

Article

Two-Factor Cancelable Biometric Key Binding via Euclidean Challenge–Response Pair Mechanism

Michael Logan Garrett , Mahafujul Alam , Michael Partridge and Julie Heynssens

School of Informatics and Computing, Northern Arizona University, Flagstaff, AZ 86001, USA; ma3755@nau.edu (M.A.); mcp292@nau.edu (M.P.); julie.heyssens@nau.edu (J.H.)

* Correspondence: michael.garrett@nau.edu

Abstract

This work proposes a lightweight biometric key-binding scheme that adapts a PUF-style challenge–response mechanism to face geometry: a two-factor password and session nonce generate random challenge points, Gray-coded Euclidean distances to facial landmarks form responses, and a random key is bound by discarding selected positions so only a reduced subset, the nonce, and a key hash are stored. At authentication, a fresh response set is compared to the subset with a Hamming-distance tolerance, and bounded local search corrects residual errors; each successful session rotates the nonce and refreshes the ephemeral key. We frame this as a conceptual exploration of an interpretable, on-device, controlled-capture design niche—a per-session nonce-driven cancelable biometric key-binding mechanism—and we quantify the resulting security–usability trade-offs. Empirically, the scheme works under stable capture conditions with carefully tuned thresholds, and it is naturally suited to tightly controlled deployments (e.g., access kiosks) where it can also incorporate user-driven micro-gestures as an extra behavioral factor. While the construction is fragile under broader variability and leans on the second factor for security, it offers an alternative to existing mechanisms and a clear niche, and we present it as a conceptual exploration showing how CRP mechanisms can inform cancelable biometrics with per-session revocability.

Keywords: revocable biometrics; key binding; challenge–response pairs; physical unclonable functions; facial landmarks

1. Introduction

Biometrics offer the convenience of using the body as an identifier, but they come with an uncomfortable property: once compromised, they cannot be replaced. A fingerprint, an iris, or a face image that is leaked is lost forever as a secret. This permanence makes revocability one of the central challenges in biometric system design. Cancelable biometrics were introduced to answer that challenge by applying transformations that turn raw biometric data into renewable templates, such that a compromise can be addressed by issuing a new template instead of discarding the user’s biometric altogether.

In a different domain, PUFs have shown the power of CRPs as an interface for extracting secrets from noisy but differentiable systems. We connect those two ideas by presenting a biometric key-binding scheme that uses facial landmarks as the response mechanism in a PUF-style framework. A password and a session nonce generate random challenge points within a facial image frame, and the responses are defined as the Euclidean distances from these challenges to the detected landmarks. A random key is bound by



Academic Editor: Steven Furnell

Received: 12 January 2026

Revised: 15 February 2026

Accepted: 16 February 2026

Published: 2 March 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

discarding responses according to the key, and authentication reconstructs the key by comparing a fresh full set of responses against the stored subset. Each successful session generates a new nonce, thereby refreshing the challenges and re-enrolling the user with a new ephemeral key.

1.1. Research Question

Can a PUF-style challenge–response interface applied to facial landmark geometry support a two-factor key-binding protocol under tightly controlled capture conditions, and what security–usability trade-offs does such a design impose?

1.2. Contributions

1. We propose a two-factor biometric key-binding construction that adapts a PUF-style challenge–response interface to facial landmark Euclidean geometry, using a password and session nonce to derive challenges and responses.
2. We introduce an enrollment/key recovery flow that yields per-session template/key renewal without storing a long-lived similarity template, keeping only a reduced response subset, a nonce, and a key hash.
3. We adapt an existing recovery procedure that reconstructs an ephemeral key from noisy responses, and we characterize its computational cost as a function of the uncertainty budget.
4. We empirically map operating points induced by protocol parameters (window size/tolerance/uncertainty bounds) to key-recovery reliability and false-accept behavior under controlled capture.

2. Background

2.1. Cancelable Biometrics

The idea of cancelable biometrics arises in response to the awkward truth that you only get a couple of fingers, irises, and a single face, but you might need to revoke and reissue templates more than once in a lifetime. Early framing from the International Business Machines (IBM) group laid out the weak links in end-to-end biometric systems and sketched distort-then-match as a method to address them [1,2].

Faces and fingerprints were the initial focus for transformations, with key-based filtering for faces showing that you could convolve with a user-specific kernel, build a correlation filter, and still match while having something you could reissue if compromised [3], and with biotokens proposing to do matching in a protected space so that a lost template is just a token you can replace rather than your one and only face [4–6].

A large branch of the prior literature explores BioHashing and its variants, where you mix a biometric vector with a user token and project, quantize, and hash to obtain a compact and revocable representation; this mechanism was realized in Random Multispace Quantization [7,8]. The multispaces approach introduced diversity by applying different projections based on user-specific keys, while follow-ons tuned robustness and accuracy under less-than-clean conditions [9]. At the same time, the community also put a spotlight on the assumptions, thereby showing, for example, how reported near-perfect accuracy can collapse if the secret token or the specific transformation applied is not actually secret [10].

There is also a recurring theme in the existing body of work of using the protected template as a stepping stone toward a cryptographic key or an anonymous identifier rather than a match-score endpoint. Examples include key generation from noninvertible minutiae transforms, anonymous verification protocols for documents, and e-passport-flavored deployments that keep the template revocable [6,11,12].

Alongside construction papers, the state of the art built a much more thorough security understanding with models, attacks, and evaluation criteria: surveys and analyses catalog what “noninvertibility”, “diversity”, and “unlinkability” should mean and how to test them, and they document inversion, cross-matching, and record multiplicity threats that appear when multiple transformed records of the same user leak across systems [13–17].

Attack-centric works have kept us honest about risks such as pre-image attacks in known-algorithm settings, while theory publications pushed toward provable or even information-theoretic guarantees for noninvertible transforms that still support correlation-style matching, thereby expanding the space of what “matching in a protected space” can look like [18–20]. Taken together, recent and ongoing efforts in the field continue to reinforce the same core principle: if a template leaks, then you should not have to worry about the thief stealing your biometric identity, and you should be able to roll a new template without changing the person underneath [21–23].

2.2. PUFs and CRPs

PUFs popularized an interface where you do not store a secret so much as you query a noisy but consistent physical system, and read out a response indexed by a challenge [24]. While our setting is biometric rather than silicon, the interface (lots of CRPs, each repeatable with some mild noise) maps well onto a protocol of (1) presenting user plus auxiliary secret; (2) posing a challenge; and (3) harvesting a response, with some additional mechanisms layered on top to cope with the inevitable variation. In spirit, this is close to how biometric cryptosystems and cancelable transforms view the world, specifically that you transform or bind to something you can reissue and you handle noise with helper data, projection, quantization, sparse reconstruction, and/or error correction [14].

In our construction, the “challenge” is a location in the image frame derived from a second-factor password plus a nonce, and the “response” is the vector of quantized Euclidean distances from the challenge to a fixed set of landmarks, so the CRP analogy is literal at the API level even if the physics underneath is a human face rather than delay lines or Static Random Access Memory (SRAM) cells, which puts us in familiar territory with respect to design: you want lots of distinct challenges, you want responses to carry discriminative power while being tolerably noisy, and you want a revocation mechanism to change the challenge space without touching the underlying anatomy, themes consistent with cancelable templates, key-binding, and privacy-aware deployments [7–9,15,17,18,25–32].

Because biometrics are not unclonable functions in the strict hardware sense and because the adversary may know a lot about the transformation, security expectations must be stated modestly. In this proposed construction, the biometric itself is not an entropy source, but rather a noisy but reproducible authentication factor required to be present in order to reconstruct the key, while the second-factor password or token bears the burden of security. The point is not to claim that CRP-style indexing magically solves biometric privacy, but rather to suggest that it may offer a clean method to separate revocable indexing (the challenge set) from the variability of the response, and then to lean on an existing subset/key binding and recovery toolbox developed by Cambou and Ghanai et al. for diversity, revocability, and noise management [13,14,16,17,33,34] to design novel constructions worthy of study.

2.3. Positioning and Related Work

Our construction is best viewed as a biometric cryptosystem in the key-binding setting, with an explicit renewability mechanism: the stored binding record is revoked and refreshed on each successful authentication by regenerating the nonce-derived challenge set and re-binding a newly generated key. This differs from work that focuses primarily

on improving unconstrained biometric recognition accuracy; here, operating points arise from protocol parameters that trade off key-recovery reliability, false accepts, and bounded ternary recovery cost.

Prior work mostly falls into three overlapping categories. First, biometric cryptosystems and key-binding mechanisms such as fuzzy commitment/helper-data constructions bind a cryptographic key to a noisy biometric and recover it within tolerance [14,35,36]. Second, cancelable biometrics and template-protection methods emphasize non-invertible, renewable transforms to support the revocation and re-issuance of protected templates, a requirement formalized in standards like ISO/IEC 24745 [37]. Third, some systems explicitly combine these ideas and describe *cancelable key binding* (e.g., via cancelable transforms plus key binding) [38,39], which is the category that best fits the construction proposed herein. Some recent prominent work in these categories is briefly summarized in Table 1.

Table 1. Positioning this work against some existing biometric-crypto mechanisms.

Mechanism	Primitive	What It Uses	Typical Renewal/Re-Issuance
Fuzzy commitment [40]	Key binding	Biometric	Recommit/reissue helper data
Fuzzy vault [41,42]	Key binding	Biometric	Re-enroll; regenerate vault
Fuzzy extractors [43]	Key generation/key derivation	Biometric	Rerun GEN to reissue (R, P)
Cancelable biometric transforms [1]	Template protection	Biometric + transform parameter	Change transform/parameters
Cancelable iris key binding [38]	Cancelable key binding	Biometric + transform parameter	Change transform/parameters
This work	Two-factor CRP-based key binding	Biometric + user-held password/token	Per-authentication key/template refresh

Note: Many constructions optionally add a user secret for hardening; the “What it uses” column lists the base mechanism. Our construction explicitly requires a user secret.

