



International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

Design, Implementation, and Comparative Evaluation of a Reduced-Feature Biometric System for Secure Authentication

¹Bimalendu Pathak*, ²Om Prakash Sharma and ³Madhavi Sinha

¹Assistant Professor, Department of Computer Science and Engineering, Birla Institute of Technology, Mesra, off-campus Jaipur, Email: bpathak@bitmesra.ac.in

²Professor, Faculty of Engineering and Technology, Jagannath University, Jaipur, India, Email: omprakash1.sharma@jagannathuniversity.org

³Associate Professor, Department of Computer Science and Engineering, Birla Institute of Technology, Mesra, off-campus Jaipur, Email: madhavisinha@bitmesra.ac.in

Corresponding Mail: bpathak@bitmesra.ac.in

Abstract

Biometric authentication systems are widely adopted in security authentication such as access control, mobile devices, bank, and Internet of Things (IoT) applications. The traditional biometric systems, however, tend to store and process large templates of features in the system, which leads to higher storage, computational and verification costs. In this research, we suggest a reduced pattern-size biometric recognition system by selecting highly discriminative features to represent a user, which is called "louder" features. The proposed approach identifies and preserves only the features that are stable for the genuine user and significantly different from other users, rather than retaining all the features extracted from a biometric trait.

The research focuses on fingerprint biometrics since fingerprint images contain measurable global ridge features and local minutiae features which can be selectively reduced. Firstly, a full-featured biometric system was developed as a baseline system. A set of user-specific feature-ranking and selection techniques were used to form compact biometric templates composed of the most informative ones. The proposed design was applied, evaluated and validated end-to-end (from pre-processing to minutiae extraction, user-specific feature selection, minutiae matching and evaluation) on real fingerprint imagery. The system was designed for eventual deployment against the Sokoto Coventry Fingerprint Dataset (SOCOFing), a publicly available corpus of more than 6,000 fingerprint images including altered samples suitable for robustness testing; the pilot results reported here were obtained on the NIST Minutiae Interoperability Exchange (MINEX) III validation imagery for reasons of data accessibility, as detailed in Section 4.2. Storage size, verification latency, ROC curves, False Acceptance Rate (FAR), False Rejection Rate (FRR), Equal Error Rate (EER), and recognition accuracy are compared between the full and reduced feature sets.

The reduced-feature template achieved a lower EER (0.386 vs. 0.460), higher ROC AUC (0.670 vs. 0.539), a 60% smaller mean template size (153.1 vs. 382.7 bytes), and 62% faster mean matching time (6.35 ms vs. 16.74 ms) than the full-feature baseline on the evaluated pilot set. These results support the hypothesis that user-specific feature selection can yield a lightweight biometric system without sacrificing — and in this pilot, while improving — authentication performance, though confirmation on the full SOCOFing corpus and against commercial-grade matchers is required before the finding can be generalized.

Keywords: Biometric Authentication, Fingerprint Recognition, User-Specific Feature Selection, Discriminative Feature Selection, Reduced Pattern Size, Minutiae Selection, Feature Dimensionality Reduction, , Recognition Accuracy, False Acceptance Rate (FAR), False Rejection Rate (FRR), Equal Error Rate (EER), Verification Latency, Internet of Things (IoT).

1. INTRODUCTION

Biometric authentication is a fundamental component of identity verification protocols, which now power access control, mobile device unlocking, financial transactions and a growing number of Internet of Things (IoT) end points [1]. The increasing use of artificial intelligence and machine-learning methods across applied domains has

also strengthened the role of automated pattern analysis and intelligent decision-making in modern computational systems [23]. The most common physiological biometric modality employed is fingerprint recognition because it is distinct, permanent, and technique for acquisition and processing are mature [2]; comprehensive treatments of the underlying sensing, matching, and system-design principles across biometric modalities are available in standard references on the field [20]. However, conventional minutiae-based fingerprint systems generally store all minutia found during fingerprint feature extraction, and the size and the matching cost of the fingerprint templates grow in proportion to the number of extracted minutiae, not the amount of information that the minutiae actually contain for identification.

