirical evaluation

In the theoretical evaluation of recognition performance, the experimenter formally demonstrates the recognition error rate. In the theoretical evaluation of the security and privacy protection provided by a BTP algorithm, the experimenter formally demonstrates the attack success probability or the advantage of a specific adversary.

In the empirical evaluation of recognition performance, the experimenter shall estimate recognition error rates, in conformity with ISO/IEC 19795-1 and ISO/IEC 19795-2:2007. However, ISO/IEC 19795-1:2006 and ISO/IEC 19795-2:2007 do not cover the empirical evaluation of the security and privacy protection provided by a BTP algorithm.

In general, theoretical evaluation gives a higher assurance of the target performance than the corresponding empirical evaluation. However, if there exists evidence that the attack assumed in the empirical evaluation is the strongest possible attack for an adversary with specified capabilities, the empirical evaluation gives assurance equivalent to the theoretical evaluation.

7.5 Threat models

Evaluation of the security and privacy protection performance of a biometric template protection system can be based on measurements using, for example, information-theoretic properties such as conditional entropy, conditional min-entropy, or mutual information. The evaluation results represent the level of security and privacy protection performance independent of the information and resources available for an adversary.

EXAMPLE If the conditional entropy of the biometric data given the protected template is 20 bits, then the uncertainty about the biometric data is at least 20 bits, even when the attacker with access to the protected template has knowledge of the biometric template protection system.

Thus, information-theoretic quantities specify the “unconditional” security and privacy protection performance achieved by a biometric template protection algorithm. However, it is not always possible to measure unconditional performance. For example, measurements may need comprehensive estimation of joint probability distributions of biometric data and protected templates, which may be infeasible in practice. Alternatively, the security of an algorithm relies on the hardness of reversing encryption systems for which unconditional guarantees are not available. Additionally, an unconditional performance metric does not represent the effort that an adversary expends to achieve their goal. Therefore, it is important to measure “conditional performance” in the terms of computational complexity. Here, information and resources available to an adversary are important pre-conditions for evaluation.

In [7.5.1](#) to [7.5.3](#), three main threat models are given, which shall be used to classify the adversaries. [Subclauses 7.5.1](#) and [7.5.3](#) respectively described the naive and general models, which are comparable with the models in traditional cryptanalysis. [Subclause 7.5.2](#) describes the collision model, which is derived based on inherent properties of biometric systems. Specifying a threat model is a prerequisite

for quantifying security and privacy. In practice, an evaluator can derive their own refined threat model from these threat models based on the security and privacy requirements of the target applications.

7.5.1 Naive model

In this model, an adversary has neither information of the underlying algorithm in a biometric template protection system, nor owns a large biometric database. They only have access to the RBRs. The protected system is considered as a black box. Attacks that can be performed on biometric information that can be obtained in such a scenario are very restricted.

7.5.2 Collision model

In this model, an adversary possesses a large amount of biometric data. They can use this information to exploit inaccuracies in biometric systems, make an exhaustive search in their own database and find biometric data, which generates a PI that has sufficient similarity to the data of an individual in the biometric enrolment database.

7.5.3 General models

In this model, per Kerckhoffs' principle^[12], an adversary is assumed to possess full knowledge of the underlying template protection scheme algorithms, PIE, PIR and PIC described in [Figure 2](#). In addition, the adversary may have access to one or more protected templates from one or more databases. The adversary may also possess knowledge of the statistical properties of biometric features. However, the adversary is not allowed to alter sample acquisition algorithms or feature generation algorithms such as segmentation, feature extraction, or quality control as illustrated in [Figure 2](#).

In this case, the security and privacy protection may depend only on the presence of a secret parameter which is known to a legitimate subject of the system and a system administrator, but not to the adversary. Examples of such secret parameters include transformation parameters in cancellable biometrics, projection matrices in biohashing, and decryption keys in biometric cryptosystems based on homomorphic encryption.

It is possible to consider an adversarial model that is even stronger than the general model described above. In this stronger variant, it may be assumed that the adversary has gained access to a subset of the secret parameters. For example, in two-factor fuzzy commitment schemes, the protected template of an individual may be compromised, while the secret keys assigned to legitimate individuals are not compromised. As another example, the index of the secret transformation used for warping face features in cancellable biometrics may be compromised by an adversary. As expected, these are very powerful attacks, which can result in a dramatic reduction in the secrecy offered by a biometric template protection system. In some cases, (e.g. for secure sketch and fuzzy commitment systems based on error correcting codes), it is possible to quantify the loss of secrecy when such powerful attacks occur. In other cases, (e.g. for cancellable biometrics), quantification of the loss of secrecy may not be possible. These considerations lead to three important subcases of the general model as described in [7.5.3.1](#) to [7.5.3.3](#).