3. Proposed Construction

3.1. System Overview

The system has two stages, enrollment and key recovery, and both rely on a common sequence of operations that take a two-factor password together with a fresh session nonce and a facial frame; it turns them into a response, which is a structured set of bitstrings (Figure 1). The shared steps are as follows:




Figure 1. A face capture is aligned, landmarked, and challenged to produce a complete geometric response set.

- 1. Gather inputs.** Collect the auxiliary password P , generate a random session nonce ν , and capture a frame or set of frames $F = \{I^{(t)}\}_{t=1}^q$ of the user’s face (each $I^{(t)}$ is a 256×256 grayscale image in our implementation).

$$(P, \nu, F) \longrightarrow \text{pipeline inputs.}$$

2. **Challenge generation.** Derive N integer pixel challenges in the frame using a PRF ϕ keyed by (P, ν) ; each challenge is a coordinate $c_i = (x_i, y_i)$ in a 256×256 grid.

$$C = \{c_i\}_{i=1}^N, \quad c_i = (\phi_x(i; P, \nu), \phi_y(i; P, \nu)), \quad x_i, y_i \in \{0, \dots, 255\}.$$

3. **Landmarking.** Run a landmark predictor Λ (e.g., Ref. [44]) on F to obtain k facial landmarks in frame coordinates and write the (unshuffled) shape as $S = \{\ell_j\}_{j=1}^k$ with $\ell_j = (u_j, v_j)$.

$$S = \Lambda(F) = \{(u_j, v_j)\}_{j=1}^k, \quad u_j, v_j \in \{0, \dots, 255\}.$$

4. **Landmark shuffle.** Use a PRF seeded by (P, ν) to produce a permutation π of $\{1, \dots, k\}$ and reorder landmarks accordingly (the protocol always operates on the shuffled order).

$$\pi = \pi_{P, \nu} \in S_k, \quad S^\pi = (\ell_{\pi(1)}, \dots, \ell_{\pi(k)}).$$

5. **Compute raw responses.** For each challenge c_i and shuffled landmark $\ell_{\pi(j)}$, compute Euclidean distances and assemble the k -response list $\mathbf{d}_i$.

$$d_{i,j} = \|c_i - \ell_{\pi(j)}\|_2 = \sqrt{(x_i - u_{\pi(j)})^2 + (y_i - v_{\pi(j)})^2}, \quad \mathbf{d}_i = (d_{i,1}, \dots, d_{i,k}) \in \mathbb{R}_+^k.$$

6. **Response scaling.** Map each distance to $\{0, \dots, 2^l - 1\}$ using a fixed, monotone, clipping normalization function σ .

$$\tilde{d}_{i,j} = \sigma(d_{i,j}; \theta) \in \{0, \dots, 2^l - 1\}, \quad \tilde{\mathbf{d}}_i = (\tilde{d}_{i,1}, \dots, \tilde{d}_{i,k}),$$

where θ denotes the scaling parameters $\{d_{min}, d_{max}\}$ computed from (C, S^π) for the session.

7. **Gray coding.** Convert each $\tilde{d}_{i,j}$ to its l -bit Gray code segment $\mathbf{g}_{i,j} \in \{0, 1\}^l$ via $g(x) = x \oplus (x \gg 1)$.

$$\mathbf{g}_{i,j} = g(\tilde{d}_{i,j}) \in \{0, 1\}^l, \quad \mathbf{G}_i = (\mathbf{g}_{i,1}, \dots, \mathbf{g}_{i,k}).$$

8. **Forming response block set R .** Concatenate the k Gray-coded segments to obtain a block $r_i \in \{0, 1\}^m$ of length $m = l \cdot k$, and collect all N blocks as the response set.

$$r_i = \mathbf{g}_{i,1} \parallel \mathbf{g}_{i,2} \parallel \dots \parallel \mathbf{g}_{i,k} \in \{0, 1\}^m, \quad R = \{r_1, \dots, r_N\}.$$

The two core ideas underlying the enrollment phase (binding by partitioning the response positions based on an ephemeral key) and the recovery phase (reconstructing the key with a guided, ternary recovery algorithm plus bounded local search) follow recent formulations by Ghanai et al. and Cambou et al. in distributed key reconstruction and subset-driven matching; we adopt those patterns here for a new application in a biometric setting while keeping a design that is comparable to prior analyses [33,34].

3.2. Enrollment

Enrollment begins by running the shared sequence with (P, ν, F) to produce the full response set R . A random N -bit key is then generated from a custom Key Derivation Function (KDF), which before binding imposes a constraint parameter ω that limits the maximum run length of consecutive zeros. If the initial random key contains a run longer than 0^ω , the KDF randomly flips one or more bits within that run to break it. This ensures that the recovery algorithm always knows an upper bound on the size of any consecu-

tive zero segment, which provides an initial window size for aligning responses during key recovery.

Once the key satisfies the ω constraint, then binding is performed by discarding all responses in positions where the key bits are set, leaving a subset of R (Figure 2). This subset, together with the current nonce, is stored as the user's template. The key itself is used to encrypt the user's data and is then discarded. As a correctness oracle for later verification, the system also stores a hash of the key alongside the subset and nonce. As a result, the stored record cannot by itself decrypt the data, and the key cannot be reconstructed without the auxiliary password, the correct nonce, and a matching biometric; the user's face geometry cannot be easily recovered from the enrollment data alone.




Figure 2. The set bits in the ephemeral key produce a subset of responses from R , which we use later to recover the key.

3.3. Key Recovery

At authentication time, the user again supplies their password and face, and the system uses the stored nonce to regenerate a fresh full set of responses; this new set is compared against the stored subset by the recovery matching algorithm (Algorithm 1, Figure 3), which reconstructs the key position by position, assigning 0 or 1 when a comparison is confident and -1 when it is not. Comparisons are made using a normalized Hamming distance on response blocks: for blocks $x, y \in \{0, 1\}^m$ with $m = l \cdot k$, define $\delta_H(x, y) \triangleq d_H(x, y)/m$, and declare a match when $\delta_H(x, y) \leq T$, where $T \in [0, 1]$ is the tolerance (e.g., $T = 0.10$ treats $\leq 10\%$ bit differences as a match), with the initial output therefore a ternary vector $\hat{K} \in \{-1, 0, 1\}^N$, whose -1 entries mark unresolved bits.




Figure 3. The recovery algorithm attempts to reconstruct the ephemeral key from the enrollment subset.

The use of the ω -constraint during enrollment ensures that the recovery algorithm has a predictable bound on the maximum run length of zeros, which ensures a fixed search window when aligning responses and reduces the number of positions that may remain uncertain in the first place by limiting how far the alignment can drift before a mismatch must be declared under the same $\delta_H \leq T$ rule.

The subsequent error correction step is then targeted: rather than brute-forcing over all nearby bitstrings, the search explores only the subset of possible true keys defined by the uncertain (-1) positions, flipping them all exhaustively between 0 and 1; the search is bounded by a maximum number of corrections, which acts as the radius of exploration, and each candidate binary key K' is checked against the stored hash H (the correctness oracle), after which a passing K' is accepted and used to decrypt the protected user data. No candidate within the bound leads to failure, and a successful session triggers a new nonce and a fresh re-enrollment with a new ephemeral key.

Algorithm 1: Ternary key recovery.

Input: Enrollment subset f , recovery set R' , max run parameter ω
Output: Uncorrected ternary key $\hat{K}$

```

 $\hat{K} \leftarrow \{0\}^N$  // Empty (zero) key
 $\alpha \leftarrow 0$  // Window start position
 $w \leftarrow \omega$  // Initial window size
foreach  $r'_i$  in  $f$  do
     $W \leftarrow R'[\alpha : \alpha + w]$  // Set the window
     $M \leftarrow \{j \mid \delta(r'_i, W_j) \text{ is TRUE } \forall W_j \in W\}$  // Indices of matches
    // Zero match case
    if  $|M| == 0$  then
        foreach  $m \in M$  do  $\hat{K}[\alpha + m] \leftarrow -1$ 
         $\alpha \leftarrow \alpha + 1$  // Move the window
         $w \leftarrow w + \omega - 1$  // Expand window size
    end
    // Single match case
    else if  $|M| == 1$  then
         $i_m \leftarrow M$  // Index of only match
         $\hat{K}[\alpha + i_m] \leftarrow 1$ 
         $\alpha \leftarrow \alpha + i_m + 1$  // Move the window
         $w \leftarrow \omega$  // Reset window size
    end
    // Collision case
    else if  $|M| > 1$  then
        foreach  $m \in M$  do  $\hat{K}[\alpha + m] \leftarrow -1$ 
         $\alpha \leftarrow \alpha + \min(M) + 1$  // Move the window
         $w \leftarrow w + \omega - 1$  // Expand window size
    end
end
return  $\hat{K}$ 
  
```

3.4. Formal Specification

Let P denote the auxiliary password and F a captured facial frame. For each session, we also generate a random nonce v . The pair (P, v) is used to derive a set of N challenges $C = \{c_1, \dots, c_N\}$, where each challenge $c_i = (x_i, y_i)$ is an integer coordinate with $x_i, y_i \in \{0, \dots, 255\}$. We also extract a shuffled set of $k = 68$ facial landmarks $S = \{\ell_1, \dots, \ell_{68}\}$, with each landmark $\ell_j = (u_j, v_j)$ an integer coordinate in the same 256×256 frame.

