

Defining Security Limits in Biometrics

Axel DURBET
axel.durbet@uca.fr

UCA, LIMOS (UMR 6158 CNRS), Clermont-Ferrand,
France

Pascal LAFOURCADE
pascal.lafourcade@uca.fr

UCA, LIMOS (UMR 6158 CNRS), Clermont-Ferrand,
France

Paul-Marie GROLLEMUND
paul_marie.grollemund@uca.fr

UCA, LMBP (UMR 6620 CNRS), Clermont-Ferrand,
France

Kevin THIRY-ATIGHEHCHI
kevin.atighehchi@uca.fr

UCA, LIMOS (UMR 6158 CNRS), Clermont-Ferrand,
France

Abstract

Biometric systems are widely used for *authentication* (i.e., one-to-one matching to verify a claimed identity) and *identification* (i.e., one-to-many matching to find an identity in a database). The *False Match Rate* (FMR) quantifies the probability of matching a biometric template to a non-corresponding template and serves as an indicator of the system robustness against security threats. We analyze biometric systems through two main contributions. First, we study *untargeted attacks*, where an adversary aims to impersonate any user in the database. We compute the number of trials needed for a successful impersonation and derive the *critical population size* (i.e., the maximum database size) and *critical FMR* required to maintain security against *untargeted attacks* as the database grows. Second, we address the *biometric birthday problem*, which quantifies the probability that there exists two distinct users that collide (i.e., can impersonate each other). We compute approximate and exact probabilities of collision and derive the associated *critical population size* and *critical FMR* to bound the risk of biometric collisions, particularly in large-scale databases. These thresholds provide actionable insights for designing biometric systems that mitigate the risks of impersonation and biometric collisions, particularly in large-scale databases. Nevertheless, our findings show that current systems fail to meet the required security level against *untargeted attacks*, even in small databases, and face significant challenges with the *biometric birthday problem* as databases grow.

Keywords

Biometric Security; Minimal Leakage Model; Birthday Problem ; False Match Rate ; Untargeted Attacks ; Collisions

ACM Reference Format:

Axel DURBET, Paul-Marie GROLLEMUND, Pascal LAFOURCADE, and Kevin THIRY-ATIGHEHCHI. 2025. Defining Security Limits in Biometrics. In . ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Conference'17, July 2017, Washington, DC, USA

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 Introduction

Biometric systems, operating in either authentication or identification modes, leverage unique physical or behavioral characteristics to confirm or establish an individual's identity within a database. Authentication compares a fresh biometric data with an enrolled biometric template, such as when unlocking a smartphone or logging into an account. Identification, on the other hand, entails searching for a match between an individual's biometric data and a database of enrolled records. This process is commonly employed in forensic investigations or large-scale governmental programs, such as national identification systems and border control, where rapid and accurate identification is essential. Biometric technologies are now widely adopted in everyday devices and critical applications, such as online financial platforms, passport verification, and national identification systems (e.g., India's Aadhaar [42], Estonia's e-Residency [51], or Malaysia's MyKad [32]).

Biometric templates are structured representations of raw biometric data (*feature vectors*), designed to enable efficient matching and to facilitate integration with security technologies, including Biometric Template Protection (BTP) schemes [38, 47]. Despite their convenience, biometric data present unique security challenges, as they are inherently not revocable once compromised, unlike passwords. This lack of revocability makes them highly sensitive, and their sensitivity is further amplified by vulnerabilities inherent to biometric systems. For instance, *untargeted attacks* exploit the probabilistic nature of biometric matching to compromise any template in the database without targeting a specific user, as opposed to *targeted attacks*, where the adversary's goal is to impersonate a particular individual by matching their specific template. In this work, we adopt the *minimal information leakage* model, which ensures that the comparison process reveals only the binary outcome (i.e., acceptance or rejection), completely preventing adversaries from exploiting intermediate computations or states. Such attacks achieve two objectives: (i) impersonating any user in the database, effectively granting unauthorized access, and (ii) compromising the template of any legitimate user by generating an attacker's input that closely collides with the stored template. These collision inputs expose template details, which could be sensitive if the template is not revocable, thereby violating data protection regulations like the General Data Protection Regulation (GDPR). In the identification mode, *untargeted attacks* exploit the absence of limits on the number of attempts, unlike the authentication mode, which often enforces restrictions similar to password-based mechanisms. Analyzing these attacks under worst-case assumptions—such as

unlimited attempts—helps uncover systemic vulnerabilities and informs security improvements. This approach provides valuable insights into the weaknesses of biometric systems, even when practical constraints, such as enforced attempt limits, are in place.

A key metric for assessing the security of biometric systems is the False Match Rate (FMR), which quantifies the likelihood of an incorrect acceptance or identification. We demonstrate how the security level depends on parameters such as the FMR, which directly influences the probability of successful attacks, particularly as the density of enrolled templates in the database increases. This parameter is especially useful for determining the system's vulnerability to *untargeted attacks* that leverage large database sizes to increase the likelihood of success. Moreover, the FMR is also of interest for understanding the *biometric birthday problem*, which highlights the heightened risk of template collisions in large databases, analogous to the classic *birthday problem*, where the likelihood of shared birthdays increases rapidly with group size. Our analysis of these issues and their implications on attack complexity provides useful findings for mitigating vulnerabilities. For instance, by quantifying the impact of database size and FMR, we align our recommendations with principles outlined in international standards.

We draw an analogy with cryptographic hash functions to support this constructive approach. Cryptographic hash functions are evaluated based on parameters such as the output length, which directly influence resistance to preimage and collision attacks [16, 45]. Similarly, the FMR serves as a key parameter for quantifying attack complexities and deriving system-level recommendations. Just as collision resistance in hash functions informs design guidelines, the relationship between FMR, database size, and security level can guide the construction of robust biometric systems. However, as with cryptographic hash functions, correctly dimensioned parameters alone do not guarantee security. Design flaws or insufficient entropy in biometric templates can still compromise system robustness. The guidelines proposed here are therefore intended to complement broader efforts in building secure and reliable systems.

Although biometric templates are stored in a form that preserves their secrecy and data comparisons are performed securely within biometric systems, their compromise remains possible under the *minimal leakage* model, in both *online* and *offline attack* scenarios, which are relevant to the attacks studied in this paper. In the *online attack scenario*, adversaries interact with the remote biometric system to test inputs of their choice and identify matches. These attacks rely on the ability to make repeated queries and, when constrained in number, may resemble *password spray attacks* [56], where an attacker minimizes the number of guesses per user while maximizing the number of users targeted.

In the *offline attack scenario*, adversaries gain access to encrypted or transformed templates after a biometric database breach. Such attacks are particularly effective when the transformation relies solely on biometric data as the secret, which is the case for cryptographic obfuscations of evasive functions [18, 19] and biometric cryptosystems (BC [49, 53]). BC schemes, a subcategory of BTP schemes, often use primitives like secure sketches, fuzzy vaults, fuzzy commitments, or fuzzy extractors to either protect an existing cryptographic key or derive one directly from biometric data. These primitives can validate the correctness of reconstructed data or keys using a cryptographic hash function: the hash of the original key

or biometric data is stored during enrollment and later compared to the hash of the reconstructed key or data. If the hashes match, the reconstruction is validated. However, possessing both the helping data and the hash enables offline attacks, where adversaries can brute-force or reconstruct biometric data or keys without system interaction. While the search space is limited by the entropy of the biometric data, offline attacks remove network latency, allowing unlimited, high-speed guesses. In some cases, the secrecy of non-revocable templates depends on an independent secret key, whose compromise can lead to the complete compromise of the database.

Biometric recognition systems follow standards like ISO/IEC 24745 [27] and ISO/IEC 30136 [28], which define four security criteria: *irreversibility*, *unlinkability*, *revocability*, and *performance preservation*. Secure biometric systems achieve these goals through a chain of treatments. Below, we contextualize our work in relation to these principles:

- **Irreversibility and revocability in protected templates:** Secure biometric templates must ensure *irreversibility*, preventing reconstruction of the original biometric data from the (protected) template, and *revocability*, enabling replacement of compromised templates. Revocability requires randomizing features using a replaceable and secret user-specific token (e.g., stored secret or password)¹. However, not all irreversible methods allow revocability¹, even in secure biometric systems—where data are stored in a form that preserves their secrecy, and comparisons are performed securely. The sequence of treatments may include cancelable biometrics (CB [34, 41]), a subcategory of BTP schemes, that map biometric features into a *new metric space* using a secret-based one-way transformation applied client-side. This process enables comparisons of transformed features instead of raw data, improving entropy, supporting template revocation, and mitigating risks such as impersonation and reconstruction of raw features. Non-revocable templates (*i.e.*, those not randomized client-side) are not uniformly distributed, making them more susceptible to exhaustive search attacks. In this work, we demonstrate that *untargeted attacks* can be performed efficiently when the sequence of treatments excludes such a revocable transformation. This underscores the importance of client-side randomization.
- **Unlinkability across systems:** This criterion ensures that biometric templates from different systems cannot be correlated, preserving user privacy. However, reversing part of the templates through *untargeted attacks* may act as an initial step toward such linkability.
- **Efficient performance preservation:** Security measures should preserve the biometric system's *accuracy*, *speed*, and *scalability*. Our findings characterize this trade-off, showing how larger databases increase collision risks, thereby reducing the *accuracy* in the *identification mode*.

¹For example, a hashed social security number is irreversible (up to its inherent entropy) but not revocable because the SSN serves as a personal identifier that is hard to change—even if compromised. Since the hash is entirely determined by this sensitive input, it cannot be replaced if exposed. Adding salt prevents attackers from reusing precomputed tables (e.g., rainbow tables) but does not eliminate vulnerability to exhaustive search attacks, as each salt is public and each salted hash can still be brute-forced independently.