7.5.3.1 Standard model

In this model, the adversary has the full knowledge of the algorithms used for template extraction, template protection and comparison, following Kerckhoffs' principles, but cannot execute the submodules of the system that make use of the secrets (if any). In particular, this implies that the adversary knows none of the secrets.

NOTE This is related to known-ciphertext attack concept in cryptography.

7.5.3.2 Advanced model

In this model, the standard model is augmented with the capability of the adversary to execute part of or all submodules that make use of the secrets (if any).

NOTE This is related to the concept of chosen-plaintext attack and chosen-ciphertext attack in cryptography.

7.5.3.3 Full disclosure model

In this model, the standard model (or the advanced model) is augmented by disclosing the secrets (if any) to the adversary.

NOTE In the two latter models, it is possible to consider variants where only part of the secrets are known or executables.

8 Performance metrics for biometric template protection systems

8.1 General

This clause provides guidelines about metrics that should be used to evaluate various aspects of a biometric template protection system. The immediate goal of this document is to provide metrics to evaluate a certain class of systems, namely biometric template protection systems described in ISO/IEC 24745. At the same time, it is desirable to specify metrics that are as general as possible, and that apply to a generalized biometric verification architecture described in [Clause 6](#). Throughout this document, the usual distinction is made between biometric verification, which involves a one-to-one comparison operation, and biometric identification which involves a one-to-many comparison operation.

All metrics shall be estimated with respect to each applicable threat model as discussed in [7.5](#). This is needed to clearly state which threat models are considered during an evaluation. Indeed the adversary or the experimenter will have a different knowledge (on the underlying data, algorithms and secrets) that depends on the exact threat model. To be meaningful, the threat model or the threat models (from [7.5](#)) used by the evaluator shall be explicitly reported with the evaluation results. In particular, the metrics for protection performance described in [8.4](#) shall be evaluated based on certain threat models including the general models defined in [7.5.3](#).

Test design, assembling an appropriate test corpus, and reporting shall be compliant with ISO/IEC 19795-2:2007, 6.1, 6.2, and 6.4, respectively, in the evaluation of the metrics, False Match Rate, False Non-Match Rate, Failure-To-Enrol (FTE) Rate, Failure-To-Acquire (FTA) Rate, and in the empirical evaluation of Successful Attack Rate, Template Diversity, Irreversibility, and Unlinkability.

ISO/IEC 19795-2:2007, 6.1, 6.2, and 6.4 refer to technology tests. Often, template protection systems are implemented such that the biometric sample processed by the template protection system needs to come from a sensor (as opposed to a previously-collected image stored on a file system). In such cases, technology tests – which rely on a previously collected corpus – cannot be executed. The experimenter may need to obtain a version of the template protection system that decouples the sensor and algorithm.

8.2 Case of multiple biometric access control systems

This clause discusses the impact on metrics for evaluation in the scenario where a given individual has enrolled in multiple biometric template protection systems for purposes of biometric verification and/or identification. For example, an individual may enroll their right index finger in biometric template protection systems at their gym, bank, workplace, and mobile device.

Several metrics, such as False Non-Match Rate, False Match Rate, or Storage Requirements, are straightforward to generalize to the multi-system scenario and are defined in exactly the same way as for each individual system. However, some metrics, such as Irreversibility, must explicitly handle this multi-system scenario, since such metrics are defined in terms of compromised biometric information. An adversary who attacks multiple biometric template protection systems can conceivably obtain a

greater amount of verifying or identifying information about an individual than a simple adversary who attacks only a single template protection system. Thus, it is expected that the successful attack rate and the privacy leakage in the context of multiple compromises are at least as large as – and possibly larger than – the respective quantities in the context of a single system compromise.

A vital point to consider in designing and deploying biometric access control systems for individuals who have enrolled in multiple such systems is as follows: even if an adversary has not compromised a particular biometric access control device, their ability to attack that device and to compromise the irreversibility of one of its enrolled individuals, can be enhanced by attacking other devices in which the individual is enrolled.