For each challenge c_i and landmark ℓ_j , we define the raw response distance

$$d_{i,j} = \sqrt{(x_i - u_j)^2 + (y_i - v_j)^2}.$$

Each $d_{i,j}$ is normalized to an l -bit integer $\tilde{d}_{i,j} \in \{0, \dots, 2^l - 1\}$ and then Gray-coded into a bitstring $g_{i,j} \in \{0, 1\}^l$. For challenge c_i , we concatenate across all landmarks to obtain a response block

$$r_i = g_{i,1} \parallel g_{i,2} \parallel \dots \parallel g_{i,k}.$$

The full response set is

$$R = \{r_1, r_2, \dots, r_N\}.$$

During enrollment, we sample a random key $K \in \{0, 1\}^N$ subject to the constraint that no run of more than ω consecutive zeros appears in K . If a generated key violates this constraint, runs longer than 0^ω are broken by randomly flipping one or more bits inside them. Binding is then performed by discarding each block r_i whenever $K_i = 1$, producing a stored subset

$$f = \{r_i : K_i = 0\}.$$

We then compute $H = h(K)$ using a cryptographic hash $h(\cdot)$, and the enrollment record for this session is (f, H, ν) . The key K is used once to encrypt user data M , producing ciphertext $C = \text{Enc}_K(M)$, and then discarded.

At authentication, the same procedure generates a fresh full response set $R' = \{r'_1, \dots, r'_N\}$; the recovery algorithm compares R' with the stored subset f to build a candidate key $\hat{K} \in \{-1, 0, 1\}^N$, where values 0 and 1 are confident assignments and -1 marks “uncertain” positions that could not be decided. We let $m = l \cdot k$ be the block length in bits and define $\delta_H(x, y) \triangleq d_H(x, y) / m$, and say that blocks x and y match if $\delta_H(x, y) \leq T$, with $T \in [0, 1]$ a fixed tolerance. The algorithm aligns R' to f using an ω -bounded window and applies this $\delta_H \leq T$ rule to assign bits (unique matches yield $\{0, 1\}$, ambiguous or failed matches yield -1), after which error correction explores substitutions only at the -1 positions and only up to a maximum number of corrections (the search radius), verifying each candidate binary key K' via $h(K') = H$, accepting on success or otherwise failing if no candidate within the radius passes, and upon success generating a new nonce ν' and restarting with a fresh enrollment based on R' . An average modern computer can easily correct a radius of about 10% of a 256-bit key, or about 25 ternary positions, before correction cost starts to become prohibitive.

3.5. Security Implications

This construction is two-factor in the strict sense that successful key recovery requires both (i) the second factor (password or token-derived secret that determines challenges and recovery context) and (ii) a biometric presentation that yields responses within the protocol’s tolerance setting. The biometric factor is not intended to contribute entropy to the session key. Instead, it provides the noisy but reproducible signal needed for key recovery to succeed.

We discuss security implications under the four unique cases of compromise. Throughout, the protocol’s “accept” decision is always a binary outcome: authentication succeeds if and only if the session key is recovered and matches the stored key hash. If the recovery process fails to recover the correct key, authentication fails. There is no “confidence” threshold to set, as we would expect in a traditional classifier; either the exact correct key is recovered or it is not. In all cases below, we assume the attacker can read the currently valid stored enrollment record (the unlabeled response subset, the nonce, and the key hash), and we distinguish whether the attacker also possesses the second factor and/or can present a biometric sample that matches within tolerance.

3.5.1. Case A: Biometric Data Compromised, Second Factor Remains Secret

If an adversary obtains biometric information about a user, but does not obtain the second factor, the adversary cannot derive the correct challenges needed to reproduce the responses corresponding to the enrolled response subset and recover the session key. In this case, biometric leakage alone is insufficient for offline reconstruction because the missing second factor prevents the adversary from even formulating the correct reconstruction attempt. This case captures the intended separation of roles: the biometric constrains reconstruction, while the second factor supplies the secrecy that gates access to the reconstruction procedure.

3.5.2. Case B: Second Factor Compromised, Biometric Remains Secret

If an adversary obtains the second factor but cannot present a biometric sample that yields responses within tolerance, then the adversary can compute the challenge material and can attempt offline key recovery against stored records. The security that remains in this case is governed by the false-accept behavior at the chosen operating point. The adversary succeeds only when a non-matching biometric presentation, or an attempted surrogate, nonetheless causes the recovery procedure to output a key whose hash matches the stored value. This is a stolen-token setting: the relevant limitation is that record multiplicity and repeated trials increase the adversary's cumulative success probability, so we must treat the measured FAR as the effective risk metric under compromise of the second factor. If the attacker in this case can mount t independent offline attempts against the currently valid record during its lifetime, then the cumulative success probability is $P_{\text{succ}}(t) = 1 - (1 - \text{FAR})^t$.

3.5.3. Case C: Both Biometric and Second Factor Compromised

If an adversary obtains both the second factor and a biometric sufficient to generate responses within tolerance, then the adversary can recover session keys and impersonate the user. In this case, the scheme provides no meaningful cryptographic protection. Mitigation is still possible via the revocation mechanism: the system can restore security by revoking and re-issuing the second factor and re-enrolling with a fresh biometric under a new nonce, subject to the constraints discussed above on record multiplicity and offline checking.

3.5.4. Case D: No Compromise

If neither factor is compromised, an attacker lacks both the second factor and a biometric presentation that yields responses within tolerance, so they cannot carry out a meaningful offline key-recovery attempt. Without the second factor, the attacker cannot reproduce the correct challenge set and recovery context, and, without a matching biometric, they cannot produce responses close enough to the enrolled responses for key recovery to succeed. In this baseline setting, the relevant security question is online impersonation risk under normal, reasonable rate-limiting assumptions and physical security conditions.

These cases clarify the intended interpretation of our empirical results. Under second-factor compromise, security reduces to the false-accept characteristics of the key recovery process given by the chosen protocol parameters (mostly the matching threshold). Under biometric compromise alone, security depends on the secrecy of the second factor and the resulting inability to make correct response reconstruction attempts. Under compromise of both, only revocation and re-issuance can restore security. Under compromise of neither, an adversary lacks both the correct second factor needed to instantiate the enrollment-specific challenges and a biometric sample that matches the enrolled user closely enough to permit key recovery, so they cannot mount a meaningful offline key-recovery attempt in our model and the probability of successful attack is negligible.

3.5.5. Record Lifetime and Refresh Policy

The stored hash of the session key functions as a correctness oracle for offline key-recovery attacks: an attacker who can run the key-recovery procedure can immediately test whether a candidate recovery output is correct by hashing it and comparing. This does not reveal the key itself, but it does mean that key recovery is a verifiable offline decision problem. Accordingly, the attacker's practical success probability for an offline key-recovery attack with a stolen enrollment is governed by the false-accept behavior of the recovery procedure at the chosen operating point and by the number of independent attempts they can mount before the enrollment they stole is invalidated by a legitimate authentication.

The intended deployment maintains a single active enrollment record per user. Upon successful authentication, the protocol cancels the prior enrollment and immediately creates a fresh one by generating a new nonce (hence new challenges) and a new session key, using the just-captured biometric sample as the basis for the refreshed response set and subset and binding the new session key thereto. This gives the mechanism epoch isolation with respect to the bound session key: an enrollment stolen in one epoch does not remain usable for recovering the current session.

3.5.6. Multi-Epoch Compromise and Biometric Accumulation Risk

The refresh policy limits the lifetime of any one enrollment record and its associated session key, but it does not eliminate the possibility of biometric accumulation if an attacker can exfiltrate enrollment records across many epochs. In that setting, the security question is no longer only whether the session key remains unrecoverable, but whether the repeated exposure of randomized challenge–response information can progressively constrain (and ultimately reveal) the underlying landmark geometry.

Recall that an epoch is defined by the pair (P, ν) : the second factor P together with the nonce ν deterministically generates (i) the full set of N challenge points $C = \{c_i\}_{i=1}^N$ and (ii) the landmark shuffle permutation used to order the k landmarks within each response block. From a captured face frame, the protocol produces the full response set $R = \{r_i\}_{i=1}^N$, where each block $r_i \in \{0, 1\}^m$ is a Gray-coded, quantized encoding of the distance vector from challenge point c_i to all shuffled landmarks (i.e., r_i is a compact representation of $(\|c_i - \ell_{\pi(1)}\|, \dots, \|c_i - \ell_{\pi(k)}\|)$ after scaling and Gray coding). Enrollment then binds a random N -bit key by subsetting: blocks at positions corresponding to unset key bits are discarded, and the stored template consists of the remaining subset $f \subseteq R$ together with the nonce ν and a hash of the key.

If an attacker steals records from E epochs, they obtain $(\nu^{(e)}, f^{(e)}, H^{(e)})$ for $e = 1, \dots, E$. The biometric accumulation question depends critically on whether the attacker also learns the second factor P :

Case A: Records Only (no P)

Without P , the attacker cannot regenerate the corresponding challenge coordinates $C^{(e)}$ or the shuffle $\pi^{(e)}$, so the observed blocks $f^{(e)}$ are not readily interpretable. In this case, direct geometric reconstruction from $f^{(e)}$ is substantially impeded because the attacker lacks the coordinate system (the challenges) in which the distances were measured.

Case B: Records Plus P

If the attacker obtains P (and thus can regenerate $C^{(e)}$ and $\pi^{(e)}$ for each stolen epoch), then each retained block $r_i \in f^{(e)}$ becomes a set of quantized distances anchored at a known challenge point $c_i^{(e)}$ with a known landmark ordering. In this setting, the subset operation does not prevent leakage; it only reduces the number of anchored distances revealed per epoch. Over many epochs, the attacker can aggregate distances from many independent

challenge points. Since the stored subset reveals only a fraction of the N blocks each epoch (roughly half for an unbiased key, subject to the ω -constraint), the attacker's per-epoch geometric information is reduced, but across E epochs the total number of anchored blocks grows approximately as $E \cdot |f|$. In principle, with sufficiently many anchored distance observations, the landmark geometry can be increasingly revealed (up to the quantization taken by the distance encoding and the operating tolerance), yielding an accumulation risk that is distinct from session-key compromise.