Contributions. This paper investigates the security of biometric systems, focusing on two critical aspects: untargeted attacks and the *biometric birthday problem*. By analyzing the False Match Rate (FMR) and its confidence interval—a range of values that likely contains the true FMR with a specified probability—we derive indicators to support the design of secure biometric systems. These indicators are: (i) *the critical population*, which represents the largest number of entries a biometric system can handle while satisfying a predefined security level, and (ii) *the critical FMR*, which is the highest FMR required to uphold that security level. Both indicators explicitly account for operational constraints to ensure their practical applicability. The specific contributions are detailed below:

- (i) **Untargeted Attack Analysis Based on FMR.** This part focuses on the security implications of *untargeted attacks* against biometric systems. Using the FMR and its confidence interval, we derive the complexity of such attacks and their impact on system design, particularly in relation to the size of the user databases:
 - We analyze the feasibility of untargeted attacks in biometric systems, focusing on a *non-adaptive attacker*, as this approach enables the derivation of general recommendations applicable across different types of biometric systems and modalities. Adaptivity, in contrast, requires case-by-case studies specific to the system and modality, which are beyond the scope of this analysis.
 - From this analysis, we derive the *critical population i.e.*, the maximum database size that guarantees a predefined security level against untargeted attacks.
 - We introduce the *critical FMR*, the maximum acceptable FMR required to maintain the predefined security level for a given database size.
 - These results emphasize the relationship between FMR, database size, and system security, providing practical recommendations for designing and deploying secure biometric systems at scale.
 - The analysis is complemented by detailed graphs that illustrate the relationships between the parameters.
- (ii) **Analysis of the Biometric Birthday Problem.** This part delves into the *biometric birthday problem*, exploring both its approximate and exact formulations. By integrating the confidence interval on the FMR, we provide a refined analysis of collision probabilities and their impact on biometric system design:
 - We provide a rigorous explanation of the approximate *biometric birthday problem*, originally introduced by John Daugman [12–14]. His analysis relies on the strong independence assumption that the absence of a collision between a user *A* and another user *B* does not influence the probability of a collision between *A* and any other user *C*. We refine his approximation by incorporating the confidence interval on the FMR.
 - From this analysis, we compute the approximate *critical population*, where weak collisions—two different biometric templates being incorrectly matched, as in the *birthday problem*—become significant, given a predefined security level and the confidence interval on the FMR.

- We define the *critical FMR* for weak collisions, offering guidance to reduce collision risks in large-scale deployments.
- For the first time, we introduce the exact *biometric birthday problem*, which overcomes the strong independence assumption made by Daugman. This analysis provides a precise estimation of collision probabilities and confirm the approximation of Daugman.
- These results are complemented by detailed graphs visualizing collision probabilities, *critical populations*, and *critical FMR* values, offering practical insights for system designers.

Related Works. Attacks on biometric systems are often analyzed on a case-by-case basis, reflecting the wide variety of secure biometric systems proposed in the literature. However, there has been limited exploration of generic attack methods or their effectiveness across diverse systems. An example of such generic attacks is *targeted attacks*, where an adversary seeks to impersonate a specific individual within a database. Pagnin *et al.* [39] introduced the "center search attack," a method that iteratively refines queries using minimal information leakage, such as a single bit indicating a match or mismatch, to locate specific templates. Their framework is built around binary templates and the Hamming distance as the underlying similarity metric. While this study provides valuable insights, its scope is constrained to specific template representations and similarity measures, limiting its relevance to more generalized biometric systems.

In contrast, *untargeted attacks*, where an outsider aims to impersonate any user within a database, remain unexplored in the literature. Unlike targeted attacks, these attacks exploit the probabilistic nature of biometric matching, leveraging statistical properties such as the FMR to increase the likelihood of unauthorized access. The FMR quantifies the probability that a random biometric input will match a stored template, making it a critical parameter for assessing the vulnerabilities of biometric systems. Furthermore, due to the non-uniform distribution of biometric data, *untargeted attacks* are expected to perform significantly better than if the adversary were testing against uniformly distributed data. The increasing availability of biometric data, driven by advancements in generative artificial intelligence and synthetic data generation, further complicates these vulnerabilities.

Among the earliest tools for generating synthetic biometric data, SFinGe [9], has been widely used to generate synthetic fingerprint datasets, enabling the evaluation of fingerprint recognition algorithms under controlled conditions. Its ability to simulate diverse fingerprint patterns and variations has made it a valuable tool for testing and benchmarking in biometric research. Building on these foundations, Bahmani *et al.* [4] proposed a method for generating high-fidelity fingerprints, focusing on improving quality, uniqueness, and privacy. These advancements have enhanced the utility of synthetic data for fingerprint recognition and system evaluation.

The emergence of generative adversarial networks (GANs) has further broadened the scope of synthetic biometric data generation. Engelsma *et al.* [17] introduced PrintsGAN, which leverages GANs to generate fingerprints with improved realism and diversity, addressing limitations in earlier methods. Beyond fingerprints, generative models have extended to other biometric modalities.

Kim *et al.* [31] proposed DcFace, a synthetic face generator based on a dual-condition diffusion model, capable of producing highly realistic and diverse facial images. These tools provide valuable resources for augmenting training datasets and testing machine learning models under varied conditions. While these synthetic data generators have advanced the state of biometric system evaluation, they also introduce potential security risks. The ease of creating realistic biometric data lowers the barrier for malicious actors to conduct attacks, such as exhaustive searches or the creation of forged templates. These developments underscore the importance of studying how system parameters, such as the FMR, interact with synthetic data and affect the robustness of biometric systems.

Biometric deployments at national or global scales introduce significant challenges, particularly regarding the *biometric birthday problem*. Analogous to the classical *birthday problem*, this problem quantifies the likelihood of collisions between biometric templates as the size of the database increases. For example, nationwide systems like Aadhaar in India and facial recognition initiatives in China [1, 50], each managing biometric data for over 1.4 billion individuals, face substantial collision risks that threaten uniqueness. Addressing these risks requires robust entropy management and careful parameterization, especially for large-scale deployments where we should maintain the balance between accuracy and security.

The *biometric birthday problem*, first introduced by Daugman [11–14], was briefly discussed in the context of iris recognition. Daugman’s work, while insightful, is limited in scope and relies on simplifying assumptions, such as the independence of collisions between templates. These assumptions, while practical for mathematical modeling, fail to capture structural dependencies and non-uniform distributions that often occur in real-world biometric systems. Building on Daugman’s contributions, we provide both an exact formulation of the *biometric birthday problem* and a refined analysis of its approximate formulation. By removing the independence assumption, our study generalizes the problem to a wider range of biometric systems, addressing practical complexities often overlooked. This broader analysis provides a more precise understanding of collision risks in large databases and aims to offer practical guidance for mitigating these risks across diverse biometric modalities.

Outline. The paper is structured as follows. Section 2 introduces the key notations and concepts necessary for understanding the subsequent sections. Section 3 presents the analysis of *untargeted attacks*. Section 4 explores the problem of collisions in biometric systems. Finally, Section 5 concludes the paper.

2 Preliminaries

This section introduces the foundational concepts and theorems necessary for the remainder of this paper. First, we detail the calculation of the False Match Rate (FMR) and discuss the construction and interpretation of its confidence intervals. Next, we present the attack models, focusing on both *targeted* and *untargeted* scenarios. Finally, we address the issue of collisions and detail their different types, including those related to the *biometric birthday problem*.

2.1 False Match Rate

In the context of authentication systems, specifically in a one-to-one system, the predominant metric utilized is the False Match Rate (FMR). This rate represents the probability of a biometric sample being incorrectly recognized by the matcher. In other words, this is the probability that the matcher incorrectly determines that a newly collected template matches the stored reference. The practical signification of the FMR lies in its impact on the security and usability of biometric systems. A high FMR implies a greater risk of unauthorized access, which is problematic for high-security applications and security-sensitive environments such as financial institutions or military facilities. As the FMR is directly linked to the security of the biometric system, it is important to quantify its value correctly.

Formal Definition. According to the Face Recognition Technology Evaluation (FRTE) 1:1 Verification [21], given a vector of n imposter scores v and T a threshold, an estimation of the FMR is

$$\overline{\text{FMR}} = \frac{1}{n} \sum_{i=1}^n H(T - v_i) = \frac{\text{Number of false matches}}{\text{Total number of comparisons}} \quad (1)$$

with $H(\cdot)$ the unit step function, and $H(0)$ taken to be 1. $H(T - v_i)$ can only take two values *i.e.*, 0 or 1 meaning that either the corresponding pair of user match or not.

It is important to highlight that, based on the FMR calculation method, the maximum security a system can achieve is limited to the inverse of the total number of comparisons. In other words, if we consider the entire human population (*i.e.*, 8×10^9) with one sample per individual, the smallest observable FMR is approximately $10^{-19.5}$. While this may seem counterintuitive, it underscores the fact that certain systems may have an accuracy level so sharp that it cannot be effectively measured using this method. For security-critical applications, it may be theoretically desirable to achieve an FMR this small; however, current technologies are not yet capable of reliably measuring or observing such a low FMR. Indeed, the best-performing biometric systems, such as those tested by the National Institute of Standards and Technology (NIST), can achieve extremely low FMRs. For example, state-of-the-art systems, including facial recognition technologies, have been recorded with FMRs as low as 10^{-8} and even $10^{-9.6}$, corresponding to a false match probability of one in 40 billion comparisons [21–23]. On the other hand, widely used commercial systems, such as Apple’s FaceID [2] and TouchID [3], and Google’s AndroidID [20], typically report an FMR on the order of 10^{-6} .

Experimental Estimation and Confidence Interval. As shown by Schuckers [48], it is possible to compute confidence interval in biometric metrics such as FMR. An empirical estimate of the FMR, denoted by $\overline{\text{FMR}}$, can be calculated by dividing the number of false matches by the number of matches tested for a given system. By convention, if we consider that matching is a success then, $H(T - v_i) \sim X_i$ a random variable which follows a Bernoulli distribution of probability FMR. The above equation can be rewritten as $\overline{\text{FMR}} = \frac{1}{n} \sum_{i=1}^n X_i$. By using the Central Limit Theorem (see Theorem A.1 in Appendix A.1 for more details), we deduce that $\overline{\text{FMR}}$ follows a normal distribution of mean μ and variance σ^2/n . The estimation of μ is given by $\bar{\mu} = \frac{1}{n} \sum_{i=1}^n X_i = \overline{\text{FMR}}$. The estimation of the standard

deviation σ is given by

$$\bar{\sigma} = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (X_i - \overline{FMR})^2} = \sqrt{\frac{n}{n-1} \overline{FMR} (1 - \overline{FMR})}.$$

Using the confidence interval (see Theorem A.2 in Appendix A.2 for more details) on the normal distribution, we have $FMR = \overline{FMR} \pm c_\alpha \sqrt{\overline{FMR}(1 - \overline{FMR})/(n-1)}$ with c_α the quantile of the Student's distribution with $n-1$ degrees of freedom. Throughout the rest of the paper, we use the following notations for respectively the Confidence Interval Lower bound (CIL) and the Confidence Interval Upper bound (CIU): (i) $CIL_{n,\alpha}(\overline{FMR}) = \overline{FMR} - c_\alpha \sqrt{\frac{\overline{FMR}(1-\overline{FMR})}{n-1}}$ and (ii) $CIU_{n,\alpha}(\overline{FMR}) = \overline{FMR} + c_\alpha \sqrt{\frac{\overline{FMR}(1-\overline{FMR})}{n-1}}$.

2.2 Attack Models

This section provides a detailed description of the two main types of attacks on biometric systems, alongside an explanation of biometric collisions. The first type, known as *targeted attacks*, focuses on impersonating a specific user in a database. The second, *untargeted attacks*, targets any user within the database. Lastly, we address the concept of collisions within a biometric database to present a comprehensive overview of security-related concerns.