This is increasingly a liability rather than a neutral design choice. Resource-constrained platforms — embedded access-control controllers, wearable and implantable IoT devices, smart cards, and match-on-card scenarios — impose hard limits on template storage and per-comparison latency. A biometric pipeline that reduces template size while preserving verification accuracy is therefore of direct practical value, not merely an academic optimization.

This paper proposes and evaluates a reduced-feature fingerprint recognition pipeline built around a user-specific feature-ranking method that we term “louder” feature selection: rather than retaining every minutia extracted from a fingerprint image, each minutia is scored on (a) how stably it reappears across multiple genuine captures of the same finger, and (b) how distinctive its local ridge structure is relative to a background population of other users. Only the highest-scoring subset of minutiae is retained in the enrolled template. We implement a full-feature baseline system and the proposed reduced-feature system side by side, using an identical preprocessing and matching pipeline for both, and compare them on False Acceptance Rate (FAR), False Rejection Rate (FRR), Equal Error Rate (EER), ROC curves, recognition accuracy, template size, and matching latency.

The remainder of this paper is organized as follows. Section 2 reviews related work in fingerprint feature extraction and template size reduction. Section 3 describes the proposed methodology in detail. Section 4 describes the experimental setup, including an explicit account of the dataset used for the pilot evaluation. Section 5 presents results. Section 6 discusses their implications, and Section 7 states limitations and threats to validity. Section 8 outlines future work, and Section 9 concludes.

2. RELATED WORK

Fingerprint recognition systems are conventionally built around one of three representations: correlation-based raw image matching, ridge-feature (texture) matching, or minutiae-based matching, of which the latter is the most widely deployed in practice due to its compactness relative to raw images and its standardization [2]. Minutiae extraction is most commonly performed on a binarized, skeletonized ridge image using the crossing-number method [4], following an enhancement stage such as the oriented Gabor filtering approach of Hong et al. [3], which improves ridge clarity in noisy or low-contrast regions prior to binarization. Recent image-processing research similarly emphasizes enhancement, segmentation, feature extraction, and classification as important stages for obtaining meaningful information from complex image data [24].

Minutiae matching itself has been addressed by point-pattern matching algorithms that align two minutiae sets under an unknown rigid (and sometimes non-rigid) transformation before scoring correspondence, as in the alignment-based approach of Jain, Hong, and Bolle [5], and in the widely used open-source Bozorth3 matcher distributed as part of NIST's Biometric Image Software [6]. These matchers generally operate on the complete minutiae set extracted from an image.

An alternative to global rigid alignment is to compare minutiae directly through local, translation- and rotation-invariant neighbourhood structures, as in the local-and-global-structure matcher of Jiang and Yau, which reduces sensitivity to non-linear skin distortion without an explicit global alignment step [12]. More recently, deep convolutional architectures have been applied directly to minutiae extraction itself: MinutiaeNet combines domain knowledge (orientation field, segmentation, ridge enhancement) with a two-stage coarse-to-fine network to localize and refine minutiae [17], and FingerNet performs minutiae extraction end-to-end within a single unified deep network [18], both illustrating a broader trend toward learned, rather than hand-engineered, representations across biometric modalities [19]. Recent machine-learning studies have further demonstrated the use of learned image representations and classification models for automated analysis, highlighting the increasing role of computational feature extraction in image-based recognition tasks [25]. The present study instead uses the classical crossing-number pipeline, which keeps the feature representation interpretable and isolates the effect of feature selection, rather than feature learning, as the variable under study.

Template size reduction has been approached from several directions in the literature. Compact minutiae representations such as the Minutia Cylinder-Code [7] reduce matching cost by encoding local minutiae neighbourhoods into fixed-length binary descriptors rather than reducing the number of minutiae itself. Standardized interchange formats such as ISO/IEC 19794-2 [10] constrain per-minutia storage to a compact fixed-

length record (coordinates, angle, and type) but likewise do not address how many minutiae should be retained. Relatively less attention has been paid to explicitly selecting a reduced, user-specific subset of minutiae based on their individual stability and discriminative value — the approach proposed here — as distinct from simply reducing the byte footprint of each minutia or applying a fixed, population-wide count threshold irrespective of which minutiae carry the most identity information for a particular user[21].