A simple way to view this risk is in terms of how many anchored challenge blocks leak over time. We let N denote the total number of challenge blocks per epoch, and let $s = |f|$ denote the number of blocks retained in an epoch after subsetting by the key (typically $s \approx N/2$). If an attacker compromises E epochs and (in the worst case) also knows the corresponding second-factor secrets so that the challenge coordinates are known, then the attacker obtains approximately $E \cdot s$ anchored distance blocks in total. If the retained blocks behave approximately like a uniform random subset of size s in each epoch, then repeats occur and the number of distinct anchored blocks observed grows with E , starting near s for small E and approaching N as E becomes large. After observing E epochs, the expected distinct fraction of blocks seen is about

$$1 - \left(1 - \frac{s}{N}\right)^E,$$

so the attacker's geometric information increases steadily with the number of compromised epochs, even though each individual epoch reveals only a fraction of the blocks.

For this reason, we consider multi-epoch compromise a biometric privacy risk distinct from session key compromise. The refreshment policy bounds the window for attacking any single record for key recovery, but it does not substitute for template confidentiality if an attacker can repeatedly steal records over time.

3.5.7. Effect of the Runs Constraint on Key Entropy

The key generation procedure restricts valid N -bit keys to those that never contain more than ω consecutive zeros. In effect, long uninterrupted runs of zeros are disallowed, giving the recovery algorithm a default window size within which one matching block ought to be. This constraint creates a structured subset of all possible N -bit strings, and the size of this subset affects how much entropy the ephemeral key retains.

To understand how many valid keys exist under this rule, we want to count them. We can define $A_N^{(\omega)}$ as the number of N -bit strings that satisfy the constraint (i.e., no substrings of $0^{\omega+1}$). The basic observation is that we can build larger valid strings from smaller ones: any valid n -bit string ending in 0, 00, 000, and so on up to 0^ω can be extended by either appending a 1 or by appending another 0 (as long as we have not already used ω zeros in a row). This leads to the Fibonacci-like recurrence

$$A_N^{(\omega)} = A_{N-1}^{(\omega)} + A_{N-2}^{(\omega)} + \cdots + A_{N-\omega-1}^{(\omega)}.$$

Each term corresponds to choosing how long the final run of zeros is (from 0 zeros up to ω zeros) before a 1 appears. The initial conditions are

$$A_0^{(\omega)} = 1, \quad A_N^{(\omega)} = 2^N \text{ for } 1 \leq N \leq \omega,$$

where $A_0^{(\omega)} = 1$ counts the trivial case, which says that the empty string ($N = 0$) does not violate the runs constraint, and is therefore superficially valid, while $A_N^{(\omega)} = 2^N$ for $1 \leq N \leq \omega$ accounts for the fact that, for string lengths up to ω , the run-length constraint never comes into play because the strings are too short to violate it, so all 2^N binary strings are allowed.

So, the complete definition expressing the number of valid N -bit keys for a given ω is

$$A_N^{(\omega)} = \begin{cases} 1, & N = 0, \\ 2^N, & 1 \leq N \leq \omega, \\ A_{N-1}^{(\omega)} + A_{N-2}^{(\omega)} + \cdots + A_{N-\omega-1}^{(\omega)} & N > \omega. \end{cases}$$

and the total number of valid keys of length N is

$$|\mathcal{K}_N^{(\omega)}| = A_N^{(\omega)}$$

and the effective entropy of this constrained keyspace is therefore

$$H_{\text{const}}(N, \omega) = \log_2 |\mathcal{K}_N^{(\omega)}|.$$

This quantity answers the question, “How many bits of unpredictability do the bound keys have, once the structural rule is taken into account?” [45]. Evaluating this recurrence with a computer, e.g., for $N = 256$ and a fairly restrictive $\omega = 5$, gives

$$H_{\text{const}}(256, 5) = \log_2 |\mathcal{K}_{256}^{(5)}| = 253.0054,$$

indicating that eliminating all 256-bit keys containing runs longer than 0^5 costs about 2 bits of entropy for the valid keys that remain.

In summary, even though the ω -run rule restricts which binary strings are allowed to be used as cryptographic material in this construction, the number of valid keys is still enormous, and key entropy loss due to the constraint is therefore negligible. This suggests that the keys used by the system retain nearly all of their unpredictability while satisfying the structural property needed for the binding mechanism to be most effective.

4. Materials and Methods

We implemented a unified testing system to make sure that all experiments follow a consistent evaluation protocol (Figure 4). The test runner standardizes dataset access through a generalized interface, applies a fixed configuration, and generates genuine and impostor test cases for enrollment and recovery. It records configurations, outcomes, and other metrics for each trial and aggregates the data for external analysis.




Figure 4. The testing system unifies binding mechanisms with biometric datasets, and applies a set of configuration parameters to generate test cases and log performance data.

To evaluate the fundamental capability of the construction for recovering keys from facial geometric enrollments, we ran a multi-parameter configuration sweep over three protocol parameters: the key-recovery matching threshold T , the max-run parameter ω (reported as w), and the uncertainty bound $u_{\max}$ that gates uncertainty correction. This sweep was chosen to reveal a clear performance envelope for recovery accuracy rates and enrollment and recovery latencies, as well as the computational cost impact of bounded uncertainty correction.

We report results on two datasets. First, we use a proprietary Artificial Intelligence (AI)-generated facial dataset designed to minimize inter-frame variation, which provides a controlled environment for characterizing the mechanism in the absence of excessive confounding biometric noise. Second, we evaluate on the Faces94 dataset [46–48], which we selected because it provides comparatively stable frontal face captures and thus aligns most closely with our intended deployment regime of tightly controlled acquisition. While these datasets do not constitute an unconstrained stress test, they support a clearer attribution of effects to specific design choices while still exposing performance limits under real-image capture.

In real-world deployments, higher capture variability due to pose, expression, lighting, and blur—factors that affect both the geometric measurement process and landmarking accuracy—usually necessitates higher matching thresholds T to maintain usability. This, in turn, will raise the average uncertainty $\bar{u}$ for genuine users and may eventually increase False-Accept Rate (FAR) for impostors. The operating points reported in Section 5.1 should therefore be interpreted as optimistic for well-aligned, low-noise capture conditions, with Faces94 providing a more realistic reference for controlled real-image acquisition.

4.1. Configurations/Operating Points

We evaluate the mechanism with sweeps across the three main protocol parameters: tolerance T , the max-run parameter ω , and an uncertainty cutoff $u_{\max}$ that bounds the ternary correction step. For each configuration $(T, \omega, u_{\max})$, we measure key-recovery outcomes over genuine trials (enrolled subject equals recovery subject) and impostor trials (different subjects), treating authentication as a binary decision: success if and only if a binary session key is recovered and its hash matches the stored hash of the true key.

4.1.1. Correction Attempt Gating

Uncertainty correction is attempted only when the internal uncertainty score u satisfies $u \leq u_{\max}$. Otherwise, the protocol fails immediately. Because this gating can affect observed error rates, we report ternary correction attempt rates separately for genuine and impostor trials:

$$\alpha_g = \Pr[u \leq u_{\max} \mid \text{genuine}], \quad \alpha_i = \Pr[u \leq u_{\max} \mid \text{impostor}].$$

We also report the mean uncertainty $\bar{u}$ for genuine and impostor trials at each operating point.

When $u = 0$, recovery reduces to an oracle check of the candidate key from the initial recovery step, and no ternary correction is invoked. More generally, even when $u \leq u_{\max}$, ternary correction is invoked only if the candidate key from the initial recovery step has $u \neq 0$.

4.1.2. Error Rates and Timing Metrics

We define $\text{FAR} = \Pr[\text{accept} \mid \text{impostor}]$, $\text{FRR} = \Pr[\text{reject} \mid \text{genuine}]$, $\text{TAR} = 1 - \text{FRR}$, and $\text{TRR} = 1 - \text{FAR}$. We report mean enrollment time $\bar{\tau}_{\text{enr}}$. We also report mean initial ternary recovery time $\bar{\tau}_{\text{rec}}$ and mean correction time $\bar{\tau}_{\text{corr}}$, where $\bar{\tau}_{\text{corr}}$ is the time spent in the

bounded ternary correction step (zero when correction is not invoked). Total key-recovery latency is therefore $\bar{\tau}_{\text{rec}} + \bar{\tau}_{\text{corr}}$.

5. Results

5.1. Operating-Point Behavior Under the Multi-Parameter Configuration Sweep



We are reporting results using operating-point “slices” that fix $\omega = 5$ with $T \in \{0.15, 0.20, 0.25, 0.30, 0.35, 0.40\}$, and then report outcomes over the uncertainty bound $u_{\text{max}} \in \{0, 16, 28\}$ for both genuine and impostor trials (Table 2, Figure 5 and Table 3, Figure 6). In our evaluation, the second factor is held fixed per trial, so the reported impostor FAR should be interpreted as conditional on second-factor compromise (i.e., impostors are tested under the same second-factor-derived challenges as the target); we do this intentionally to isolate the core face-geometry mechanism and assess discriminability under a consistent challenge set rather than confounding results with “wrong-challenge” failures. This presentation retains the full structure of the two-parameter sweep that matters most empirically, and it reports separately the two latency components—ternary recovery time $\bar{\tau}_{\text{rec}}$ and uncertainty correction time $\bar{\tau}_{\text{corr}}$ —alongside the corresponding error rates and correction attempt rates.

Table 2. NAUAI dataset operating points for fixed empirically best $\omega = 5$. Each panel fixes T and reports metrics across $u_{\text{max}} \in \{0, 16, 28\}$, across $N = 96$ samples per trial type for each configuration.