Further, analyzing the multi-system case illuminates a subtle trade-off between the multi-system Successful Attack Rate (SAR) and the multi-system privacy leakage^[3]. Specifically, for many constructions of biometric template protection systems, such as fuzzy vaults or secure sketches, it is not possible to simultaneously reduce the multi-system SAR and the multi-system privacy leakage.

8.3 Metrics for enrolment and verification performance

8.3.1 General

In addition to the metrics defined here, the metrics FMR, FNMR, FTE Rate, and FTA Rate, which are defined in ISO/IEC 19795-1, are also applicable to biometric template protection systems.

8.3.2 Accuracy degradation

8.3.2.1 General

As noted earlier, most biometric template protection algorithms are implemented by extracting suitable biometric features and processing the features using a template protection primitive, such as a public-key cryptosystem, an error correcting code, or a secret transformation. Using the information processing inequality [CT06], it is possible to show that such processing of feature vectors cannot increase the information content in them. A practical consequence of this property is that the incorporation of a template protection primitive cannot improve the recognition accuracy. A well-designed biometric template protection system attempts to minimize the accuracy degradation.

8.3.2.2 Theoretical definition

Degradation in accuracy is defined as the difference between a metric used to measure the recognition accuracy (such as a FNMR at a fixed FMR) when template protection is not applied and the same metric when template protection is applied^[13]. This degradation is a characteristic of the template protection scheme.

8.3.2.3 Operational definition

For a given biometric dataset, biometric feature extraction scheme, and template protection scheme, the degradation in accuracy is defined as the difference between a metric used to measure the recognition accuracy (such as the operational FNMR at a fixed operational FMR) when template protection is not applied and the same metric when template protection is applied. This reduction may be expressed as a percentage. This degradation depends not only on the template protection scheme, but also on the way in which the biometric feature extraction algorithm is implemented.

8.3.2.4 Implications

Because the enrolment and comparison algorithms used in biometric template protection systems typically differ from those used in conventional biometric systems, it may be difficult to measure the impact of template protection on performance within a single system. In other words, one typically cannot simply run two tests on a given system, one with biometric template protection functionality enabled and one with biometric template protection disabled, and then compare the performance

results. Instead, an experimenter may need to conduct a technology test to generate results for multiple conventional biometric algorithms, thereby obtaining a baseline for FMR/FNMR/FTE/FTA against which biometric template protection performance can be compared.

8.3.3 Template diversity

8.3.3.1 General

Template diversity is the expected value of the number of independent protected templates that can be extracted from a generative biometric data by a biometric template protection algorithm. More precisely, template diversity is calculated as a reciprocal of the cross match rate (CMR), namely as $1/\text{CMR}$, where CMR is computed by using the following formula:

$$\text{CMR} = \frac{1}{n} \sum_{i=1}^n a_i$$

where

n number of test subjects;

a_i for the i th subject, the rate of tuples $(x, (\alpha, \pi), (\alpha', \pi'))$ of an i th subject's generative biometric data, x , and two different RBRs, (α, π) and (α', π') , generated from the i th subject's biometric characteristic to satisfy $\text{PIC}(\pi, \text{PIR}(\alpha', x)) = \text{match}$.

NOTE In two-factor template protection systems as described in 6.4.2 or 6.4.3, CMR can be also regarded as recognition accuracy metric after re-enrolling new AD and PI and reissuing a new smart card or password. More precisely, in a two-factor template protection system based on smart cards, CMR can be regarded as the probability that a genuine individual who presents their own biometric characteristic and old card is falsely accepted. In a two-factor template protection system based on password, CMR can be regarded as the probability that a genuine individual who presents their own biometric characteristic, ID and old (wrong) password is falsely accepted.

Template diversity directly ensures renewability and revocability. In particular, when a template is known to have been compromised, it can be revoked, i.e. cancelled by the system administrator, and a new template can be assigned, if available. A large template diversity implies that the stored template can be revoked and renewed a larger number of times. While this quality is, in itself, desirable, it may come at the cost of increased SAR or increased reversibility.

For each biometric template protection algorithm tested, the experimenter shall demonstrate the template diversity via the theoretical or empirical evaluation methods described in 8.3.3.2 or 8.3.3.3, respectively.

8.3.3.2 Theoretical evaluation

The experimenter shall specify a theoretical method of demonstrating template diversity. If (the lower bound of) template diversity has been theoretically proved, evidence of the reliability of the theoretical evaluation shall be reported. For example, the experimenter may report as evidence a description of the proof of diversity published in conferences or journals listed in Annex A.