Targeted Attacks. An external attacker (*i.e.*, an individual not registered in the system) attempts to impersonate a specific user in the database. To achieve this, the attacker has access to all the information necessary for the attack, including the user ID and other relevant details. As previously noted, the probability of successfully impersonating a specific user is equal to the system's FMR. On average, the number of biometric attempts required for the attacker to succeed is $1/\text{FMR}$. The FMR should be minimized, ideally conforming to the standards of cryptographic systems, which today require at least 128 bits for a moderate level of security [6], and in the near future will correspond to the minimum level [5]. This would correspond to an FMR smaller than 2^{-128} . However, to the best of our knowledge, no existing biometric systems achieve this level of security. Moreover, *targeted attacks* can be performed both *online* and *offline*. In *online attacks*, the attacker interacts directly with a remote biometric system to test guesses. In *offline attacks*, the attacker tests a chosen biometric data against a locally available obfuscated database, allowing simulations without interacting with a remote system.

Untargeted Attacks. An external attacker seeks to impersonate any user within the database rather than a specific individual. As in *targeted attacks*, the attacker has access to any additional information required to carry out the attack. This scenario is commonly associated with identification systems [10, 15, 29, 30, 40, 44, 55], which reject users who are not known to the system. To illustrate *untargeted attacks*, consider a device (*e.g.*, a smartphone) with a biometric recognition system (*e.g.*, fingerprint recognition). A user may enroll multiple fingerprints or even those of other individuals on their device. Consequently, an attacker attempting to unlock the device does not need to match a specific enrolled biometric but only one of the enrolled templates. This makes *untargeted attacks* significantly easier than targeted ones. *Untargeted attacks*, however, are not limited to identification systems. They are equally

applicable to authentication systems, where the attacker attempts to match a guessed template with a specific enrolled identity. In this mode, the attacker must test each guessed template against the enrolled biometric data of a claimed identity, introducing additional constraints compared to the identification mode. Specifically, in authentication systems, the attacker must know or guess the claimed identity (*e.g.*, the username or account ID) associated with the enrolled biometric template yielding additional costs. As for the previous attack, *untargeted attack* may be performed in both *online* and *offline* setup.

Collisions. In the context of a biometric database containing N users, we distinguish between two types of collisions: *strong collisions* and *weak collisions*.

- A *strong collision* occurs when a specific user in the database can be impersonated by another individual. Formally, for a user u with biometric template b_u , a strong collision happens if there exists another user $v \neq u$ such that the biometric matching function satisfies $\text{Match}(b_u, b_v) = \text{True}$. The probability of a strong collision for a single user is equal to the FMR, which quantifies the likelihood of incorrectly matching two distinct biometric templates.
- A *weak collision* occurs when at least two users in the database can impersonate each other. Formally, for a database of N users with biometric templates b_1, b_2, \dots, b_N , a weak collision exists if there exist two distinct users $i \neq j$ such that $\text{Match}(b_i, b_j) = \text{True}$. The probability of weak collisions increases with the size of the database and is related to the "*biometric birthday problem*" [13].

The probability of weak collisions depends on the FMR and the database size. A low weak collision probability indicates robust system accuracy, reducing risks of misidentification and impersonation by malicious users. Conversely, a high weak collision probability highlights vulnerabilities, particularly in large-scale databases. Estimating these probabilities through probabilistic modeling or simulations helps optimize system parameters and supports the design of secure, scalable biometric systems.

Ideally, both probabilities should be minimized to ensure robust security. While strong collisions can often be controlled by maintaining a sufficiently low FMR, the weak collision probability may become significant as the database size increases. In practical scenarios, a weak collision probability below 50% is often considered sufficient to maintain acceptable performance and reliability, but additional measures may be required to address risks associated with strong collisions in sensitive applications.

2.3 Experimental Setup

The numerical evaluations were performed in R [26] with a precision level of 200 bits, utilizing the `Rmpfr` [33] package to ensure high accuracy. The results were subsequently stored in CSV files, which were then processed using Python 3 [54] for graphical representation. The Python libraries NumPy [24], Pandas [37], and Matplotlib [25] were employed to generate the visualizations. The selection of the parameter values was made in such a way as to both reflect the current state-of-the-art in biometric systems and provide insight into the performance levels required to achieve high security and scalability.

3 Untargeted Attacker

In this section, we address two key topics. First, we derive the complexity of *untargeted attacks*, analyzing the resources required under various configurations. Second, leveraging this analysis, we compute two critical metrics: the maximum population size a system can securely accommodate at a specified FMR, and the minimum FMR required to support a given population while maintaining a desired security threshold.

Our study focuses on the identification mode, where an attacker attempts to match a query against a large database. In the authentication mode, however, the complexity of attacks scales by a factor of N , where N represents the total number of login attempts. In this scenario, the attacker repeatedly uses the same guess for each login attempt before proceeding to a new guess. This adjustment is conceptually similar to *password spray attacks*, where attackers exploit a large user base to attempt a single password guess per user, thus bypassing lockout mechanisms. Consequently, the *critical population size* and *critical FMR* derived here are left to the reader for recalculation to reflect this adjustment. Each of these findings is supported by extensive experiments, with tables and graphs illustrating the results and complementing the derived theoretical bounds.

3.1 Untargeted Attack Complexity

We analyze the median number of attempts required for a *non-adaptive attacker* to successfully gain acceptance in an identification system. In this context, the attacker aims to match one of the valid users from the database. This scenario can be modeled as follows: the attacker submits a query to the system, and, in this same round, this input is compared against all entries in the database.

To formalize this, let A_i represent the event that the attacker's input matches the i -th valid user in the database. The probability of A_i is given by the system's FMR. To account for practical variability, our analysis incorporates a confidence interval on FMR, reflecting the uncertainty or fluctuation in this metric due to measurement errors, unavoidable sampling errors, or inadequacies in the sampling frame or reference database.

In this context, we consider a *non-adaptive attacker*, which refers to an attacker that does not adapt its strategy based on the outcomes of previous attempts. This type of attacker submits queries independently across rounds, making it a general model applicable to any biometric system. While this assumption provides broad applicability, it is worth noting that in practice, for a specific biometric system and modality, an adaptive attacker could refine its guesses based on previous unsuccessful attempts. For example, the attacker might use feedback from the system to narrow down possible matches. However, such adaptive strategies depend on system-specific characteristics and are beyond the scope of this general analysis.

In this paper, we consider two distinct scenarios regarding the statistical dependence of the events A_i within the same round:

- **Independence of events A_i :** In this case, the events A_i are assumed to be mutually independent, meaning that the success of matching one user does not affect the probability of matching another user within the same round.

- **Dependence of events A_i :** Here, we introduce the notion of *collision dependence*, where the success of matching one user might influence the probability of matching others. It is important to note that this dependence is restricted to the same round, as we assume a *non-adaptive attacker* throughout the paper. This implies no auxiliary information is gained or exploited between rounds, ensuring independence across rounds.

Under this framework, and noting that the probability of not matching user i is given by $1 - \text{FMR}$, bounded within its confidence interval, we establish the following theoretical result regarding the median number of attempts required.

THEOREM 3.1 (BOUNDS ON THE COMPLEXITY OF THE UNTARGETED ATTACKER). *Let FMR denote the False Match Rate of a biometric recognition system, with an estimated value $\overline{\text{FMR}}$ obtained from n comparisons at a significance level α . The median number of attempts required for an untargeted attacker to successfully match one of the N users is bounded as follows:*

- (1) *Under the assumption of independence, the median number of attempts satisfies:*

$$\Omega\left(2^{-\log_2(\overline{\text{FMR}} + \overline{\text{FMR}}^2) - \log_2(N)}\right) \quad \text{and} \quad O\left(2^{-\log_2(\overline{\text{FMR}}) - \log_2(N)}\right).$$

- (2) *Under the assumption of non-independence, and for $\text{FMR} \leq \frac{1}{2N}$, the bounds are:*

$$\Omega\left(2^{-\log_2(\overline{\text{FMR}}) - \log_2(N^2 \cdot \overline{\text{FMR}}^2)}\right) \quad \text{and} \quad O\left(2^{-\log_2(\overline{\text{FMR}})}\right).$$

Accounting for the confidence interval, $\overline{\text{FMR}}$ is replaced by $\text{CIU}_{n,\alpha}(\overline{\text{FMR}})$ in the lower bound and by $\text{CIL}_{n,\alpha}(\overline{\text{FMR}})$ in the upper bound.

PROOF. Case 1: Independence Assumption. Let A_i represent the event that the attacker does not match the i -th user. Given a biometric system with performance FMR, the probability of A_i is $1 - \text{FMR}$. The events A_i are not necessarily independent, and by the chain rule of probability, we have:

$$\mathbb{P}(A_1 \cap A_2 \cap \dots \cap A_N) = \mathbb{P}(A_1) \prod_{i=2}^N \mathbb{P}(A_i | A_1 \cap A_2 \cap \dots \cap A_{i-1}).$$

Assuming independence between the events A_i , the probability that the attacker does not match any user is:

$$\mathbb{P}(A_1 \cap \dots \cap A_N) = \prod_{i=1}^N (1 - \text{FMR}) = (1 - \text{FMR})^N.$$

Let V_i be the event that the attacker matches one of the N users at the i -th round. Since the attacker is non-adaptive, the sequence of V_i events are independent Bernoulli trials, and we are interested in the first success. The median of a geometric distribution, representing the law of first success, is given by $m_{out} = \left\lceil \frac{-1}{\log_2(1-p)} \right\rceil$. The probability of success in a single trial is given by $p = 1 - (1 - \text{FMR})^N$. Using the series expansion of $\ln(1 - x)$ for $0 \leq x \leq 1/2$, we have:

$$\frac{-x - x^2}{\ln 2} \leq \log_2(1 - x) \leq \frac{-x}{\ln 2}.$$

From this, we find that $\log_2(1-p) = \log_2((1-\text{FMR})^N) = N \log_2(1-\text{FMR})$. Hence, we obtain the following bounds:

$$\frac{-N \cdot (\text{FMR} + \text{FMR}^2)}{\ln 2} \leq \log_2(1-p) \leq \frac{-N \cdot \text{FMR}}{\ln 2},$$

which lead to:

$$\frac{\ln 2}{N} (\text{FMR} + \text{FMR}^2)^{-1} \leq m_{out} \leq \frac{\ln 2}{N} \text{FMR}^{-1}.$$

Then, considering the fact that $\text{FMR} = \overline{\text{FMR}} \pm c_\alpha \sqrt{\frac{\overline{\text{FMR}}(1-\overline{\text{FMR}})}{n-1}}$, where n is the number of comparisons used to estimate FMR and c_α is the quantile of the Student's distribution for significance level α , the result follows.