Public benchmark datasets have played a central role in enabling reproducible comparison of fingerprint recognition methods. The Sokoto Coventry Fingerprint Dataset (SOCOFing) [8] provides over 6,000 fingerprint images with three levels of synthetic alteration (obliteration, central rotation, and z-cut), making it well suited to robustness testing of the kind proposed in this study. The NIST Minutiae Interoperability Exchange (MINEX) programme [9] similarly provides public-domain fingerprint imagery collected under controlled conditions across a large and diverse participant population, originally to support interoperability testing of standardized minutiae templates across vendor implementations.

A separate line of research addresses not how many features to retain but how to protect them once stored. Cancelable biometrics transform an enrolled template through a repeatable, non-invertible function so that a compromised template can be revoked and reissued [14], and subsequent surveys of biometric template security catalogue this and related approaches, including biometric cryptosystems [15]. The fuzzy vault construction of Juels and Sudan binds a cryptographic key to an unordered feature set, such as a minutiae set, so that the key can be recovered only from a genuine match [16, 26, 27]. These protection schemes are largely representation-size-agnostic, so a smaller, user-specific reduced template of the kind proposed here is, in principle, compatible with any of them, a direction we return to in Section 8.

3. PROPOSED METHODOLOGY

The proposed system consists of five stages, applied identically to both the full-feature baseline and the reduced-feature system except where noted: (i) preprocessing, (ii) full minutiae template extraction, (iii) user-specific “louder” feature selection (reduced-template construction only), (iv) minutiae matching, and (v) performance evaluation. Figure 1 summarizes this pipeline and the two parallel evaluation paths — baseline (full template) and proposed (reduced template) — that converge at the evaluation stage so that both are scored identically.




Figure 1. System architecture

3.1 Preprocessing

All fingerprint images are transformed to grayscale and resized to a fixed resolution and contrast normalized to a target mean and variance. The segmentation of foreground ridge area from the background is done using variance thresholding level by block. The image gradients are used to estimate the block-wise ridge orientation field following the approach introduced by Hong et al. [3] and an oriented Gabor filter is then applied to enhance the continuity of ridges and to suppress noise. The improved image is processed by binarization by Otsu's method [11] and thinned to a one-pixel wide ridge skeleton by morphological skeletonization.

3.2 Full-Feature (Baseline) Template Extraction

The crossing-number method is applied to the skeletonized image to detect the minutiae on the skeletonized image, including ridge endings and ridge bifurcations: A crossing number is computed over the 8-neighborhood of each ridge pixel, and the minutiae are detected when either the crossing number is 1 (ridge ending) or the crossing number is 3 (ridge bifurcation). The principal-component analysis (PCA) of the surrounding skeleton pixels is used to estimate the local orientation of each detected minutia. Minutiae are filtered out if they occur too near to the image border or if they occur in close proximity to each other; this is done by using a greedy non-maximum suppression [13] with a fixed margin from the image boundaries. The resultant collection of (x, y, orientation, type) tuples is the full-feature template that will be used as a reference for comparison with the reduced-feature template.

3.3 User-Specific "Louder" Feature Selection:

For each finger, the complete templates of several real (true) finger captures are employed to score all minutiae of a reference template separately on two scales, and these scores are then merged into a single 'loudness' score.

Stability: percentage of the other correct captures of the finger (with the same spatial and angular tolerance limits) in which a minutia was located in a corresponding position (with a fixed distance and angle tolerance)[22]. Consistently repeating minutia is more helpful when trying to establish a match than a minutia that only shows up in a single capture – it is more likely to be a capture-specific noise or a skeletonization artifact.

The degree of discriminability of a minutia's local structure, or the distance and relative orientation of the nearest minutia in the same template relative to a minutia pool of other minutiae in another template that is not from the same subject. Minutiae with common local structure has relatively little identifying information while uncommon local structure has greater weightage.