(a) $T = 0.15$							(b) $T = 0.20$						
Genuine							Genuine						
u_{max}	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	u_{max}	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	1.00	0.00	90.78	0.46	0.02	0.00	0	1.00	0.00	34.64	0.45	0.02	0.00
16	1.00	0.00	102.68	0.45	0.02	0.00	16	0.61	0.39	33.49	0.45	0.02	0.00
28	0.97	0.03	103.30	0.45	0.02	3.91	28	0.43	0.58	33.67	0.46	0.02	5.39
Impostor							Impostor						
u_{max}	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	u_{max}	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	0.00	0.00	255.81	0.46	0.46	0.00	0	0.00	0.00	249.29	0.45	0.29	0.00
16	0.00	0.00	255.86	0.45	0.45	0.00	16	0.00	0.00	250.40	0.45	0.26	0.00
28	0.00	0.00	255.90	0.45	0.46	0.00	28	0.00	0.00	251.20	0.51	0.31	0.00
(c) $T = 0.25$							(d) $T = 0.30$						
Genuine							Genuine						
u_{max}	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	u_{max}	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	0.99	0.01	13.55	0.45	0.02	0.00	0	1.00	0.00	17.76	0.45	0.02	0.00
16	0.26	0.74	14.02	0.45	0.02	0.00	16	0.46	0.55	16.92	0.45	0.02	0.00
28	0.12	0.91	14.41	0.46	0.02	0.33	28	0.12	0.93	17.42	0.47	0.02	5.52
Impostor							Impostor						
u_{max}	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	u_{max}	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	0.00	0.00	214.08	0.45	0.07	0.00	0	0.00	0.00	130.65	0.45	0.03	0.00
16	0.00	0.00	215.09	0.45	0.07	0.00	16	0.01	0.01	113.60	0.45	0.03	0.00
28	0.00	0.00	206.51	0.48	0.07	0.00	28	0.05	0.06	117.79	0.54	0.03	0.35
(e) $T = 0.35$							(f) $T = 0.40$						
Genuine							Genuine						
u_{max}	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	u_{max}	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	1.00	0.00	34.07	0.45	0.02	0.00	0	1.00	0.00	131.61	0.45	0.08	0.00
16	0.99	0.02	35.85	0.45	0.02	0.00	16	1.00	0.00	125.10	0.45	0.08	0.00
28	0.72	0.31	34.03	0.46	0.02	17.92	28	1.00	0.00	128.21	0.45	0.07	0.00
Impostor							Impostor						
u_{max}	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	u_{max}	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	0.00	0.00	63.56	0.45	0.02	0.00	0	0.00	0.00	123.67	0.45	0.05	0.00
16	0.00	0.00	64.34	0.45	0.02	0.00	16	0.00	0.00	136.06	0.45	0.07	0.00
28	0.02	0.03	73.04	0.51	0.02	1.85	28	0.00	0.00	126.03	0.45	0.06	0.00

Table 3. Faces94 dataset operating points for fixed empirically best $\omega = 5$. Each panel fixes T and reports metrics across $u_{\max} \in \{0, 16, 28\}$, across $N = 66$ samples per trial type for each configuration.

(a) $T = 0.15$							(b) $T = 0.20$						
Genuine							Genuine						
$u_{\max}$	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	$u_{\max}$	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	1.00	0.00	197.09	0.08	0.08	0.00	0	1.00	0.00	98.79	0.09	0.03	0.00
16	1.00	0.00	206.00	0.08	0.10	0.00	16	0.94	0.06	99.20	0.09	0.03	0.00
28	1.00	0.00	200.52	0.08	0.09	0.00	28	0.91	0.09	129.20	0.09	0.04	0.38
Impostor							Impostor						
$u_{\max}$	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	$u_{\max}$	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	0.00	0.00	256.00	0.09	0.47	0.00	0	0.00	0.00	253.27	0.09	0.37	0.00
16	0.00	0.00	255.98	0.08	0.47	0.00	16	0.00	0.00	254.48	0.09	0.37	0.00
28	0.00	0.00	256.00	0.08	0.47	0.00	28	0.00	0.00	254.45	0.09	0.40	0.00
(c) $T = 0.25$							(d) $T = 0.30$						
Genuine							Genuine						
$u_{\max}$	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	$u_{\max}$	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	1.00	0.00	33.42	0.11	0.02	0.00	0	1.00	0.00	28.92	0.09	0.02	0.00
16	0.76	0.24	43.91	0.13	0.03	0.00	16	0.70	0.32	32.11	0.09	0.02	0.00
28	0.59	0.41	52.64	0.11	0.03	2.52	28	0.29	0.73	25.86	0.11	0.02	12.96
Impostor							Impostor						
$u_{\max}$	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	$u_{\max}$	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	0.00	0.00	230.35	0.11	0.14	0.00	0	0.00	0.00	179.02	0.09	0.04	0.00
16	0.00	0.00	234.94	0.13	0.17	0.00	16	0.00	0.00	173.79	0.09	0.04	0.00
28	0.00	0.00	233.91	0.11	0.14	0.00	28	0.00	0.00	174.02	0.13	0.05	0.00
(e) $T = 0.35$							(f) $T = 0.40$						
Genuine							Genuine						
$u_{\max}$	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	$u_{\max}$	FRR	α_g	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	1.00	0.00	37.24	0.09	0.02	0.00	0	1.00	0.00	124.68	0.09	0.07	0.00
16	0.95	0.05	39.36	0.09	0.02	0.00	16	1.00	0.00	128.73	0.09	0.08	0.00
28	0.74	0.27	35.95	0.11	0.02	8.59	28	1.00	0.00	127.23	0.09	0.08	0.00
Impostor							Impostor						
$u_{\max}$	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$	$u_{\max}$	FAR	α_i	$\bar{u}$	$\bar{\tau}_{\text{enr}}$	$\bar{\tau}_{\text{rec}}$	$\bar{\tau}_{\text{corr}}$
0	0.00	0.00	112.48	0.09	0.02	0.00	0	0.00	0.00	122.92	0.09	0.04	0.00
16	0.00	0.00	87.15	0.09	0.02	0.00	16	0.00	0.00	127.89	0.09	0.05	0.00
28	0.02	0.02	110.56	0.11	0.03	0.03	28	0.00	0.00	122.38	0.09	0.05	0.00

**(a)** Error rates vs. T ($\omega = 5$)**(b)** α rates vs. T ($\omega = 5$)**Figure 5.** Full data for fixed $\omega = 5$ on NAUAI dataset.




Figure 6. Full data for fixed $\omega = 5$ on Faces94 dataset.

5.1.1. The Role of Matching Tolerance

Across both datasets, very small tolerances yield a “no-recovery” regime: genuine trials exhibit $FRR \approx 1$ and $\alpha_g \approx 0$, accompanied by large mean uncertainty $\bar{u}$. On NAUAI at $T = 0.15$, $FRR = 1.00$ for $u_{\max} \in \{0, 16\}$ with $\alpha_g = 0$ and $\bar{u} \approx 90$ – 103 ; even at $u_{\max} = 28$, the system remains in a failure area with $FRR = 0.97$, $\alpha_g = 0.03$, and $\bar{u} = 103.30$ (Table 2, Figure 5). Faces94 fails even more clearly here, at $T = 0.15$, where $FRR = 1.00$ and $\alpha_g = 0$ for all three $u_{\max}$ values, with $\bar{u} \approx 197$ – 206 (Table 3, Figure 6). It is interesting to note that the level of uncertainty is still significantly higher for impostors than it is for genuine users, suggesting that the mechanism is still reasonably discriminative in principle, but in this particular construction still contains too much error at such tight tolerance to successfully recover keys.

As T increases into the midrange, genuine error rates improve sharply, but only when correction is permitted. On NAUAI at $T = 0.25$, raising $u_{\max}$ from $0 \rightarrow 16 \rightarrow 28$ reduces FRR from $0.99 \rightarrow 0.26 \rightarrow 0.12$ and increases α_g from $0.01 \rightarrow 0.74 \rightarrow 0.91$, while keeping $\bar{u}$ in a low regime (about 13–14) (Table 2, Figure 5). At $T = 0.30$, NAUAI shows a similar pattern: $FRR = 1.00$ at $u_{\max} = 0$ versus $FRR = 0.46$ at $u_{\max} = 16$ and $FRR = 0.12$ at $u_{\max} = 28$ (Table 2, Figure 5). This is evidence that the error correction phase of the construction is a significant help.

Faces94 exhibits the same transition but at generally worse error rates for comparable T . For example, at $T = 0.30$, $FRR = 1.00$ for $u_{\max} = 0$, it improves to $FRR = 0.70$ at $u_{\max} = 16$, and reaches $FRR = 0.29$ at $u_{\max} = 28$ (Table 3, Figure 6). This reinforces the expected sensitivity of facial-geometry recovery to capture variability: a real-people dataset moves the usable envelope toward requiring more permissive correction (larger $u_{\max}$) at a given T .

Finally, at large T (here, $T = 0.40$), both datasets return to a failure regime on the genuine side: $FRR = 1.00$ and $\alpha_g = 0$ for all $u_{\max}$, with $\bar{u} \approx 125$ – 132 (NAUAI) and $\bar{u} \approx 124$ – 129 (Faces94) (Table 2, Figure 5 and Table 3, Figure 6). We observe that, within the explored range, usable performance occurs only in a midrange band of T , and the band is narrower on Faces94 than on the tightly controlled AI dataset.

5.1.2. The Role of Maximum Uncertainty

For fixed T (and here fixed $\omega = 5$), $u_{\max}$ primarily determines how often the protocol is allowed to proceed beyond gating (higher α_g), which in turn largely controls the observed genuine error rate. The NAUAI slice at $T = 0.25$ shows the intended behavior clearly: $u_{\max} = 0$ produces essentially complete failure ($FRR = 0.99$, $\alpha_g = 0.01$), while $u_{\max} = 16$ yields usable behavior ($FRR = 0.26$, $\alpha_g = 0.74$) and $u_{\max} = 28$ improves further ($FRR = 0.12$, $\alpha_g = 0.91$) (Table 2, Figure 5). Similar improvements appear at $T = 0.20$ and $T = 0.30$.