8.3.3.3 Empirical evaluation

The experimenter shall specify an empirical method of demonstrating template diversity. If the probability is empirically estimated, evidence of the reliability of the empirical evaluation by the employed experiments shall be reported. Example of such evidence is publication of the description of the employed empirical estimation in conferences or journals listed in Annex A.

8.3.4 Storage requirement per registered individual

The storage requirement is defined as the number of bits required per enrolled individual in the biometric enrolment database. For a template protection system, the biometric enrolment database consists of the AD and PI databases, as depicted in [Figures 2 and 3](#). Thus, an equivalent definition of the storage requirement is the sum of the bits required per enrolled individual in the AD and PI databases of a template protection system.

The AD and PI databases may not both reside at the same place, as indicated by the examples in [6.3](#). Therefore, when the storage requirement is reported, the number of bits needed for the AD database and the PI database shall be reported separately. Some examples are provided below, to indicate how the storage requirement is calculated.

In biometric template protection systems based on the fuzzy vault, e.g. Reference [7], the storage requirement per individual for the PI database is the number of bits required to store a cryptographic hash of an individual-specific secret string. The storage requirement per individual for the AD database is the number of bits required to store the point set, namely the set of genuine feature points and chaff points.

In biometric template protection systems based on the secure sketch implemented using ECCs, e.g. Reference [3], the storage requirement per individual for the PI database is the number of bits required to store a cryptographic hash of an individual's biometric feature vector or a perturbation of the biometric feature vector. The storage requirement per individual for the AD database is the number of bits required to store the syndrome of an ECC.

In biometric template protection systems based on cancellable biometrics, e.g. Reference [10], the storage requirement per individual for PI database is the number of bits required to store a transformed template. The storage requirement per individual for the AD database is the number of bits required to store the transformation parameters.

More examples of the PI and AD databases for various implementations of biometric template protection systems are given in ISO/IEC 24745:2011, Table D.1.

8.4 Metrics for security and privacy protection performance

8.4.1 Irreversibility

8.4.1.1 General

Irreversibility is the property of a template protection algorithm that creates a biometric reference from the generative biometric data such that knowledge of the transformed biometric reference cannot be used to determine any information about the generative biometric data. More precisely, irreversibility is the difficulty of determining, from an AD and/or a PI generated from the extracted biometric features, biometric features which are closer to the generative biometric data than to randomly drawn biometric features.

In conventional biometric systems, partial or total compromise of the biometric enrolment database necessarily results in partial or total compromise of the irreversibility of generative biometric data. In a template protection system, this may or may not be the case. For example, using a two-factor template protection scheme, it is possible to ensure that, even if the adversary compromises the information stored in the AD database, the amount of information revealed about the generative biometric data of an individual remains zero[3].

Full reversibility corresponds to the case where the adversary retrieves exactly the generative biometric data. In some situations, the difficulty to retrieve an approximation of the generative biometric data could be easier than being able to achieve full reversibility; this case corresponds to the notion of partial reversibility. Consequently, the evaluator may distinguish both cases if the target is only resistant to full reversibility. If not made explicit, resistance to partial reversibility shall be considered as the default property to achieve, as it is stronger than resistance to full reversibility.

Strength of irreversibility may be classified depending on how much the adversary is empowered, in particular depending on how much they know about the template protection algorithm. To this aim, the strength of irreversibility shall be evaluated following certain threat models including the general models defined in 7.5.3.

For each biometric template protection algorithm tested, the experimenter shall demonstrate irreversibility by the methods described in 8.4.1.2 or 8.4.1.3, namely by theoretically or empirically estimating the adversary's success probability or the adversary's advantage over a random guess when the adversary attempts to retrieve, from a given AD and/or a PI, biometric features sufficiently similar to the original biometric features.

8.4.1.2 Theoretical evaluation

The experimenter shall specify a theoretical method of demonstrating irreversibility. For example, irreversibility may be demonstrated by the following methods:

- a) proving a computational property limiting the success probability or the advantage of any computationally-bounded adversary;
- b) proving an information-theoretic property limiting the success probability or the advantage of any adversary having unbounded computational power.