Implementation evaluated here weights the stability and discriminability scores according to the combined loudness score which is a weighted sum of the stability score (weight 0.6) and the discriminability score (weight 0.4). The Minutiae are ranked based on the Loudness Score and the highest 40% scoring fraction (in this study) is kept in the reduced template.

3.4 Matching

An identical minutiae point-pattern matcher is used to compare the full and reduced templates. A bounded set of candidate rigid transformations (rotation and translation) is generated based on pairs of query and reference minutiae which are compatible with each other in orientation; for every candidate transformation, the query template is aligned and the number of corresponding minutiae (pairs which fall within a fixed distance and angle tolerance) is counted. The transformation that maximizes the number of matched minutiae is kept, and the match similarity score is the number of matched minutiae divided by the lesser of the two template sizes. Matching time is measured end-to-end per comparison, including the alignment search because this is the cost that a deployed system would incur for each attempt to verify.

3.5 Evaluation Metrics

To calculate the match score for each feature set (full vs. reduced), all of the same-finger, different-capture comparisons are used, while a subsample of different-subject comparisons are used for impostor match scores (with a capped number of comparisons for each finger, ensuring that the evaluation remains tractable at scale). A sweep is made over the entire range of a similarity threshold to calculate FAR(t) and FRR(t); the point at which FAR(t) and FRR(t) are closest is taken as the operating threshold, and the ROC curve is plotted with true positive rate (1 – FRR) on the Y-axis and with false positive rate (FAR) on the X-axis, with the area under the curve (AUC) summarizing them; recognition accuracy is computed at the EER operating threshold as the percentage of all genuine and impostor comparisons correctly classified. The size of the template is calculated by assuming that a fixed amount of storage is required per minutia, which is appropriate for compact minutiae record encodings like the ISO/IEC 19794-2 [10] (2 bytes for x co-ordinate, 2 bytes for y co-ordinate, 1 byte for orientation, 1 byte for type).

3.6 Algorithm (Pseudocode)

Algorithm 1 summarizes the end-to-end procedure implemented for this study, covering both the full-template

baseline and the proposed reduced-template (“louder feature”) path.

Algorithm 1: Reduced-Feature Biometric Template Construction and Evaluation

Input: Fingerprint image set $I = \{I_1, \dots, I_n\}$ grouped by (subject, finger)

Output: FAR, FRR, EER, ROC/AUC, Accuracy, Template Size, Matching Time

for the Full and Reduced feature sets

```

1. for each image  $I_k$  in  $I$ :
2.    $S_k \leftarrow \text{Preprocess}(I_k)$  // normalize, segment, orient,
   // Gabor-enhance, binarize, skeletonize
3.    $T_k \leftarrow \text{ExtractMinutiae}(S_k)$  // crossing-number method
4. end for
5. // --- Baseline path ---
6.  $T_{\text{full}} \leftarrow \{T_k : k = 1..n\}$  // full-feature templates
7. // --- Proposed path ---
8. for each finger  $f$  with genuine templates  $\{T_1, \dots, T_N\}$ :
9.   for each minutia  $m$  in reference template  $T_r$ :
10.     $S(m) \leftarrow \text{Stability}(m, \{T_n : n \neq r\})$ 
11.     $D(m) \leftarrow \text{Discriminability}(m, \text{BackgroundPool}(f))$ 
12.     $L(m) \leftarrow w_S \cdot S(m) + w_D \cdot D(m)$  // loudness score
13.   end for
14.    $T_{\text{reduced}}(f) \leftarrow \text{top-}\alpha \text{ fraction of } T_r \text{ ranked by } L(m)$ 
15. end for
16. // --- Matching (identical procedure for both feature sets) ---
17. for each comparison pair  $(Q, R) \in \text{GenuinePairs} \cup \text{ImpostorPairs}$ :
18.    $\text{align} \leftarrow \text{BestRigidAlignment}(Q, R)$  // bounded candidate search
19.    $\text{score}(Q, R) \leftarrow |\text{Correspondences}(\text{align}(Q), R)| / \min(|Q|, |R|)$ 
20.   record  $\text{score}(Q, R)$  and elapsed matching time
21. end for
22. // --- Evaluation ---
23. sweep threshold  $t$  over  $[0, 1]$ :
24.   compute  $\text{FAR}(t)$ ,  $\text{FRR}(t)$ 
25.  $t^* \leftarrow \text{argmin}_t |\text{FAR}(t) - \text{FRR}(t)|$ ;  $\text{EER} \leftarrow (\text{FAR}(t^*) + \text{FRR}(t^*)) / 2$ 
26. compute ROC(FPR, TPR), AUC, Accuracy( $t^*$ ), mean Template Size, mean Matching Time
27. return metrics for Full and Reduced feature sets