Raising $u_{\max}$, however, clearly shifts work into uncertainty correction and substantially increases $\bar{\tau}_{\text{corr}}$. On NAUAI, this cost is visible but generally modest in the best midrange slices: at $T = 0.25$, $u_{\max} = 28$ increases correction cost to $\bar{\tau}_{\text{corr}} = 0.33$ relative to $\bar{\tau}_{\text{corr}} = 0.00$ at $u_{\max} = 16$ (Table 2, Figure 5). In contrast, on Faces94, the correction-time penalty at $u_{\max} = 28$ strongly dominates: at $T = 0.30$, going from $u_{\max} = 16$ to $u_{\max} = 28$ improves FRR from 0.70 to 0.29 and increases α_g from 0.32 to 0.73, but raises $\bar{\tau}_{\text{corr}}$ from 0.00 to 12.96 (Table 3, Figure 6). We acknowledge that 12 s is a long time to wait for a successful recovery. At $T = 0.35$, the same trade-off persists (FRR = 0.95 $\rightarrow$ 0.74 and $\alpha_g = 0.05 \rightarrow 0.27$, while $\bar{\tau}_{\text{corr}}$ jumps to 8.59 at $u_{\max} = 28$) (Table 3, Figure 6). Empirically, $u_{\max}$ therefore acts as the main usability–latency trade-off knob, which is not surprising given the exponential nature of the search problem.

5.1.3. Impostor Behavior and Nonzero FAR

False accepts do occur for some operating points. On NAUAI, impostor error remains essentially absent through $T \leq 0.25$ for all $u_{\max}$ (all listed FAR = 0.00), but becomes nonzero at higher T when correction is more permissive: at $T = 0.30$, FAR = 0.01 for $u_{\max} = 16$ and FAR = 0.05 for $u_{\max} = 28$; at $T = 0.35$, FAR = 0.02 for $u_{\max} = 28$ (Table 2, Figure 5). In these same cells, the impostor attempt rate is also nonzero ($\alpha_i = 0.01$ at $(T, u_{\max}) = (0.30, 16)$ and $\alpha_i = 0.06$ at $(0.30, 28)$), indicating that the gating condition is met often enough for impostors that recovery can sometimes succeed (Table 2, Figure 5). This is an unfortunate weakness of the uncertainty gating mechanism: larger T increases inter-subject overlap and reduces uncertainty, making it easier for impostors to stay under the uncertainty bound and increasing the chance of a successful fraudulent recovery.

On Faces94, the impostor slices remain almost entirely at FAR = 0.00 across this T range, with the lone exception at $T = 0.35$, $u_{\max} = 28$, where FAR = 0.02 and $\alpha_i = 0.02$ (Table 3, Figure 6). We observe that, within these slices, Faces94's primary limitation is usability (high FRR unless correction is permissive), whereas NAUAI's limitation at higher T becomes security (nonzero FAR) when $u_{\max}$ is large enough to admit impostor attempts.

5.1.4. Timing and Latency Decomposition

Across both datasets, ternary recovery time $\bar{\tau}_{\text{rec}}$ remains small and comparatively stable in the midrange where genuine outcomes are nontrivial (typically ≈ 0.02 on NAUAI and ≈ 0.02 – 0.03 on Faces94), while uncertainty correction time $\bar{\tau}_{\text{corr}}$ is the dominant contributor when correction is frequently invoked (Table 2, Figure 5 and Table 3, Figure 6). This separation is clearest on Faces94: at $T = 0.30$, $u_{\max} = 28$, $\bar{\tau}_{\text{rec}} = 0.02$ but $\bar{\tau}_{\text{corr}} = 12.96$; at $T = 0.35$, $u_{\max} = 28$, $\bar{\tau}_{\text{rec}} = 0.02$ but $\bar{\tau}_{\text{corr}} = 8.59$ (Table 3, Figure 6). On NAUAI, correction time is usually near-zero except in the same cells where α_g is high and correction is permissive (e.g., $\bar{\tau}_{\text{corr}} = 5.39$ at $T = 0.20$, $u_{\max} = 28$ and $\bar{\tau}_{\text{corr}} = 5.52$ at $T = 0.30$, $u_{\max} = 28$) (Table 2, Figure 5). This suggests unsurprisingly that ternary recovery is “cheap,” uncertainty correction is the bottleneck, and the bottleneck activates mainly when $u_{\max}$ is set to admit borderline samples.

Enrollment time differs by dataset and is relatively stable within each dataset's slices: on NAUAI, $\bar{\tau}_{\text{enr}} \approx 0.45$ across most listed operating points; on Faces94, it is smaller (roughly 0.08–0.13) (Table 2, Figure 5 and Table 3, Figure 6). This effect is probably attributable to implementation-level differences such as image dimensions and preprocessing cost rather than protocol parameters (a result of misplaced timing measurements), since the data collection sweeps affect recovery/correction much more than enrollment.

5.1.5. Practical Operating Regions Within the $\omega = 5$ Slices

The slice tables suggest a concise picture of the usable region and its trade-offs. On NAUAI, the most attractive region for usability and security lies around $T = 0.25$ with

$u_{\max} \in \{16, 28\}$. At $T = 0.25$, $u_{\max} = 16$ achieves $\text{FRR} = 0.26$ with negligible correction cost ($\bar{\tau}_{\text{corr}} = 0.00$) and $\text{FAR} = 0.00$; $u_{\max} = 28$ improves usability ($\text{FRR} = 0.12$) at a modest correction penalty ($\bar{\tau}_{\text{corr}} = 0.33$) while still maintaining $\text{FAR} = 0.00$ in this slice (Table 2, Figure 5). Moving to $T = 0.30$, it continues to perform well with genuine FRR at $u_{\max} = 28$ (0.12) but introduces small impostor error ($\text{FAR} = 0.05$) and nonzero impostor correction attempt rate ($\alpha_i = 0.06$), making $T = 0.30$ materially less attractive from a security perspective in this dataset (Table 2, Figure 5).

On Faces94, the only slice exhibiting meaningfully usable genuine error rates is $T = 0.30$ with $u_{\max} = 28$, but this setting carries substantial correction cost ($\bar{\tau}_{\text{corr}} = 12.96$) (Table 3, Figure 6). Reducing $u_{\max}$ to 16 eliminates the correction-time overhead but leaves the system in a high-error regime ($\text{FRR} = 0.70$) (Table 3, Figure 6). It appears that, on a much less tightly-controlled real-face dataset, the mechanism can be driven to usable error rates only by permitting correction on a larger fraction of trials, and the practical constraint becomes correction latency rather than impostor error (which remains near-zero in these areas).

Overall, the results are fairly clear to summarize: T and $u_{\max}$ jointly determine whether the mechanism operates in a (i) no-recovery regime (high $\bar{u}$, $\alpha_g \approx 0$), (ii) usable regime (moderate $\bar{u}$, large α_g , low FRR), or (iii) permissive regime where impostor attempts can pass gating and FAR becomes non-negligible. In general, empirical results on both datasets suggested the theoretical potential for such a mechanism to work, but exhibited relatively low overall usability, especially as capture conditions moved outside of extremely tight controls.

6. Discussion

The system proposed herein was designed as an illustration of how a PUF-style challenge–response mechanism can be mapped onto a biometric setting. It demonstrates that one can take random challenges, tie them to a biometric structure such as facial landmarks, and then bind a key by discarding responses, all while keeping revocability through a second factor and an automatically refreshed nonce. This gives us an interesting platform for thinking about key binding, revocation, and noise correction in biometrics, even if this particular construction by itself is somewhat fragile.

It is important to recognize that the security of the scheme as defined here really sits on the shoulders of the auxiliary password. If an attacker steals the enrollment record, what they receive is the reduced subset of responses, the hash of the bound key, and indirectly the Hamming weight of that key. The entropy of the key itself remains largely intact: for a sufficiently long key, knowing the Hamming weight and/or the ω parameter does not meaningfully help an attacker mount a feasible brute-force search (see Section 3.5.7 [33,34]). The subset of responses is harder to interpret because each response is a concatenation of Gray-coded distances from a randomly chosen challenge point to the set of landmarks. Without the password and nonce, the challenge points cannot be regenerated and the landmark order remains unknown, so reconstructing the facial geometry from the stored template is not straightforward. In this sense, the subset alone is not an obvious privacy leak.

The situation changes if the password is compromised. With the second factor in hand, an attacker can regenerate the challenge points and recover the landmark permutation, which makes the stored responses directly interpretable. At that point, reconstructing the underlying face geometry becomes trivial. This highlights an imbalance in the design: there is revocability, but it comes exclusively from leaning on the password as the revocation lever. The biometric portion of the scheme adds complexity but not much security on its own.

There are other limitations worth noting. The noise in landmark detection and capture conditions makes the reconstruction problem brittle, which is why bounded local search is necessary in the first place. That bounded search grows more expensive as noise increases, and at some point the cost becomes intolerable. The nonce refresh prevents long-lived template reuse, but it does not eliminate cross-matching risk entirely, since the repeated compromise of different session records for the same user could eventually leak correlations. Finally, the system has not been analyzed against modern attack models that are well-documented in the cancelable biometrics literature, and it is not therefore immediately clear how it would hold up under adversarial scrutiny.

Beyond limitations, there are several contexts where a geometric and CRP-based system is a feature rather than a liability. In tightly controlled capture regimes (e.g., badge-and-gate building access, lab cleanrooms, kiosk check-ins, or operator consoles) pose, distance, and lighting can realistically be constrained, at which point the deterministic geometry becomes a stable signal, the tolerance T can be tuned once and held fixed, and the nonce-driven refresh gives you per-session revocability without model retraining or cross-site template reuse. Unlike a learned embedding that generalizes “who you are” across many appearances, our response pipeline is deliberately narrow: it performs best when the environment is predictable, which is exactly what many access-control deployments can guarantee.