Case 2: Non-Independence. The Bonferroni inequality [7] provides an upper bound for the probability of the union of events. Given N events A_1, A_2, \dots, A_N , the inequality states:

$$\mathbb{P}(A_1 \cap A_2 \cap \dots \cap A_N) \geq 1 - N + \sum_{i=1}^N \mathbb{P}(A_i).$$

Moreover, note that we have for all $1 \leq i \leq N$:

$$\mathbb{P}(A_1 \cap A_2 \cap \dots \cap A_N) \leq \mathbb{P}(A_i).$$

In this case, applying the Bonferroni inequality and the above inequality gives the following bounds:

$$\begin{aligned} 1 - N \cdot \text{FMR} &\leq 1 - p &\leq 1 - \text{FMR} \\ \log_2(1 - N \cdot \text{FMR}) &\leq \log_2(1 - p) &\leq \log_2(1 - \text{FMR}) \end{aligned}$$

if $\text{FMR} \leq 1/N$. Then, using the expansion series, we have:

$$\begin{aligned} \frac{-N \cdot \text{FMR} - N^2 \cdot \text{FMR}^2}{\ln 2} &\leq \log_2(1 - p) &\leq \frac{-\text{FMR}}{\ln 2} \\ \frac{-N \cdot \text{FMR} + N^2 \cdot \text{FMR}^2}{\ln 2} &\leq m_{out} &\leq \frac{\ln 2}{\text{FMR}} \end{aligned}$$

if $\text{FMR} \leq 1/2N$. Then, considering the fact that $\text{FMR} = \overline{\text{FMR}} \pm c_\alpha \sqrt{\frac{\overline{\text{FMR}}(1-\overline{\text{FMR}})}{n-1}}$, where n is the number of comparisons used to estimate FMR and c_α is the quantile of the Student's distribution for significance level α , we have:

$$\frac{\ln 2}{N \cdot \text{CIU}_{n,\alpha}(\overline{\text{FMR}}) + N^2 \cdot \text{CIU}_{n,\alpha}(\overline{\text{FMR}})^2} \leq m_{out} \leq \frac{\ln 2}{\text{CIL}_{n,\alpha}(\overline{\text{FMR}})}$$

and the result follows. \square

The results presented in Theorem 3.1 show that the security of biometric systems is primarily determined by two factors: the False Match Rate ($\overline{\text{FMR}}$) and the number of users N . It is clear that as expected, systems with lower FMR and fewer users are inherently more secure.

When analyzing the lower bounds, the first term in the exponent is $-\log_2(\text{CIU}_{n,\alpha}(\overline{\text{FMR}}) + \text{CIU}_{n,\alpha}(\overline{\text{FMR}})^2)$ for the first case and $-\log_2(\text{CIU}_{n,\alpha}(\overline{\text{FMR}}))$ for the second case. Considering only this first term, the median number of rounds required for an attacker is smaller in the first case than in the second. Examining the second term, $-\log_2(N)$ in the first case and $-\log_2(N^2 \text{CIU}_{n,\alpha}(\overline{\text{FMR}}))$ in the second case, we observe that when $\text{CIU}_{n,\alpha}(\overline{\text{FMR}}) \leq 1/N$ (a scenario likely to hold as in this case $\text{FMR} \leq 1/2N$), the number of rounds for the first case remains smaller than that of the second. Consequently,

for the lower bound, the first case grants the attacker a greater advantage, requiring fewer attempts to succeed. Conversely, for the upper bound, it is clear that the number of rounds in the first case is consistently smaller than in the second. Therefore, when seeking to establish a pessimistic security bound, it is appropriate to rely on the lower bound derived from the first case.

REMARK 3.1.1. No matter which specific case is considered, asymptotic bounds on the median number of rounds required for an untargeted attacker to succeed are as follows:

$$\Omega\left(2^{-\log_2(\overline{\text{FMR}} + \overline{\text{FMR}}^2) - \log_2(N)}\right) \quad \text{and} \quad O\left(2^{-\log_2(\overline{\text{FMR}})}\right).$$

Incorporating the confidence interval, $\overline{\text{FMR}}$ is replaced by $\text{CIU}_{n,\alpha}(\overline{\text{FMR}})$ in the lower bound and by $\text{CIL}_{n,\alpha}(\overline{\text{FMR}})$ in the upper bound.

3.2 Critical Population for the Untargeted Attack

In this section, we derive the *critical population* size for a given biometric system. Specifically, we establish the maximum number of users that a database can securely manage while satisfying both security constraints and system parameters, such as the FMR and the collision probability.

THEOREM 3.2 (CRITICAL POPULATION AGAINST UNTARGETED ATTACKER). Consider a biometric system operating with an error rate FMR, estimated as $\overline{\text{FMR}}$ based on n comparisons at a significance level α . The critical population N , such that the median number of rounds required by an untargeted attacker to successfully impersonate any user is smaller than a given threshold S , is given by:

$$N = \frac{\ln 2}{S \left(\overline{\text{FMR}} + \overline{\text{FMR}}^2 \right)}.$$

When accounting for the confidence interval, $\overline{\text{FMR}}$ is replaced by the upper bound $\text{CIU}_{n,\alpha}(\overline{\text{FMR}})$.

PROOF. Let S be the security desired (e.g., 2^{128}) for a biometric system with a performance estimated by $\overline{\text{FMR}}$ as before. Thus, using the smallest lower bound from Theorem 3.1 (i.e., the lower bound in the first case), we require:

$$S \leq \frac{\ln 2}{N \cdot \left(\text{CIU}_{n,\alpha}(\overline{\text{FMR}}) + \text{CIU}_{n,\alpha}(\overline{\text{FMR}})^2 \right)}$$

which may be rewritten as

$$N \leq \frac{\ln 2}{S \left(\text{CIU}_{n,\alpha}(\overline{\text{FMR}}) + \text{CIU}_{n,\alpha}(\overline{\text{FMR}})^2 \right)}$$

and the result follows. \square

The theorem shows that the *critical population* size N is inversely proportional to the security threshold S and the term $\text{CIU}_{n,\alpha}(\overline{\text{FMR}}) + \text{CIU}_{n,\alpha}(\overline{\text{FMR}})^2$. Increasing S enhances system security by requiring more attempts for an attacker to successfully impersonate a user. However, this comes at the expense of a reduced N , imposing stricter limits on the system's scalability. While higher values of

S may be offset by reducing $\text{CIU}_{n,\alpha}(\overline{\text{FMR}}) + \text{CIU}_{n,\alpha}(\overline{\text{FMR}})^2$, practical constraints on the FMR limit this compensation.

REMARK 3.2.1. *If the critical population size N falls below 1, the desired security threshold S is unattainable with the system's current accuracy, indicating the need for significant improvements in system performance or a relaxation of the security requirements. When N lies between 1 and 2, the system ensures adequate security for single-user scenarios and should be consistent with the targeted attack. Notably, if N ranges from 3 to 10, the system is capable of securely managing small-scale deployments. Such scenarios are typical in consumer devices like smartphones, where biometric systems frequently support the enrollment of 5 to 10 fingerprints for user authentication or identification.*

3.3 Critical $\overline{\text{FMR}}$ for the Untargeted Attack

As shown in the previous section, the FMR plays a key role in ensuring that the median number of attempts required for an attacker to succeed aligns with the desired security level for a system supporting a given number of users. In this section, we shift our focus to determining the necessary error rate (FMR) to ensure that a biometric system managing N users achieves S bits of security against *untargeted attacks*. This analysis provides practical guidelines for configuring system performance parameters, particularly for large-scale deployments where security and scalability must be carefully balanced.

THEOREM 3.3 (CRITICAL ESTIMATED FMR). *Given a database with N users and a security level S against the untargeted attack, the biometric error rate FMR, estimated as $\overline{\text{FMR}}$ from n comparisons with a significance level α , must satisfy: $\overline{\text{FMR}} \leq \sqrt{\frac{1}{4} + \frac{\ln 2}{N \cdot S}} - \frac{1}{2}$. Accounting for the confidence interval, $\overline{\text{FMR}}$ is replaced by $\text{CIU}_{n,\alpha}(\overline{\text{FMR}})$.*

PROOF. Using Theorem 3.1, the security level S satisfies $S \leq \frac{\ln 2}{N(\text{FMR} + \text{FMR}^2)}$, leading to the quadratic inequality $\text{FMR}^2 + \text{FMR} - \frac{\ln 2}{N \cdot S} \leq 0$. Studying the function $f(x) = x^2 + x - \frac{\ln 2}{N \cdot S}$ and solving for its roots, we find $\text{FMR} = \frac{-1 \pm \sqrt{1 + \frac{4 \ln 2}{N \cdot S}}}{2}$. Since $\text{FMR} \geq 0$, we select the positive root, giving $\text{FMR} \leq \sqrt{\frac{1}{4} + \frac{\ln 2}{N \cdot S}} - \frac{1}{2}$. Hence, the result follows. \square

The inequality shows that the upper confidence interval for the FMR, $\text{CIU}_{n,\alpha}(\overline{\text{FMR}})$, is fundamentally constrained by both the desired number of users, N , and the desired security level, S . As the product $N \cdot S$ increases, $\text{CIU}_{n,\alpha}(\overline{\text{FMR}})$ diminishes. Importantly, N and S have equivalent influence on reducing $\text{CIU}_{n,\alpha}(\overline{\text{FMR}})$; neither parameter dominates the other, as both contribute equally to the requirement for a smaller FMR. Consequently, achieving higher N or higher S necessitates reducing $\text{CIU}_{n,\alpha}(\overline{\text{FMR}})$, which depends on improving the FMR and the statistical confidence of its estimation.

REMARK 3.3.1 (CONFIDENCE PARADOX). *It is important to emphasize the problem associated with obtaining a tight confidence interval on the FMR. For the upper bound $\text{CIU}_{n,\alpha}(\overline{\text{FMR}})$ to be meaningful, it must be smaller than the right-hand side of the inequality in Theorem 3.3. Simple approximations reveal that the number of comparisons, n , required to achieve this is roughly proportional to $N \times S$, where N is the population size and S the desired security level.*

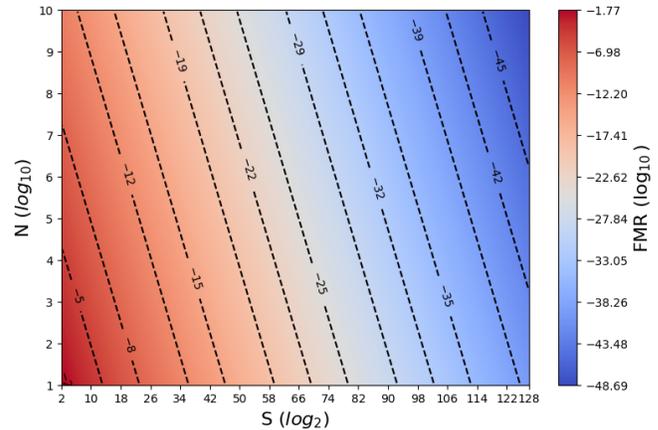


Figure 1: Critical FMR as a function of the population size N and the security level S against *untargeted attacks*.