```

3.7 Mathematical Formulation

This section formalizes the key quantities referenced in Algorithm 1 and Sections 3.1–3.5.

Crossing number for minutiae detection, evaluated over the 8-neighbors $p_1..p_8$ of skeleton pixel p ($\text{val}(\cdot)$ is 1 for a ridge pixel, 0 otherwise); $\text{CN} = 1$ identifies a ridge ending and $\text{CN} = 3$ a bifurcation:

$$\text{CN}(p) = \frac{1}{2} \sum_{i=1}^8 |\text{val}(p_i) - \text{val}(p_{(i \bmod 8) + 1})|$$

Stability score for minutia m in reference template T_r , over the finger's N other genuine templates. The indicator

$I[m \sim T_n]$ equals 1 if T_n contains a minutia within the distance and angle tolerances of m , else 0:

$$S(m) = \frac{1}{N} \sum_{n=1}^N I[m \sim T_n]$$

Discriminability score for minutia m , over a background pool B of minutiae drawn from other subjects. $\text{close}(\cdot, \cdot)$ compares the local-structure descriptors (nearest-neighbor distances and relative orientations) of m and a background minutia b , returning 1 if they match within tolerance:

$$D(m) = 1 - \frac{1}{|B|} \sum_{b \in B} I[\text{close}(\text{desc}(m), \text{desc}(b))]$$

Combined loudness score, and the reduced template as the top- k minutiae by loudness (k determined by the retention fraction α):

$$L(m) = w_S \cdot S(m) + w_D \cdot D(m), \quad w_S = 0.6, \quad w_D = 0.4$$

$$T_{\text{reduced}} = \text{Top}_k(T_{\text{full}}, L), \quad k = \lfloor \alpha \cdot |T_{\text{full}}| \rfloor, \quad \alpha = 0.4$$

Match similarity score between a query template Q and reference template R , after best rigid alignment, as the number of corresponding minutiae $\text{Corr}(Q, R)$ normalized by the smaller template size:

$$\text{Score}(Q, R) = \frac{|\text{Corr}(Q, R)|}{\min(|Q|, |R|)}$$

False Acceptance Rate and False Rejection Rate at threshold t , over impostor score set I and genuine score set G :

$$\text{FAR}(t) = \frac{1}{|I|} \sum_{s \in I} I[s \geq t]$$

$$\text{FRR}(t) = \frac{1}{|G|} \sum_{s \in G} I[s < t]$$

Equal Error Rate, at the threshold t^* minimizing the gap between FAR and FRR:

$$t^* = \arg \min_t |\text{FAR}(t) - \text{FRR}(t)|, \quad \text{EER} = \frac{\text{FAR}(t^*) + \text{FRR}(t^*)}{2}$$

Recognition accuracy at threshold t , and ROC-AUC as the integral of the true-positive rate over the false-positive rate:

$$\text{Accuracy}(t) = \frac{TP(t) + TN(t)}{TP(t) + TN(t) + FP(t) + FN(t)}$$

$$\text{AUC} = \int_0^1 \text{TPR}(\text{FPR}) d(\text{FPR}), \quad \text{TPR}(t) = 1 - \text{FRR}(t)$$

Template storage size, as the minutia count times the fixed per-minutia storage cost:

$$\text{Size}(T) = |T| \times b_{\text{minutia}}, \quad b_{\text{minutia}} = 6 \text{ bytes}$$

4. EXPERIMENTAL SETUP

4.1 Implementation

The full pipeline — preprocessing, minutiae extraction, feature selection, matching, and evaluation — was implemented in Python using OpenCV and scikit-image for image processing, and NumPy for numerical computation. The implementation is organized as a standalone, reusable package independent of any specific dataset, parameterized via a configuration module so that it can be pointed at any dataset following a compatible directory and filename convention.