A second, more interesting advantage is that geometry makes it easy to incorporate user-driven micro-gestures (most simply, a facial expression) as a private factor that is never written down and cannot be brute-forced in the classical sense. In this scenario, the enrollment can capture a neutral pose plus a user-chosen expression (or subtle muscle pattern), and the same expression is required at authentication, so the matcher’s responses are tied not only to who the user is but to what the user intentionally does in that moment, which is closer to a password than to a passive biometric. A conventional modern neural feature extractor will usually recognize “you” whether you are smiling or not, but a geometric CRP mechanism can be parameterized to make that expression matter, effectively turning an otherwise static trait into a reproducible, revocable behavior that raises the bar for lookalike spoofing.

There are also pragmatic advantages: the whole pipeline is interpretable end-to-end (challenge points, landmark shuffle, distances, bit-level masks), which makes failure triage easier; it has a small computational footprint with no large models to store, no training data to curate; it can run comfortably on-device, it avoids retaining raw images, and the nonce guarantees that a successful recovery yields a short-lived key that is replaced automatically at session close, allowing per-door or per-session keys and straightforward revocation. Taken together, these traits sketch a viable niche for the construction: controlled capture, minimal infrastructure, interpretable security posture, and a simple, human-in-the-loop “gesture factor” that other templates typically do not exploit.

Our empirical FAR/FRR results should be interpreted in the context of the specific evaluation protocol and datasets we used. First, the reported operating points are conditional on the evaluated datasets and their capture conditions; we make no claim that the same error rates will hold under substantially different capture conditions. Second, the reported error rates are specific to the particular genuine/impostor pairings sampled from the datasets; different pairing or sampling strategies could certainly yield different empirical FAR/FRR. Third, our threat model is limited to the impostor setting evaluated in this study; we did not experimentally assess the ability of stronger adversaries to invert the templates or fraudulently recover the ephemeral keys (e.g., presentation attacks, deepfakes, adaptive attacks, or attackers with additional knowledge).

In summary, the contribution here is broadly conceptual rather than immediately practical, but it carves out a clean design point that is useful to consider. The construction shows that CRP-style interfaces can be adapted to biometrics, that revocability can still be achieved by mixing in a second factor and discarding parts of the response, that keys can be refreshed on every session via the nonce, and that key recovery is possible even under noisy conditions. At the same time, the design as specified carries clear weaknesses: privacy and security rest almost entirely on the second factor, error tolerance is limited, and revocation is only as strong as the secrecy of the second factor. Still, the pipeline is functional, lightweight, and easy to deploy in controlled capture environments, and it naturally accommodates user-driven factors (e.g., expressions) that conventional learned embeddings tend to ignore. These trade-offs should temper expectations of robustness, but they do not diminish the value of the idea as a practical niche solution in constrained settings and as a seed for further exploration.

7. Future Work

While the CRP-style feature extraction mechanism is an interesting notion, its performance as a purely geometric feature is inevitably limited in discriminative power. The feature representation used here is highly sensitive to capture conditions and provides only very modest intra-subject consistency and inter-subject separability, which constrains our achievable FAR/FRR even when the binding mechanism behaves very well and as-intended. A natural extension of this work is therefore to replace the proposed geometric CRP feature extraction mechanism with more discriminative and capture-robust feature extractors (e.g., learned embeddings). It would also be informative and valuable to study the same binding protocol with alternative biometric modalities (e.g., fingerprint, iris, vascular patterns). Such future analyses would better isolate the security and usability properties of the binding mechanism from the limitations of a less-than-ideal feature extractor, and would likely lead to a more accurate and practical technology.

8. Conclusions

We presented a conceptual exploration of a two-factor, per-session refreshable cancellable biometric key-binding scheme that borrows a PUF-style challenge–response interface and applies it to facial landmark geometry, binding and recovering a key according to recently proposed novel cryptographic constructions. The scheme works, in principle, as follows: given a second factor, a session nonce, and a cooperating user, it is possible to bind a key to the biometric, discard the key, and later recover the key reliably enough to decrypt protected data from a matching biometric from a genuine subject. Our experiments confirm that the mechanism can function reasonably well under well-controlled capture conditions. The addition of a nonce ensures that every successful session produces a fresh enrollment and a new ephemeral key, so the system is continually refreshed rather than relying on a long-lived template. At the same time, the trade-off between usability and security remains somewhat severe. To make the system tolerant enough to handle natural biometric noise, one must accept error margins wide enough that false acceptances can become uncomfortably likely. The contribution here should therefore be understood as a conceptual demonstration that PUF-style challenge–response mechanisms, combined with the subset key binding and recovery mechanisms previously proposed in the literature, can in principle be ported into biometrics, with the hope that the idea itself may inspire stronger constructions that do not suffer from the same security and practicality constraints. Overall, the paper’s contribution is the articulation and empirical characterization of this design niche, together with a clear accounting of the security constraints and where the approach degrades as capture conditions loosen.

9. Patents

This article describes research related to subject matter disclosed in the following U.S. patent applications. These disclosures are acknowledged as relevant prior art and to clarify the relationship between this manuscript and the disclosed inventions; nothing herein should be interpreted as claiming or assigning rights to any intellectual property.

1. **U.S. Patent Application Publication No. US 2026/0005876 A1** (published 1 January 2026), *Pseudo-Homomorphic Authentication of Users with Biometry* (filed 5 September 2025) [49].
2. **U.S. Patent Application Publication No. US 2025/0167997 A1** (published 22 May 2025), *Protocols with Noisy Response-Based Cryptographic Subkeys* (filed 13 September 2024) [50].

Author Contributions: Conceptualization, M.L.G., M.A. and M.P.; methodology, M.L.G., M.A. and M.P.; software, M.L.G.; validation, M.L.G., M.A. and J.H.; formal analysis, M.L.G.; investigation, M.L.G., M.A. and J.H.; resources, J.H.; data curation, M.L.G.; writing—original draft preparation, M.L.G.; writing—review and editing, M.L.G., M.A. and J.H.; visualization, M.L.G.; supervision, J.H.; project administration, J.H.; funding acquisition, J.H. All authors have read and agreed to the published version of the manuscript.

Funding: Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-23-2-0014. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: The AI dataset used for biometric protocol testing was sourced from Generated Photos (<https://generated.photos/>, accessed on 15 February 2026). The authors used generative AI tools (ChatGPT 5.2, 2025) to assist in the articulation of ideas and clarity of expression; all conceptual content, analysis, and final revisions were developed and verified by the authors. The authors gratefully acknowledge Bertrand Cambou, Ian Burke, Saloni Jain, Jack Garrard, and Amisha Bagri for their support in developing, implementing, and analyzing the construction proposed herein.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

PUF	Physical Unclonable Function
IBM	International Business Machines Corporation
CRP	Challenge–Response Pair
API	Application Programming Interface
SRAM	Static Random-Access Memory
KDF	Key Derivation Function
TAR	True-Accept Rate
FAR	False-Accept Rate
TRR	True-Reject Rate
FRR	False-Reject Rate
PRF	Pseudorandom Function