This dependency leads to an exponential growth in the number of comparisons as N or S increases. For example, a system designed to handle 1000 users while achieving 50 bits of security would require approximately 2^{498} or 10^{150} comparisons, which is computationally infeasible.

Consequently, in such cases, confidence intervals on the FMR must be disregarded. This limitation implies that no guarantees on the required FMR can be reliably established, regardless of the chosen confidence level α .

3.4 Numerical Evaluations

This section presents a detailed numerical evaluation of critical parameters for biometric systems under untargeted attacks. We estimate the FMR required to maintain a specified level of security for a fixed population size, N . Note that, confidence intervals on the FMR are omitted, in alignment with earlier observations (see Appendix B for an insight on the impact of the confidence interval).

Most contemporary systems operate with an FMR around 10^{-6} . Figure 1 provides further insights. In the absence of specific security constraints, an attacker making four attempts to impersonate an individual limits the maximum manageable population size to approximately 10^6 . For small databases with only 10 users, the maximum achievable security level is 10 bits (2^{10} attempts required to compromise the system).

To meet cryptographic standards, systems must achieve at least 112 bits of security, with 128 bits becoming mandatory from 2030 onward [5, 6]. For small-scale databases (e.g., 10 users and 112 bits of security), achieving this standard requires an $\text{FMR} \leq 10^{-35}$. For large-scale systems accommodating 10^9 users and requiring 128 bits of security, the required FMR decreases further to 10^{-45} .

These results underscore the significant challenges biometric systems face in balancing scalability and cryptographic security. Achieving such low FMR values poses technical and computational hurdles, emphasizing the importance of advancing biometric technologies to meet these stringent requirements.

4 On the Weak Collisions

This section investigates the occurrence of collisions in biometric identification systems. Collisions are particularly problematic as they increase the probability of user misidentification, compromising the reliability of biometric-based authentication. This issue is especially critical in access control systems with hierarchical privileges, where a collision could lead to unauthorized access to sensitive areas or information.

We begin by analyzing the probability of collisions in both approximate and exact cases. For the approximate case, we build upon Daugman's framework [13], providing a rigorous mathematical estimation of collision probabilities. For the exact case, we derive the collision probability through detailed probabilistic modeling. Additionally, we identify the *critical population* size at which the collision probability surpasses a defined threshold. Furthermore, we derive the necessary $\overline{\text{FMR}}$ (*Critical FMR*) to ensure a given maximal probability of collision for a system with N users. Each of these points is supported by extensive experiments, with tables and graphs illustrating the results and complementing the derived theoretical bounds.

4.1 Approximated Biometric Birthday Problem

In this section, we derive an approximation for the probability of a collision in a biometric database, leveraging the concept of the *biometric birthday problem*, as proposed by various authors [11–14].

The *biometric birthday problem* extends the classical *birthday problem* (see [8, 35, 36] for more detailed) to biometric systems and is formally stated as follows: "Given a biometric system with a known FMR, how many randomly selected individuals must be considered before the probability of at least one pair experiencing a biometric collision (i.e., a false match) exceeds a given threshold p ?"

4.1.1 An Approximation for the Biometric Birthday Problem. First, we revisit Daugman's result [13] (Theorem 4.1) and provide a rigorous proof to substantiate his claims.

THEOREM 4.1 (APPROXIMATED BIOMETRIC BIRTHDAY PROBLEM). *Given a biometric system operating with an error rate FMR, an approximation of the probability that there is at least one collision in a database of N user is given by $\mathbb{P}(K) \approx 1 - (1 - \text{FMR})^{\mathcal{N}}$ with $\mathcal{N} = \frac{N(N-1)}{2}$.*

PROOF. Let M_i represent the event "The individuals within the i -th pair of $\overline{\mathcal{P}}$ do not falsely match," $Q(N)$ the probability of the event "There is no false match among N pairs," and $\mathbb{P}(N)$ the probability of the event "There is at least one false match among N pairs." We have $\mathbb{P}(M_i) = \text{FMR}$ and $\mathbb{P}(K) = 1 - Q(K)$.

Let $\overline{\mathcal{P}}$ denote the set of all possible unordered pairs among N individuals. Then, the total number of pairs of individuals is given by $|\overline{\mathcal{P}}| = \binom{N}{2} = \frac{N!}{2!(N-2)!} = \frac{N(N-1)}{2} = \mathcal{N}$. Then, we have:

$$Q(N) = \mathbb{P}\left(\bigcap_{i=1}^{|\overline{\mathcal{P}}|} M_i\right) = \mathbb{P}(M_1) \times \prod_{i=2}^{\mathcal{N}} \mathbb{P}\left(M_i \mid \bigcap_{j=1}^{i-1} M_j\right) \approx \prod_{i=1}^{\mathcal{N}} \mathbb{P}(M_i).$$

The above approximation is derived by assuming that M_i is independent of M_j for all $i > j$, which is an important assumption in

the result. It can be deduced that $Q(N) \approx (1 - \text{FMR})^{\mathcal{N}}$ by using the fact that $\forall i \in \{1, \dots, \mathcal{N}\}, \mathbb{P}(M_i) = 1 - \text{FMR}$. Then, we have that $\mathbb{P}(K) \approx 1 - (1 - \text{FMR})^{\mathcal{N}}$ which explains and correspond to the result given by Daugman [13] for the *biometric birthday*. \square

Then, Theorem 4.2 extend Daugman's results by incorporating confidence bounds on the FMR, enabling a more accurate estimation of this probability.

THEOREM 4.2 (APPROXIMATED SOLUTION FOR THE BIOMETRIC BIRTHDAY PROBLEM). *Given a biometric system operating with an error rate FMR estimated by $\overline{\text{FMR}}$ on n comparisons with a significance level α , an approximation of the probability that there is at least one collision in a database of N user is between*

$$\Omega\left(1 - \left(1 - \overline{\text{FMR}}\right)^{\mathcal{N}}\right) \text{ and } O\left(1 - \left(1 - \overline{\text{FMR}}\right)^{\mathcal{N}}\right)$$

with $\mathcal{N} = \frac{N(N-1)}{2}$. Accounting for the confidence interval, $\overline{\text{FMR}}$ is replaced by $\text{CIL}_{n,\alpha}(\overline{\text{FMR}})$ in the lower bound and by $\text{CIU}_{n,\alpha}(\overline{\text{FMR}})$ in the upper bound.

PROOF. Using Theorem 4.1 and Section 2.1, we have $\mathbb{P}(K) \approx 1 - \left(1 - \overline{\text{FMR}} \pm c_\alpha \sqrt{\frac{\overline{\text{FMR}}(1 - \overline{\text{FMR}})}{n-1}}\right)^{\mathcal{N}}$ with n the number of comparisons used to estimate the approximated FMR (i.e., $\overline{\text{FMR}}$) and c_α the quantile of the Student's distribution with $n - 1$ degrees of freedom. The result follows. \square

As the number of pairwise comparisons, $\mathcal{N} = N(N - 1)/2$, increases quadratically with the number of users N , the exponential terms $(1 - \text{CIL}_{n,\alpha}(\overline{\text{FMR}}))^{\mathcal{N}}$ and $(1 - \text{CIU}_{n,\alpha}(\overline{\text{FMR}}))^{\mathcal{N}}$ rapidly converge to zero as N grows. Consequently, the probability of a collision converges exponentially towards 1 when the number of user increases. This confirm the intuition that larger databases are inherently more susceptible to collisions, underscoring the importance of a small FMR in systems with a high number of user.

4.1.2 Approximated Critical Population. In this section, we investigate the notion of *critical population*, which refers to the maximum number of users that a biometric system can support while maintaining a specified FMR and collision probability. This concept provides practical guidelines for system design where collision risks must be mitigated.

Theorem 4.3 formalizes this relationship and establishes a rigorous framework for analyzing the trade-off between security requirements and user capacity.

THEOREM 4.3 (APPROXIMATED CRITICAL POPULATION). *Given a biometric system operating with an error rate FMR, estimated as $\overline{\text{FMR}}$ from n comparisons at a significance level α , the maximum number of users N such that the probability of a collision is smaller than a*

$$\text{chosen threshold } p \text{ is } N \approx \sqrt{\frac{2 \ln(1 - p)}{\ln\left(1 - \text{CIU}_{n,\alpha}(\overline{\text{FMR}})\right)}}.$$

PROOF. If we define p as the maximum probability that a collision occur, the following inequality can be written using Theorem 4.2:

$$1 - \left(1 - \text{CIU}_{n,\alpha}(\overline{\text{FMR}})\right)^{\mathcal{N}} \leq p$$

$$N^2 - N - 2 \frac{\ln(1-p)}{\ln\left(1 - \text{CIU}_{n,\alpha}(\overline{\text{FMR}})\right)} \leq 0.$$

The study of the above function in N yields

$$N \leq \frac{1}{2} + \frac{1}{2} \sqrt{1 + \frac{8 \ln(1-p)}{\ln\left(1 - \text{CIU}_{n,\alpha}(\overline{\text{FMR}})\right)}} \approx \sqrt{\frac{2 \ln(1-p)}{\ln\left(1 - \text{CIU}_{n,\alpha}(\overline{\text{FMR}})\right)}}$$

and the result follows. \square

Theorem 4.3 provides an approximation for the *critical population size*, N , given a target collision probability, p , and the conservative upper bound $\text{CIU}_{n,\alpha}(\overline{\text{FMR}})$ on the FMR. This result emphasizes the relationship between collision probabilities, system accuracy, and user capacity.

It is worth noting that FMR and p have opposite effects on the size of N . As p decreases, reflecting tighter security requirements, the *critical population* N diminishes. Conversely, looser requirements on p allow N to grow. On the other hand, reducing FMR allows the system to accommodate a larger population, while higher FMR values result in a decrease in population size. The delicate balance between security requirements and system scalability is underscored by this duality.

4.1.3 Required $\overline{\text{FMR}}$ for Target Population Sizes. A complementary approach to the collision problem involves reversing the perspective. Instead of determining the *critical population size* N for a given $\overline{\text{FMR}}$, the question becomes: for a target population size N , what upper bound on $\overline{\text{FMR}}$ is required to ensure that the probability of a collision remains below a specified threshold p ? This alternative perspective offers practical insights into the required system accuracy for handling large user bases at a given probability of collision. As a result, we formalize this relationship in Theorem 4.4.