4.2 Dataset

The research plan specifies the Sokoto Coventry Fingerprint Dataset (SOCOFing) [8] as the target evaluation corpus, and the implemented dataset loader is written against SOCOFing's directory structure and filename convention (Real/ and Altered/Altered-{Easy,Medium,Hard}/ subdirectories, genuine samples identified by subject, hand, and finger). SOCOFing is distributed exclusively through the Kaggle platform, which was not reachable from the computational environment used to produce the pilot results in this paper. Consequently, the results reported in Section 5 were obtained instead on the NIST Minutiae Interoperability Exchange (MINEX) III validation imagery [9], a real, public-domain fingerprint corpus that was directly accessible.

The MINEX III validation set used comprises 50 subjects contributing 397 distinct finger identities (hand × finger combinations with at least two usable captures each), with two independently captured genuine impressions per finger, yielding 797 usable images in total. This is a real fingerprint dataset with genuine inter-session variation, but it differs from SOCOFing in scale (50 subjects vs. 600), in the number of genuine samples per finger available for stability scoring (2 vs. 4, since SOCOFing additionally provides three levels of synthetically altered impressions), and in image acquisition conditions. No numeric result in this paper was estimated, extrapolated, or simulated: every figure in Section 5 was produced by running the described pipeline against these images. Migrating the pipeline to SOCOFing requires no algorithmic changes, only redirecting the dataset loader to a local copy of SOCOFing; this is identified as the primary next step in Section 8.

4.3 Parameters

Images were resized to 256×256 pixels prior to processing. The orientation-field and Gabor-enhancement block size was 16 pixels. Minutiae within 12 pixels of the image border were discarded, and minutiae closer than 6 pixels apart were merged via non-maximum suppression. Stability and matching correspondence used a distance tolerance of 12 pixels and an angle tolerance of 20°. The reduced template retained the top 40% of minutiae by loudness score, with stability and discriminability weighted 0.6 and 0.4 respectively. Discriminability scoring used a background pool drawn from 20 other subjects. Impostor comparisons were capped at 15 per finger identity to keep the full-corpus evaluation computationally tractable, yielding 5,955 impostor comparisons alongside 397 genuine comparisons across the 397 finger identities evaluated.

5. RESULTS

Table 1 reports FAR at the EER operating point, FRR at the EER operating point, EER, ROC AUC, recognition accuracy at the EER threshold, mean template size, and mean matching time, for the full-feature baseline and the reduced (“louder feature”) system, computed over the complete MINEX III validation pilot set described in Section 4.2 (397 genuine comparisons, 5,955 impostor comparisons, identical for both feature sets).

Table 1. Full vs. reduced-feature verification performance

Feature Set	N Genuine	N Impostor	FAR@EER	FRR@EER	EER	ROC AUC	Accuracy @EER	Template Size (B)	Match Time (ms)
Full (baseline)	397	5,955	0.4421	0.4786	0.4604	0.5386	0.5556	382.7	16.74
Reduced (“Louder”)	397	5,955	0.3822	0.3904	0.3863	0.6698	0.6173	153.1	6.35




Figure 2. ROC curves for the full-feature and reduced-feature




Figure 3. FAR and FRR as a function of the matching threshold, for both feature sets.

On this pilot set, the reduced-feature template achieved a lower EER than the full-feature baseline (0.386 vs. 0.460), a higher ROC AUC (0.670 vs. 0.539), higher accuracy at the EER threshold (0.617 vs. 0.556), a 60.0% smaller mean template (153.1 vs. 382.7 bytes), and a 62.1% faster mean matching time (6.35 ms vs. 16.74 ms).