References

1. Ratha, N.K.; Connell, J.H.; Bolle, R.M. Enhancing security and privacy in biometrics-based authentication systems. *IBM Syst. J.* **2001**, *40*, 614–634. [\[CrossRef\]](#)
2. Bolle, R.M.; Connell, J.H.; Ratha, N.K. Biometrics Perils and Patches. *Pattern Recognit.* **2002**, *35*, 2727–2738. [\[CrossRef\]](#)
3. Savvides, M.; Vijaya Kumar, B.; Khosla, P. Cancelable biometric filters for face recognition. In Proceedings of the 17th International Conference on Pattern Recognition, ICPR 2004, Cambridge, UK, 26 August 2004; Volume 3, pp. 922–925. [\[CrossRef\]](#)
4. Boulton, T.E. Robust Distance Measures for Face Recognition Supporting Revocable Biometric Tokens. In Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition (FG), Southampton, UK, 10–12 April 2006; pp. 560–566.
5. Boulton, T.E.; Scheirer, W.; Woodworth, R. Revocable Fingerprint Biotokens: Accuracy and Security Analysis. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, Minneapolis, MN, USA, 17–22 June 2007; pp. 1–7.
6. Farooq, F.; Bolle, R.M.; Jea, T.Y.; Ratha, N. Anonymous and Revocable Fingerprint Recognition. In Proceedings of the 2007 IEEE Conference on Computer Vision and Pattern Recognition, Minneapolis, MN, USA, 17–22 June 2007; pp. 1–7. [\[CrossRef\]](#)
7. Jin, A.T.B.; Ling, D.N.C.; Goh, A. Biohashing: Two factor authentication featuring fingerprint data and tokenised random number. *Pattern Recognit.* **2004**, *37*, 2245–2255. [\[CrossRef\]](#)
8. Teoh, A.B.J.; Yuang, C.T. Cancelable Biometrics Realization with Multispace Random Projections. *IEEE Trans. Syst. Man Cybern. Part B (Cybern.)* **2007**, *37*, 1096–1106. [\[CrossRef\]](#)
9. Lumini, A.; Nanni, L. An improved BioHashing for human authentication. *Pattern Recognit.* **2007**, *40*, 1057–1065. [\[CrossRef\]](#)
10. Kong, A.; Cheung, K.H.; Zhang, D.; Kamel, M.; You, J. An analysis of BioHashing and its variants. *Pattern Recognit.* **2006**, *39*, 1359–1368. [\[CrossRef\]](#)
11. Lalithamani, N.; Soman, K. An Efficient Approach for Non-Invertible Cryptographic Key Generation from Cancelable Fingerprint Biometrics. In Proceedings of the 2009 International Conference on Advances in Recent Technologies in Communication and Computing, Kottayam, India, 27–28 October 2009; pp. 47–52. [\[CrossRef\]](#)
12. Abid, M.; Kanade, S.; Petrovska-Delacr'etaz, D.; Dorizzi, B.; Afifi, H. Iris Based Authentication Mechanism for E-Passports. In Proceedings of the 2nd International Workshop on Security and Communication Networks (IWSCN), Karlstad, Sweden, 26–28 May 2010; pp. 1–6.
13. Nagar, A.; Nandakumar, K.; Jain, A. Biometric template transformation: A security analysis. In *Proceedings of the Media Forensics and Security II, San Jose, CA, USA, 17–21 January 2010*; Memon, N.D., Dittmann, J., Alattar, A.M., Delp, E.J., III, Eds.; Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series; SPIE: Bellingham, WA, USA, 2010; Volume 7541, p. 75410O. [\[CrossRef\]](#)
14. Rathgeb, C.; Uhl, A. A Survey on Biometric Cryptosystems and Cancelable Biometrics. *EURASIP J. Inf. Secur.* **2011**, *2011*, 3. [\[CrossRef\]](#)
15. Manisha; Kumar, N. Cancelable Biometrics: A comprehensive survey. *Artif. Intell. Rev.* **2020**, *53*, 3403–3446. [\[CrossRef\]](#)
16. Li, C.; Hu, J. Attacks via record multiplicity on cancelable biometrics templates. *Concurr. Comput. Pract. Exp.* **2014**, *26*, 1593–1605. [\[CrossRef\]](#)
17. Maaref, Z.; Belhadj, F.; Attia, A.; Akhtar, Z.; Jasser, M.B.; Ramly, A.M.; Mohamed, A.W. A comprehensive review of vulnerabilities and attack strategies in cancelable biometric systems. *Egypt. Inform. J.* **2024**, *27*, 100511. [\[CrossRef\]](#)
18. Dong, X.; Park, J.; Jin, Z.; Teoh, A.B.J.; Tistarelli, M.; Wong, K. On the security risk of pre-image attack on cancelable biometrics. *J. King Saud Univ.—Comput. Inf. Sci.* **2024**, *36*, 102060. [\[CrossRef\]](#)
19. Hirata, S.; Takahashi, K. Cancelable Biometrics with Perfect Secrecy for Correlation-Based Matching. In *Proceedings of the Advances in Biometrics, Alghero, Italy, 2–5 June 2009*; Tistarelli, M., Nixon, M.S., Eds.; Springer: Berlin/Heidelberg, Germany, 2009; pp. 868–878.
20. Takahashi, K.; Hirata, S. Cancelable Biometrics with Provable Security and Its Application to Fingerprint Verification. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.* **2011**, *94-A*, 233–244. [\[CrossRef\]](#)
21. Nazari, S.; Moin, M.S.; Kanan, H.R. Cancelable face using Chaos permutation. In Proceedings of the 7th International Symposium on Telecommunications (IST'2014), Tehran, Iran, 9–11 September 2014; pp. 925–928. [\[CrossRef\]](#)
22. Alam, M.T.; Shamshad, F.; Karray, F.; Nandakumar, K. FaceAnonymizer: Cancelable Faces via Identity-Consistent Latent Space Mixing. *arXiv* **2025**, arXiv:2308.05636.
23. Singh, O.; Jaiswal, A.; Kumar, N. Real-numbered singular matrix transformation for non-invertible and cancelable biometric templates. *Appl. Intell.* **2025**, *55*, 719. [\[CrossRef\]](#)
24. Dey, K.; Kule, M.; Rahaman, H. PUF Based Hardware Security: A Review. In Proceedings of the 2021 International Symposium on Devices, Circuits and Systems (ISDCS), Higashihiroshima, Japan, 3–5 March 2021; pp. 1–6. [\[CrossRef\]](#)
25. Teoh, A.; Goh, A.; Ngo, D. Random Multispace Quantization as an Analytic Mechanism for BioHashing of Biometric and Random Identity Inputs. *IEEE Trans. Pattern Anal. Mach. Intell.* **2006**, *28*, 1892–1901. [\[CrossRef\]](#)

26. Ratha, N.K.; Chikkerur, S.; Connell, J.H.; Bolle, R.M. Generating Cancelable Fingerprint Templates. *IEEE Trans. Pattern Anal. Mach. Intell.* **2007**, *29*, 561–572. [\[CrossRef\]](#)
27. Pillai, J.K.; Patel, V.M.; Chellappa, R.; Ratha, N.K. Sectored random projections for cancelable iris biometrics. In Proceedings of the 2010 IEEE International Conference on Acoustics, Speech and Signal Processing, Dallas, TX, USA, 14–19 March 2010; pp. 1838–1841.
28. Pillai, J.K.; Patel, V.M.; Chellappa, R.; Ratha, N.K. Secure and Robust Iris Recognition Using Random Projections and Sparse Representations. *IEEE Trans. Pattern Anal. Mach. Intell.* **2011**, *33*, 1877–1893. [\[CrossRef\]](#)
29. Rathgeb, C.; Breiting, F.; Busch, C. Alignment-free cancelable iris biometric templates based on adaptive bloom filters. In Proceedings of the 2013 International Conference on Biometrics (ICB), Madrid, Spain, 4–7 June 2013; pp. 1–8. [\[CrossRef\]](#)
30. Rathgeb, C.; Breiting, F.; Busch, C.; Baier, H. On application of bloom filters to iris biometrics. *IET Biom.* **2014**, *3*, 207–218. [\[CrossRef\]](#)
31. Rathgeb, C.; Busch, C. Cancelable multi-biometrics: Mixing iris-codes based on adaptive bloom filters. *Comput. Secur.* **2014**, *42*, 1–12. [\[CrossRef\]](#)
32. Lin, Y.; Liao, Y.; Wei, Y.; Zeng, W.; Erkan, U.; Toktas, A.; Zhang, Y.; Chen, D. Lightweight image encryption via four-dimensional Hénon memristor map and fast block permutation. *Nonlinear Sci. Control Eng.* **2025**, *1*, 025390012. [\[CrossRef\]](#)
33. Cambou, B.; Philabaum, C.; Hoffstein, J.; Herlihy, M. Methods to Encrypt and Authenticate Digital Files in Distributed Networks and Zero-Trust Environments. *Axioms* **2023**, *12*, 531. [\[CrossRef\]](#)
34. Ghanai Miandoab, D.; Garrett, M.L.; Alam, M.; Jain, S.; Assiri, S.; Cambou, B. Secure Cryptographic Key Encapsulation and Recovery Scheme in Noisy Network Conditions. *Appl. Sci.* **2025**, *15*, 2732. [\[CrossRef\]](#)
35. Jin, Z.; Teoh, A.B.J.; Goi, B.M.; Tay, Y.H. Biometric cryptosystems: A new biometric key binding and its implementation for fingerprint minutiae-based representation. *Pattern Recognit.* **2016**, *56*, 50–62. [\[CrossRef\]](#)
36. Sharma, S.; Saini, A.; Chaudhury, S. A survey on biometric cryptosystems and their applications. *Comput. Secur.* **2023**, *134*, 103458. [\[CrossRef\]](#)
37. ISO/IEC 24745:2022; Information Security, Cybersecurity and Privacy Protection—Biometric Information Protection. International Organization for Standardization: Geneva, Switzerland, 2022.
38. Chai, T.Y.; Goi, B.M.; Tay, Y.H.; Jin, Z. A New Design for Alignment-Free Chaffed Cancelable Iris Key Binding Scheme. *Symmetry* **2019**, *11*, 164. [\[CrossRef\]](#)
39. Sandhya, M.; Prasad, M. Cancelable Fingerprint Cryptosystem Using Multiple Spiral Curves and Fuzzy Commitment Scheme. *Int. J. Pattern Recognit. Artif. Intell.* **2017**, *31*, 1756004. [\[CrossRef\]](#)
40. Juels, A.; Wattenberg, M. A fuzzy commitment scheme. In Proceedings of the 6th ACM Conference on Computer and Communications Security, CCS'99, Singapore, 1–4 November 1999; ACM: New York, NY, USA, 1999; pp. 28–36. [\[CrossRef\]](#)
41. Juels, A.; Sudan, M. A Fuzzy Vault Scheme. In Proceedings of the ISIT'02, Lausanne, Switzerland, 30 June–5 July 2002.
42. Uludag, U.; Pankanti, S.; Jain, A.K. Fuzzy vault for fingerprints. In Proceedings of the 5th International Conference on Audio- and Video-Based Biometric Person Authentication, AVBPA'05, Hilton Rye Town, NY, USA, 20–22 July 2005; Springer: Berlin/Heidelberg, Germany, 2005; pp. 310–319. [\[CrossRef\]](#)
43. Dodis, Y.; Ostrovsky, R.; Reyzin, L.; Smith, A. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. *SIAM J. Comput.* **2008**, *38*, 97–139. [\[CrossRef\]](#)
44. Kazemi, V.; Sullivan, J. One millisecond face alignment with an ensemble of regression trees. In Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 1867–1874. [\[CrossRef\]](#)
45. Shannon, C.E. A mathematical theory of communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. [\[CrossRef\]](#)
46. So-In, C.; Rujirakul, K. wPFP-PCA: Weighted Parallel Fixed Point PCA Face Recognition. *Int. Arab J. Inf. Technol.* **2016**, *13*, 59–69.
47. Rujirakul, K.; So-In, C. P-PCC: Parallel Pearson Correlation Condition for Robust Cosmetic Makeup Face Recognitions. In *Information Science and Applications; Lecture Notes in Electrical Engineering*; Springer: Berlin/Heidelberg, Germany, 2015; Volume 339, pp. 259–266.
48. Rujirakul, K.; So-In, C.; Arnonkijpanich, B. PEM-PCA: A Parallel Expectation-Maximization PCA Face Recognition Architecture. *Sci. World J.* **2014**, *2014*, 468176. [\[CrossRef\]](#)
49. Cambou, B.F.; Garrett, M.L. Pseudo-Homomorphic Authentication of Users with Biometry. U.S. Patent 20260005876, 1 January 2026.
50. Cambou, B.F.; Garrett, M.L.; Partridge, M.; Ghanaimiandoab, D. Protocols with Noisy Response-Based Cryptographic Subkeys. U.S. Patent 20250167997, 22 May 2025.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.