NOTE Some literature^{[15][16][3]} proposes an information-theoretic security notion evaluated by using the mutual information $I(X;Y)$ between the distributions X and Y of biometric features and ADs, respectively, or the mutual information $I(X;Z)$ between the distributions X and Z of biometric features and PIs, respectively. Mutual information between two sets of quantities is always non-negative and is equal to zero if and only if the two sets are independent [CT06]. Furthermore, we can always write the mutual information between two random quantities X and Y as $I(X;Y) = H(X) - H(X|Y)$ where $H(\bullet)$ and $H(\bullet|\bullet)$ are, respectively, the entropy and conditional entropy. Thus, mutual information characterizes the reduction in uncertainty about one random quantity, X , when given knowledge of another, Y . Quantified in this fashion, irreversibility is measured as the number of bits of information about the generative biometric data revealed to the adversary.

If (the upper bound of) the probability has been theoretically proven, evidence of the reliability of the theoretical evaluation shall be reported. For example, the experimenter may report as evidence a description of the proof of irreversibility published in conferences or journals listed in Annex A.

8.4.1.3 Empirical evaluation

The experimenter shall specify an empirical method of demonstrating irreversibility, for example, empirically estimating the success probability or the advantage over a random guess for some specific adversary in an attack scenario.

If the probability is empirically estimated, evidence of the reliability of the evaluation by the employed experiments shall be reported. Example of such evidence is publication of the description of the employed empirical estimation in conferences or journals listed in Annex A.

8.4.1.4 Implications in the multi-system case

Depending on the template protection strategy adopted at each system, the following example scenarios are possible^[3] when an adversary compromises multiple biometric access control systems in which the individual is enrolled:

- a) There is perfect irreversibility, i.e. no information is leaked about the individual's biometric sample. This is possible, for example, in a two-factor template protection system in which the compromised data for any device consists of either the stored template or the secret key used to obfuscate the biometric sample, but not both.
- b) The total degree of irreversibility is no worse than the irreversibility for any one compromised system. This is possible, for example, when each of the biometric access control or identification systems has a secure sketch architecture with the same ECC.

- c) The irreversibility worsens until, eventually, when a large enough number of access control systems has been compromised, the individual's biometric sample is completely revealed to the adversary. This is possible, for example, when each of the biometric access control or identification systems uses an ECC that is partially or fully independent of the ECCs of all other devices. Because of this final possibility, it is desirable to also report the worst case irreversibility, i.e. the irreversibility for the scenario in which biometric data and secret keys (if any) stored at all biometric systems under consideration, have been compromised by the adversary. This consideration is important because, all other performance measures being equal, it may suggest a preferred implementation of biometric template protection for the multiple biometric systems under consideration.

8.4.2 Unlinkability

8.4.2.1 General

Linkage attacks can occur in situations where the same biometric characteristic is used to enroll in multiple biometric systems, e.g. on several access control devices. If an adversary compromises a subset of the devices, the compromised data can be used to attack the remaining devices. The compromised data can both leak information about the underlying biometric characteristic and can be exploited to mount a successful attack on, i.e. gain unauthorized access to, one of the remaining devices.

Unlinkability is the property that two or more RBRs cannot be linked to each other or to the subject(s) from which they were derived. More precisely, unlinkability is the difficulty of distinguishing between ADs and/or PIs of two RBRs generated from the same subject's characteristic and ADs and/or PIs of two RBRs generated from different subjects' characteristics. Unlinkability can be regarded as the difficulty of classifying RBRs over time and across systems.

Unlinkability shall be evaluated with respect to certain threat models including the general models defined in 7.5.3. The strength of unlinkability will therefore directly depend on the level of the threats that is considered.

For each biometric template protection algorithm tested, the experimenter shall demonstrate unlinkability by the methods described in 8.4.2.2 or 8.4.2.3, namely by theoretically or empirically estimating the adversary's advantage over a random guess when the adversary attempts to determine whether two given ADs and/or PIs are generated from the same subject's characteristic or from different subjects' characteristics.

8.4.2.2 Theoretical evaluation

The experimenter shall specify a theoretical method of demonstrating unlinkability. For example, unlinkability may be demonstrated by the following methods:

- a) proving a computational property limiting the advantage over a random guess of any computationally-bounded adversary;
- b) proving an information-theoretic property limiting the advantage over a random guess of any adversary having unbounded computational power.

If (the upper bound of) the probability has been theoretically proven, evidences to show the reliability of the theoretical evaluation shall be reported. For example, the experimenter may report as evidence a description of the proof of unlinkability published in conferences or journals listed in Annex A.

8.4.2.3 Empirical evaluation

The experimenter shall specify an empirical method of demonstrating unlinkability, for example, empirically estimating the adversary's advantage over a random guess for some specific adversary in an attack scenario.