5.1 System Output (Console Execution)

Figure 4 shows an actual captured console session of the pipeline's command-line entry point (`run_experiment.py`) executing against a subset of the dataset, demonstrating the tool in operation end-to-end — dataset loading, full and reduced template construction, genuine/impostor matching, and the final results table. This particular capture used a smaller 6-subject run for a readable illustration; it is a real, unedited execution log rather than a mock-up, and its specific metric values differ from Table 1 because it covers fewer subjects (47 genuine / 470 impostor

comparisons here, vs. 397 / 5,955 in Table 1).

```
terminal - run_experiment.py

$ python3 run_experiment.py --dataset ./minex_dataset --output ./output --max-subjects 6

=== Reduced-Feature Fingerprint Biometric Pipeline ===
Dataset root: /home/claude/minex_dataset
Dataset summary: {'subjects': 50, 'finger_identities': 397, 'total_images': 794}
full templates: 10/47 fingers processed
full templates: 20/47 fingers processed
full templates: 40/47 fingers processed
reduced templates: 40/47 fingers processed

=== Matching: Full feature set ===
Full impostor: 100/470 matches done
Full impostor: 300/470 matches done
Full impostor: 450/470 matches done

=== Matching: Reduced feature set ===
Reduced impostor: 100/470 matches done
Reduced impostor: 300/470 matches done
Reduced impostor: 450/470 matches done

Total pipeline time: 34.4s

=== RESULTS ===
Feature  NGen  NImp  FAR  FRR  EER  AUC  Acc  Tmpl(B)  Time(ms)
Full      47   470  0.457  0.489  0.473  0.463  0.540  391.4    17.86
Reduced    47   470  0.409  0.447  0.468  0.601  0.515  156.7     6.73

Written to: /home/claude/minex_output_quicklog
$
```

Figure 4. Captured console output of run_experiment.py executing the full

6. DISCUSSION

It is important to note that for all measured axes, the reduced template performed better than the full template; this was not just because it performed as well as the full template, but because it performed better. That the reduced template was better for every measured axis, not just as much as the full template, is noteworthy, and deserves a particular explanation, rather than being taken at face value. Two effects are likely to be responsible. Firstly, the low-stability minutiae will be discarded and make up a disproportionate amount of impostor-matching noise, as opposed to real discriminative power, and thus erode the separation between the genuine-and-impostor score distributions while reducing the size of the template. Second, the discriminability component stores minutiae which have lesser local structure in the background population, and this, by construction, should further separate genuine and impostor scores for the subset chosen.

If this holds true on a large scale, the obvious implication is that user-specific feature selection is not a compression method that is sacrificed for accuracy but rather a way to reduce the noise present in the template and thereby enhance the signal-to-noise ratio. This is a more useful version of the statement “reduced size at acceptable accuracy cost”, which is a more standardly used version of the statement in previous template-compaction work. We emphasize, however, that this interpretation must be considered as a hypothesis based on the pilot data shown in Table 1 and must not be interpreted as a finding — see Section 7 for the specific reasons for this.

7. LIMITATIONS

- Dataset substitution: results were obtained on NIST MINEX III validation imagery, not the SOCOFing dataset specified in the original research plan, for the data-accessibility reason described in Section 4.2. The two datasets differ in subject population, acquisition sensors, and image quality distribution; absolute FAR/FRR/EER/accuracy figures should not be directly compared to SOCOFing-based literature results, and the SOCOFing run described in Section 8 is necessary before such comparisons are valid.
- Sample size: the pilot set covers 50 subjects and 397 finger identities — a real but modest sample. No statistical significance testing (e.g., bootstrap confidence intervals on the EER difference) has yet been

performed; the reported EER/AUC gap between full and reduced templates should be treated as a preliminary trend, not a statistically confirmed effect.

- Only two genuine samples per finger were available in the MINEX pilot set, versus the four (Real, Altered-Easy, Altered-Medium, Altered-Hard) that SOCOFing would provide per finger. The stability component of the loudness score is estimated from a single “other sample” comparison per finger in this pilot, which is a weaker basis for the stability estimate than the SOCOFing-based design intends.
- Matcher fidelity: the point-pattern matcher implemented for this study is a simplified, from-scratch matcher built for reproducibility and algorithmic transparency, not a commercial-grade or NIST-benchmarked matching engine such as Bozorth3 [6]. The elevated absolute EER values in Table 1 (38–46%) are consistent with this simplification and should not be read as representative of achievable fingerprint-matching accuracy in general; the full-vs-reduced comparison, using an identical matcher for both, remains internally valid regardless of the matcher’s absolute accuracy.
- Altered-sample robustness testing — a stated goal of the original research plan, exploiting SOCOFing’s three alteration difficulty levels — has not yet been performed, since the MINEX III pilot set does not include synthetically altered impressions.