THEOREM 4.4 (APPROXIMATED CRITICAL $\overline{\text{FMR}}$). *Given a database with N user and p the chosen maximal probability of a collision, the biometric error rate FMR estimated by $\overline{\text{FMR}}$ on n comparisons must be such that $\overline{\text{FMR}} \leq 1 - e^{\ln(1-p)/\mathcal{N}}$.*

PROOF. Using the same notations as above and Theorem 4.2, we start with $1 - \left(1 - \text{CIU}_{n,\alpha}(\overline{\text{FMR}})\right)^{\mathcal{N}} \leq p$. Taking the natural logarithm and rearranging terms, we obtain $\ln\left(1 - \text{CIU}_{n,\alpha}(\overline{\text{FMR}})\right) \geq \frac{\ln(1-p)}{\mathcal{N}}$. By exponentiating both sides, we get: $\text{CIU}_{n,\alpha}(\overline{\text{FMR}}) \leq 1 - e^{\frac{\ln(1-p)}{\mathcal{N}}}$. By observing that $\overline{\text{FMR}} \leq \text{CIU}_{n,\alpha}(\overline{\text{FMR}})$, this implies: $\overline{\text{FMR}} \leq 1 - e^{\frac{\ln(1-p)}{\mathcal{N}}}$. Yielding $\overline{\text{FMR}} \approx 1 - e^{\frac{\ln(1-p)}{\mathcal{N}}}$ for the maximal $\overline{\text{FMR}}$ needed if n is large enough, and the result follows. \square

This result provides an upper bound on the estimated $\overline{\text{FMR}}$ required to maintain a collision probability below a specified threshold p for a given database size N . The exponential term $\exp\left(\frac{\ln(1-p)}{\mathcal{N}}\right)$ emphasizes the relationship between the collision probability, the

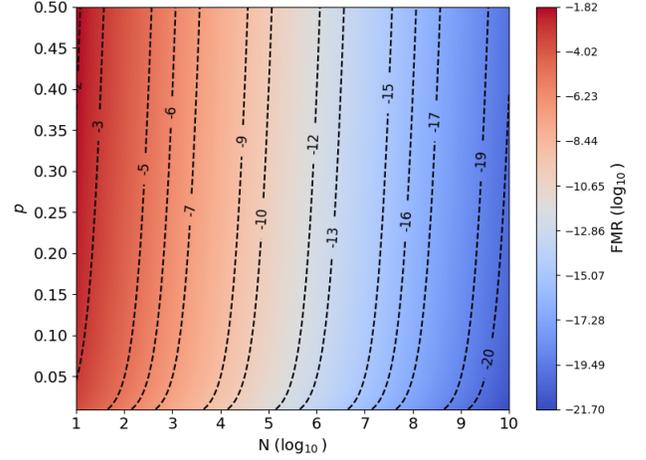


Figure 2: The maximal $\overline{\text{FMR}}$ (in \log_{10}) required to handle a population size N with a maximum allowed probability p for a collision occurrence.

population size, and the $\overline{\text{FMR}}$. As the number of pairwise comparisons $\mathcal{N} = \frac{N(N-1)}{2}$ increases, the required $\overline{\text{FMR}}$ becomes progressively smaller, highlighting the strong accuracy requirements for systems handling large user bases. For smaller values of p , reflecting higher security demands, FMR must decrease further to ensure the collision probability remains acceptably low.

4.1.4 Numerical Evaluations. To prevent weak collisions in a database with a given probability, we investigate the required FMR under specific design constraints. If the critical population size is to be derived from the FMR and the target collision probability, or conversely, if the collision probability must be determined for a given population size and FMR, these relationships are conveniently illustrated in the filled contour plot of the figure 2. To avoid redundancy, we focus below on selected numerical evaluations of particular interest.

First, let us recall that most contemporary systems operate with a false match rate (FMR) around 10^{-6} . Considering the classical *birthday problem*, where the collision probability is set to 1/2, such systems can accommodate approximately one hundred users.

In contrast, large-scale systems, such as those deployed in countries like China, must handle populations approaching 10^9 users. To maintain a low collision probability in such cases, the required FMR would need to be as low as 10^{-18} . Considering the lowest known FMR values (approximately 10^{-9}), current large-scale biometric systems inherently face a significant number of collisions.

An ideal biometric system should minimize the collision probability while supporting populations up to 10^{10} users to account for the global human population and effectively mitigate misidentification risks. Achieving this ambitious target would require an FMR of 10^{-20} , which represents a significant technological challenge for future system designs.

4.2 The Exact Biometric Birthday Problem

While the result provided in Section 4.1 gives an estimation, it is possible to compute the exact probability of the occurrence of a weak collision.

THEOREM 4.5. *Given a database with N users and operating with a False Match Rate estimated as $\overline{\text{FMR}}$ from K comparisons, the probability that there is at least one weak collision is:*

$$\mathbb{P}(N) = 1 - (1 - \overline{\text{FMR}}) \binom{(1 - \overline{\text{FMR}})\mathcal{K} - 1}{\mathcal{N}} \bigg/ \binom{\mathcal{K} - 1}{\mathcal{N}}.$$

PROOF. Let K denote the number of users used to compute the $\overline{\text{FMR}}$. Recall that we have

$$\overline{\text{FMR}} = \frac{\text{Number of pairs which falsely match}}{\text{Total number of pairs}}$$

where the total number of pairs is $\mathcal{K} = \frac{K(K-1)}{2}$,

$$\text{FMR} = \overline{\text{FMR}} \pm c_\alpha \sqrt{\frac{\overline{\text{FMR}}(1 - \overline{\text{FMR}})}{\mathcal{K}}}$$

and

$$Q(N) = \mathbb{P}\left(\bigcap_{i=1}^{\mathcal{N}} M_i\right) = \mathbb{P}(M_1) \times \prod_{i=2}^{\mathcal{N}} \mathbb{P}\left(M_i \mid \bigcap_{j=1}^{i-1} M_j\right).$$

Note that the probability $\mathbb{P}(M_{i+1} \mid \bigcap_{j=1}^i M_j)$ is the ratio of the number of remaining pairs that do not falsely match to the total number of remaining pairs. Specifically, the number of remaining pairs that do not falsely match is $\mathcal{K}(1 - \overline{\text{FMR}}) - i + 1$, while the total number of remaining pairs is $\mathcal{K} - i + 1$. Then, we have:

$$\begin{aligned} Q(N) &= (1 - \overline{\text{FMR}}) \prod_{i=2}^{\mathcal{N}} \frac{(1 - \overline{\text{FMR}})\mathcal{K} - i + 1}{\mathcal{K} - i + 1} \\ &= (1 - \overline{\text{FMR}}) \binom{(1 - \overline{\text{FMR}})\mathcal{K} - 1}{\mathcal{N}} \bigg/ \binom{\mathcal{K} - 1}{\mathcal{N}} \end{aligned}$$

and the result follows. It is important to note that if $(1 - \overline{\text{FMR}})\mathcal{K} - 1 \leq \mathcal{N}$ then, $Q(N) = 0$, as there are not enough remaining pairs to satisfy the condition. \square

The collision probability derived in Theorem 4.5 is influenced by three key parameters: the False Match Rate (FMR), the desired number of users (N), and the number of users used for estimating the FMR (K).

As the FMR increases, the probability of a collision also increases. Conversely, as FMR approaches zero, both the upper and lower bounds of the collision probability tend to zero, indicating a lower probability of collision with higher system accuracy.

It is important to note that when FMR is nonzero, and given that it is computed based on K users, as N approaches $(1 - \overline{\text{FMR}})K$, the collision probability increases accordingly.

Informally, this can be explained by the fact that the number of pairs not in a false match decreases. Thus, the probability of drawing a pair in a false match increases with N . Once all non-colliding pairs are exhausted, only colliding pairs remain, and the probability of a collision becomes one (see pigeonhole principle [52]). As a result, smaller databases tend to result in a lower collision probability,

consistent with earlier observations on the relationship between database size and security.

Finally, it is instructive to examine the behavior as $K \rightarrow \infty$, where the number of users used to compute the FMR becomes arbitrarily large. In this case, the collision probability approaches $1 - (1 - \overline{\text{FMR}})^{\mathcal{N}}$, which is the result derived in section 4.2. To be convinced, recall that $Q(N) = (1 - \overline{\text{FMR}}) \prod_{i=2}^{\mathcal{N}} \frac{(1 - \overline{\text{FMR}})\mathcal{K} - i + 1}{\mathcal{K} - i + 1}$ (see proof of Theorem 4.5). Then, as $K \rightarrow \infty$, this is equivalent to $Q(N) = (1 - \overline{\text{FMR}}) \prod_{i=2}^{\mathcal{N}} (1 - \overline{\text{FMR}})$ implying the result.

5 Conclusion

In this paper, we introduced the concept of *untargeted attacks*, wherein an adversary aims to impersonate any user within a database. We analyzed the complexity of such attacks, leading to the definition of the *critical population* (i.e., the maximum allowable database size) for a specified security level S and an estimated system performance $\overline{\text{FMR}}$. Furthermore, we derived the required $\overline{\text{FMR}}$ to achieve a desired security level for a given database size. Our experimental results reveal that current systems achieving $\overline{\text{FMR}} \approx 10^{-6}$ can offer at most 20 bits of security against *untargeted attacks* and fewer than 10 bits of security when the database exceeds 10 users. To ensure 128 bits of security for a system accommodating the global human population (i.e., approximately 10^9 users), an FMR smaller than 10^{-45} is necessary.

Additionally, we investigated the biometric *birthday problem* to study collision probabilities within a database. We determined both approximate and exact probabilities of weak collisions as functions of FMR and database size. From this analysis, we derived expressions for the *critical population* in terms of FMR and collision probability, as well as the required $\overline{\text{FMR}}$ for a given database size to maintain a specific collision probability. Experimental results indicate that systems with $\overline{\text{FMR}} \approx 10^{-6}$ can handle up to 100 users before the collision probability exceeds 50%. For databases accommodating the global human population, the $\overline{\text{FMR}}$ must be smaller than 10^{-18} to keep the collision probability below 50%.

In summary, this work provides an in-depth analysis of the FMR metric, revealing that biometric systems often considered secure may, in fact, fall short of expected security levels and more generally, to prevent both *untargeted attacks* and collisions, large scale systems must ensure $\overline{\text{FMR}} \leq 10^{-45}$.

The present work concludes that protecting biometric data relies on two important aspects. First, the design of truly secure (*zero-knowledge*) authentication protocols. Such protocols inherently satisfy the model of *minimal information leakage*. However, poorly parameterized systems or low entropy in biometric templates can still undermine security.

So, the second aspect is the enhancement of biometric data entropy. Higher entropy directly strengthens resistance to exhaustive search attacks, reducing the likelihood of successful *untargeted attacks*. Without sufficient entropy, adversaries can perform exhaustive searches to recover the original biometric features or approximate them, much like cracking weak password hashes. Techniques such as the fusion of multiple biometric modalities or the randomization of biometric data using a pseudo-random secret or a strong password can significantly mitigate these risks (see Appendix C).

We also identify techniques to slow down such attacks, inspired by strategies like password peppering and dedicated password hashing functions. These approaches include expanding the search space, effectively improving entropy, or designing obfuscation mechanisms and biometric cryptosystems that introduce artificial computational overhead, making the evaluation of a match more time-consuming.

Our theoretical model relies on the Central Limit Theorem, which assumes that the individual comparison tests used to compute the FMR are independent which may not be true in reality where multiple samples from the same individual may be used. Future work will focus on relaxing this independence assumption.

ACKNOWLEDGMENTS

We thank the reviewers for their feedback which improved the quality of this work. The authors acknowledge the support of the French Agence Nationale de la Recherche (ANR), under grant ANR-20-CE39-0005 (project PRIVABIO).

References

- [1] Shankkar Aiyar. 2017. Aadhaar: a biometric history of India's 12-digit revolution. *Westland Publications* (2017).
- [2] Apple. 2024. About Face ID advanced technology. <https://support.apple.com/en-us/102381>.
- [3] Apple. 2024. About Touch ID advanced security technology. <https://support.apple.com/en-us/HT204587>.
- [4] Keivan Bahmani, Richard Plesh, Peter Johnson, Stephanie Schuckers, and Timothy Swyka. 2021. High fidelity fingerprint generation: Quality, uniqueness, and privacy. In *2021 IEEE International Conference on Image Processing (ICIP)*. IEEE, 3018–3022.
- [5] Elaine B Barker. 2020. Recommendation for Key Management: Part 1 – General. *NIST Special Publication 800-57 Part 1 Revision 5* (2020). <https://doi.org/10.6028/NIST.SP.800-57pt1r5>
- [6] Elaine B Barker, Miles Smid, and Dennis Branstad. 2015. A Profile for US Federal cryptographic key management systems. *NIST Special Publication 800-152* (2015). <http://dx.doi.org/10.6028/NIST.SP.800-152>
- [7] Carlo Bonferroni. 1936. Teoria statistica delle classi e calcolo delle probabilità. *Publicazioni del R istituto superiore di scienze economiche e commerciali di firenze* 8 (1936), 3–62.
- [8] Mario Cortina Borja and John Haigh. 2007. The Birthday Problem. *Significance* 4, 3 (08 2007), 124–127. <https://doi.org/10.1111/j.1740-9713.2007.00246.x> arXiv:https://academic.oup.com/jrssig/article-pdf/4/3/124/49109395/sign_4_3_124.pdf
- [9] Raffaele Cappelli, Dario Maio, and Davide Maltoni. 2002. Synthetic fingerprint-database generation. In *2002 International Conference on Pattern Recognition*, Vol. 3. IEEE, 744–747.
- [10] Shaveta Dargan and Munish Kumar. 2020. A comprehensive survey on the biometric recognition systems based on physiological and behavioral modalities. *Expert Systems with Applications* 143 (2020), 113114.
- [11] John Daugman. 2003. The importance of being random: statistical principles of iris recognition. *Pattern recognition* 36, 2 (2003), 279–291.
- [12] John Daugman. 2015. Information theory and the iriscode. *IEEE transactions on information forensics and security* 11, 2 (2015), 400–409.
- [13] John Daugman. 2024. Understanding Biometric Entropy and Iris Capacity: Avoiding Identity Collisions on National Scales. *Adv. Artif. Intell. Mach. Learn.* 4, 2 (2024), 2152–2163. <https://doi.org/10.54364/AAIML.2024.42123>
- [14] John Daugman and Cathryn Downing. 2016. Searching for doppelgängers: Assessing the universality of the IrisCode impostors distribution. *IET Biometrics* 5, 2 (2016), 65–75.
- [15] Rodrigo de Luis-Garcia, Carlos Alberola-Lopez, Otman Aghzout, and Juan Ruiz-Alzola. 2003. Biometric identification systems. *Signal processing* 83, 12 (2003), 2539–2557.
- [16] W. Diffie and M. Hellman. 1976. New directions in cryptography. *IEEE Transactions on Information Theory* 22, 6 (1976), 644–654. <https://doi.org/10.1109/TIT.1976.1055638>
- [17] Joshua James Engelsma, Steven Grosz, and Anil K Jain. 2022. Printsgan: Synthetic fingerprint generator. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 5 (2022), 6111–6124.
- [18] Steven D Galbraith and Lukas Zobernig. 2019. Obfuscated fuzzy hamming distance and conjunctions from subset product problems. In *Theory of Cryptography Conference*. Springer, 81–110.
- [19] Steven D Galbraith and Lukas Zobernig. 2021. Obfuscating finite automata. In *Selected Areas in Cryptography: 27th International Conference, Halifax, NS, Canada (Virtual Event), October 21–23, 2020, Revised Selected Papers 27*. Springer, 90–114.
- [20] Google. 2024. Measuring Biometric Unlock Security. <https://source.android.com/docs/security/features/biometric/measure>.
- [21] Patrick Grother, Mei Ngan, Kayee Hanaoka, Joyce C. Yang, and Austin Hom. 2024. Face Recognition Technology Evaluation (FRTE), Part 1: Verification. <https://pages.nist.gov/frvt/html/frvt11.html>. NIST Interagency Report.
- [22] Patrick Grother, E Tabassi, GW Quinn, and W Salamon. 2009. Performance of iris recognition algorithms on standard images. *NIST Interagency Report 7629*, 1-120 (2009), 33–34.
- [23] Patrick J Grother, George W Quinn, James R Matey, Mei L Ngan, Wayne J Salamon, Gregory P Fiumara, and Craig I Watson. 2012. IREX III-Performance of iris identification algorithms. (2012).
- [24] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. 2020. Array programming with NumPy. *Nature* 585, 7825 (Sept. 2020), 357–362. <https://doi.org/10.1038/s41586-020-2649-2>
- [25] J. D. Hunter. 2007. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering* 9, 3 (2007), 90–95. <https://doi.org/10.1109/MCSE.2007.55>
- [26] Ross Ihaka and Robert Gentleman. 1996. R: a language for data analysis and graphics. *Journal of computational and graphical statistics* 5, 3 (1996), 299–314.
- [27] ISO/IEC24745:2011 2011. *ISO/IEC24745:2011: Information technology – Security techniques – Biometric information protection*. Standard. International Organization for Standardization.
- [28] ISO/IEC30136:2018(E) 2018. *ISO/IEC30136:2018(E): Information technology – Performance testing of biometric template protection scheme*. Standard. International Organization for Standardization.
- [29] AK Jain. 2007. Handbook of biometrics. *Springer google schola* 2 (2007), 114–116.
- [30] Anil Jain, Ruud Bolle, and Sharath Pankanti. 1996. *Introduction to biometrics*. Springer.
- [31] Minchul Kim, Feng Liu, Anil Jain, and Xiaoming Liu. 2023. Dcface: Synthetic face generation with dual condition diffusion model. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 12715–12725.
- [32] Wee Hong Loo, Paul HP Yeow, and Siong Choy Chong. 2009. User acceptance of Malaysian government multipurpose smartcard applications. *Government Information Quarterly* 26, 2 (2009), 358–367.
- [33] Martin Maechler, Maintainer Martin Maechler, SystemRequirementsNote MPFR, MASS Suggests, and MASS SuggestsNote. 2024. Package 'rmpfr'.
- [34] Manisha and Nitin Kumar. 2020. Cancelable biometrics: a comprehensive survey. *Artificial Intelligence Review* 53, 5 (2020), 3403–3446.
- [35] Frank H. Mathis. 1991. A Generalized Birthday Problem. *SIAM Rev.* 33, 2 (1991), 265–270. <https://doi.org/10.1137/1033051> arXiv:<https://doi.org/10.1137/1033051>
- [36] Earl H McKinney. 1966. Generalized birthday problem. *The American Mathematical Monthly* 73, 4 (1966), 385–387.
- [37] Wes McKinney et al. 2010. Data structures for statistical computing in python. In *Proceedings of the 9th Python in Science Conference*, Vol. 445. Austin, TX, 51–56.
- [38] Karthik Nandakumar and Anil K Jain. 2015. Biometric template protection: Bridging the performance gap between theory and practice. *IEEE Signal Processing Magazine* 32, 5 (2015), 88–100.
- [39] Elena Pagnin, Christos Dimitrakakis, Aysajan Abidin, and Aikaterini Mitrokovtsa. 2014. On the Leakage of Information in Biometric Authentication. In *INDOCRYPT 2014*. 265–280.
- [40] David Palma, Pier Luca Montessoro, et al. 2022. Biometric-based human recognition systems: an overview. *Recent Advances in Biometrics* 27 (2022), 1–21.
- [41] M. Vishal Patel, K. Nalini Ratha, and Rama Chellappa. 2015. Cancelable Biometrics: A review. *IEEE Signal Processing Magazine* 32, 5 (2015), 54–65.
- [42] Ursula Rao and Vijayanka Nair. 2019. Aadhaar: governing with biometrics. , 469–481 pages.
- [43] David G Rees. 2018. *Essential statistics*. Chapman and Hall/CRC.
- [44] Slobodan Ribaric and Ivan Fratric. 2005. A biometric identification system based on eigenpalm and eigenfinger features. *IEEE transactions on pattern analysis and machine intelligence* 27, 11 (2005), 1698–1709.
- [45] Phillip Rogaway and Thomas Shrimpton. 2004. Cryptographic hash-function basics: Definitions, implications, and separations for preimage resistance, second-preimage resistance, and collision resistance. In *Fast Software Encryption: 11th International Workshop, FSE 2004, Delhi, India, February 5-7, 2004. Revised Papers 11*. Springer, 371–388.
- [46] Sheldon M Ross. 2017. *Introductory statistics*. Academic Press.
- [47] Mulagala Sandhya and Munaga VNK Prasad. 2017. Biometric template protection: A systematic literature review of approaches and modalities. *Biometric Security and Privacy: Opportunities & Challenges in The Big Data Era* (2017), 323–370.
- [48] Michael E Schuckers. 2010. *Computational methods in biometric authentication: statistical methods for performance evaluation*. Springer Science & Business

Media.

[49] Shreyansh Sharma, Anil Saini, and Santanu Chaudhury. 2023. A survey on biometric cryptosystems and their applications. *Computers & Security* 134 (2023), 103458. <https://doi.org/10.1016/j.cose.2023.103458>

[50] Zhaohui Su, Ali Cheshmehzangi, Dean McDonnell, Barry L Bentley, Claudi-mar Pereira Da Veiga, and Yu-Tao Xiang. 2022. Facial recognition law in China. *Journal of Medical Ethics* 48, 12 (2022), 1058–1059.

[51] Clare Sullivan and Eric Burger. 2017. E-residency and blockchain. *Computer Law & Security Review* 33, 4 (2017), 470–481. <https://www.sciencedirect.com/science/article/pii/S0267364917300845>

[52] Wojciech A Trybulec. 1990. Pigeon hole principle. *Journal of Formalized Mathematics* 2, 199 (1990), 0.

[53] U. Uludag, S. Pankanti, S. Prabhakar, and A.K. Jain. 2004. Biometric cryptosystems: issues and challenges. *Proc. IEEE* 92, 6 (2004), 948–960. <https://doi.org/10.1109/JPROC.2004.827372>

[54] Guido Van Rossum and Fred L. Drake. 2009. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA.

[55] James L Wayman. 1997. A generalized biometric identification system model. *Conference Record of the Thirty-First Asilomar Conference on Signals, Systems and Computers (Cat. No. 97CB36136)* 1 (1997), 291–295.

[56] Stephan Wiefeling, Paul René Jørgensen, Sigurd Thunem, and Luigi Lo Iacono. 2022. Pump up password security! Evaluating and enhancing risk-based authentication on a real-world large-scale online service. *ACM Transactions on Privacy and Security* 26, 1 (2022), 1–36.

A Recall on Classical Results in Statistics

In this section, we provide a brief overview of the statistical results used in this paper.

A.1 Central Limit Theorem (CLT)

The Central Limit Theorem [46] (CLT) states that, under reasonable assumptions about the unknown distribution of the data, and provided the sample size is sufficiently large, the distribution of the sample mean converges almost surely to a normal distribution. In other words, the normal distribution acts as a universal attractor in this context.

THEOREM A.1 (CENTRAL LIMIT THEOREM (CLT)).

Let (X_1, \dots, X_n) be a sequence of independent and identically distributed random variables with a finite expected value μ (i.e., $\mathbb{E}[X_i] = \mu$) and finite variance σ^2 (i.e., $\mathbb{V}[X_i] = \sigma^2 < \infty$). Then, as n approaches infinity, the sample mean

$$X = \frac{1}{n} \sum_{i=1}^n X_i$$

converges in distribution to a normal distribution $X \sim \mathcal{N}(\mu, \sigma^2/n)$, where $\mathcal{N}(\mu, \sigma^2/n)$ denotes the normal distribution with mean μ and variance σ^2/n .

A.2 Confidence Interval for a Normal Distribution with Unknown Mean and Variance

A confidence interval [43] is defined as an estimated range within which the true value of the parameter is expected to lie, based on a given confidence level. For example, if a 95% confidence interval for the mean of a normally distributed population is computed, we can be 95% confident that the true mean lies within this interval. Confidence intervals are most often computed for confidence levels of 95% or 99%, but any significance level α can be chosen within the interval $]0, 1[$.

THEOREM A.2 (CONFIDENCE INTERVAL FOR THE MEAN OF A NORMAL DISTRIBUTION WITH UNKNOWN MEAN AND VARIANCE). Let

df	Level of Significance							
	0.25	0.20	0.15	0.10	0.05	0.025	0.01	0.005
1	1.000	1.376	1.963	3.078	6.314	12.706	31.821	63.657
2	0.816	1.061	1.386	1.886	2.920	4.303	6.965	9.925
3	0.765	0.978	1.250	1.638	2.353	3.182	4.541	5.841
4	0.741	0.941	1.190	1.533	2.132	2.776	3.747	4.604
5	0.727	0.920	1.156	1.476	2.015	2.571	3.365	4.032
6	0.718	0.906	1.134	1.440	1.943	2.447	3.143	3.707
7	0.711	0.896	1.119	1.415	1.895	2.365	2.998	3.499
8	0.706	0.889	1.108	1.397	1.860	2.306	2.896	3.355
9	0.703	0.883	1.100	1.383	1.833	2.262	2.821	3.250
10	0.700	0.879	1.093	1.372	1.812	2.228	2.764	3.169
20	0.687	0.860	1.064	1.325	1.725	2.086	2.528	2.845
30	0.683	0.854	1.055	1.310	1.697	2.042	2.457	2.750
40	0.681	0.851	1.050	1.303	1.684	2.021	2.423	2.704
50	0.679	0.849	1.047	1.299	1.676	2.009	2.403	2.678
60	0.679	0.848	1.045	1.296	1.671	2.000	2.390	2.660
70	0.678	0.847	1.044	1.294	1.667	1.994	2.381	2.648
80	0.678	0.846	1.043	1.292	1.664	1.990	2.374	2.639
90	0.677	0.846	1.042	1.291	1.662	1.987	2.368	2.632
100	0.677	0.845	1.042	1.290	1.660	1.984	2.364	2.626
500	0.675	0.842	1.038	1.283	1.648	1.965	2.334	2.586
1000	0.675	0.842	1.037	1.282	1.646	1.962	2.330	2.581
5000	0.675	0.842	1.037	1.282	1.645	1.960	2.327	2.577
10000	0.675	0.842	1.036	1.282	1.645	1.960	2.327	2.576
∞	0.674	0.842	1.036	1.282	1.645	1.960	2.326	2.576
	75%	80%	85%	90%	95%	97.5%	99%	99.5%

Table 1: Quantiles c_α of the Student's t -distribution for various degrees of freedom (df) and significance levels (α).

$X \sim \mathcal{N}(\mu, \sigma^2)$ for unknown μ and σ . Then, given an estimation $\bar{\mu}$ of μ and $\bar{\sigma}$ of σ from n samples, the confidence interval on $\bar{\mu}$ with risk α gives

$$\mu \in \left[\bar{\mu} - \frac{c_\alpha \times \bar{\sigma}}{\sqrt{n}}, \bar{\mu} + \frac{c_\alpha \times \bar{\sigma}}{\sqrt{n}} \right].$$

c_α denotes the quantile of the Student's t -distribution T with $n - 1$ degrees of freedom (i.e., $\mathbb{P}_T(-c_\alpha \leq T \leq c_\alpha) = 1 - \alpha$). This is equivalent to

$$\mathbb{P} \left(\bar{\mu} - \frac{c_\alpha \times \bar{\sigma}}{\sqrt{n}} \leq \mu \leq \bar{\mu} + \frac{c_\alpha \times \bar{\sigma}}{\sqrt{n}} \right) = 1 - \alpha.$$

The values of c_α for various degrees of freedom (df) and significance levels (α) are provided in Table 1.

B Numerical Analysis of Untargeted Attacks with Respect to the Confidence Interval

We analyze the *critical population size*, N , as a function of the *False Match Rate* (FMR), the confidence interval bounds derived from the number of comparisons n , and the desired security level S . Figure 3 illustrates how varying the number of comparisons influences the critical population size. For example, systems performing $n = 10^3$ comparisons achieve *up to 10 bits of security* with a 95% confidence level and a critical population size below 10^3 . In contrast, systems performing $n = 10^8$ comparisons can achieve *up to 20 bits of security* under the same confidence level and critical population size. These results demonstrate that the confidence interval introduces significant variability in the critical population size a system can accommodate. Specifically, increasing the number of comparisons

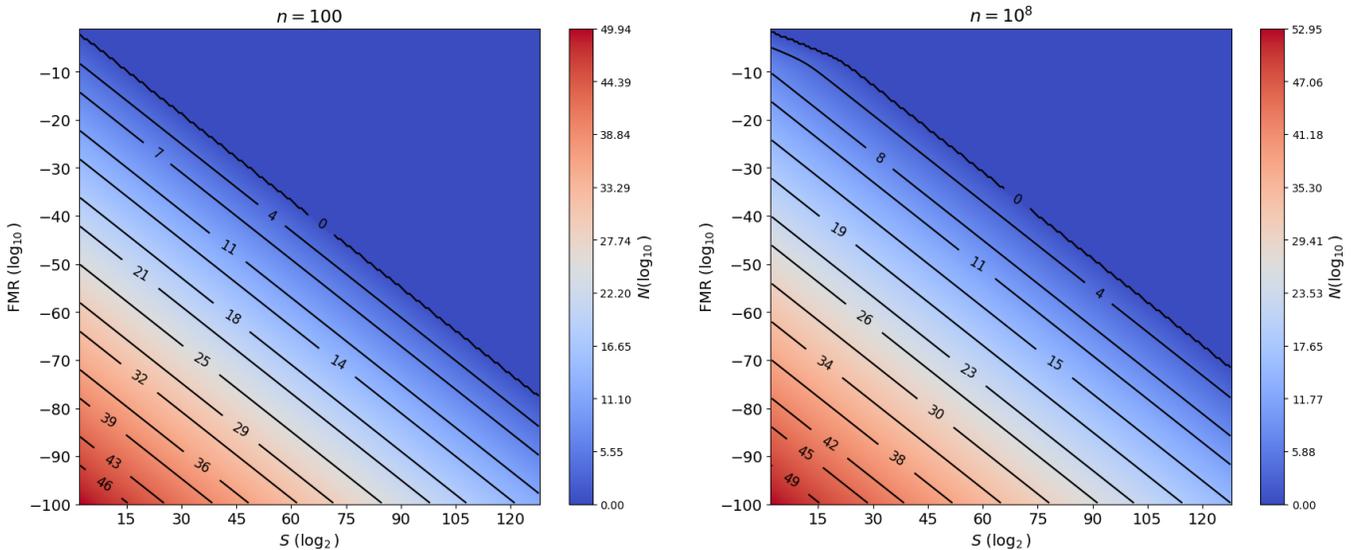


Figure 3: Critical population N as a function of the system accuracy $\overline{\text{FMR}}$ estimated on n comparisons with a significance level of 95% and the security level S against *untargeted attacks*.

by a factor of 10^5 results in a discrepancy of approximately two orders of magnitude (a factor of 100) in the critical population size.

C Multimodal Analysis

In this section, we describe how the security bounds for multimodal biometric systems can be derived from the single-modality analysis. Consider, for instance, a bimodal system that requires successful verification from both an iris scanner and a fingerprint sensor. Let p_1 denote the False Match Rate (FMR) for the iris modality and p_2 denote the FMR for the fingerprint modality. Under the assumption that the two verification processes are independent, the overall probability P_{attack} that an attacker succeeds in an untargeted attack is given by $P_{\text{attack}} = p_1 \times p_2$.

In the single-modality analysis, we relate the difficulty of an attack to the logarithm of the FMR. For example, if an attack requires an effort proportional to $\log(p_1)$ for the first modality and $\log(p_2)$ for the second, then the combined effort can be expressed as approximately $2^{\log(p_1) + \log(p_2)}$. This is equivalent to adding the efforts (in logarithmic scale) required to break each modality independently. In other words, the security bound for the multimodal system is the sum of the individual bounds.

Furthermore, if confidence intervals or other statistical measures are incorporated into the single-modality analysis, the same approach applies. The overall confidence interval for the multimodal system can be obtained by combining the confidence intervals of each modality under the assumption of independence. This additive property simplifies the extension of our theoretical framework from single to multiple modalities, providing straightforward bounds on the attack complexity in multimodal systems.