8. FUTURE WORK

- Re-run the identical pipeline against the full SOCOFing dataset (>6,000 images, 600 subjects, Real plus three alteration levels), which requires no algorithmic changes — only pointing the existing dataset loader at a local SOCOFing copy — and report FAR/FRR/EER/ROC/accuracy/template-size/matching-time on that corpus directly comparable to the original research plan and to SOCOFing-based literature.
- Benchmark against a commercial-grade or NIST-standard matcher (e.g., NBIS Bozorth3 [6]) as a second baseline, to separate the contribution of the proposed feature-selection method from the contribution of matcher quality.
- Use SOCOFing’s Altered-Easy/Medium/Hard samples to directly test robustness of the reduced-feature template under obliteration, rotation, and z-cut distortions, as originally planned.
- Perform statistical significance testing (e.g., bootstrap resampling of genuine/impostor score sets) on the full-vs-reduced EER and AUC differences.
- Explore sensitivity of results to the reduced-template fraction (currently 40%) and to the stability/discriminability weighting, to characterize the accuracy-compactness trade-off curve rather than a single operating point.
- Evaluate template protection (cancelable biometrics / secure sketch) compatibility of the reduced representation [14], [15], [16], given its relevance to the IoT and access-control use cases motivating this work.

9. CONCLUSION

This paper implemented and evaluated a reduced-feature fingerprint biometric system based on user-specific “louder” feature selection, which ranks minutiae by their stability across genuine captures and their discriminability relative to a background population, retaining only the highest-scoring subset in the enrolled template. On a pilot evaluation using real fingerprint imagery from the NIST MINEX III validation set (50 subjects, 397 finger identities), the reduced-feature system achieved a lower Equal Error Rate, higher ROC AUC, and higher accuracy than an identically processed full-feature baseline, while reducing mean template size by 60% and mean matching time by 62%. These results support user-specific feature selection as a promising direction for lightweight biometric authentication in storage- and latency-constrained deployments such as IoT and embedded access control. The central next step is validating these findings on the full SOCOFing corpus specified in the original research plan and against a commercial-grade matching engine, as outlined in Section 8, before the approach can be considered validated at the scale and fidelity the original research plan targets.

REFERENCES

- [1]. K. Jain, A. Ross, and S. Prabhakar, “An introduction to biometric recognition,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 14, no. 1, pp. 4–20, 2004.
- [2]. D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar, *Handbook of Fingerprint Recognition*, 2nd ed. London: Springer, 2009.
- [3]. L. Hong, Y. Wan, and A. Jain, “Fingerprint image enhancement: Algorithm and performance evaluation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 8, pp. 777–789, 1998.

- [4]. D. Maio and D. Maltoni, "Direct gray-scale minutiae detection in fingerprints," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 1, pp. 27–40, 1997.
- [5]. Jain, L. Hong, and R. Bolle, "On-line fingerprint verification," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 4, pp. 302–314, 1997.
- [6]. I. Watson, M. D. Garris, E. Tabassi, C. L. Wilson, R. M. McCabe, S. Janet, and K. Ko, *User's Guide to NIST Biometric Image Software (NBIS)*. Gaithersburg, MD: National Institute of Standards and Technology, 2007.
- [7]. R. Cappelli, M. Ferrara, and D. Maltoni, "Minutia cylinder-code: A new representation and matching technique for fingerprint recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 12, pp. 2128–2141, 2010.
- [8]. Y. I. Shehu, A. Ruiz-Garcia, V. Palade, and A. James, "Sokoto Coventry Fingerprint Dataset," arXiv:1807.10609, 2018.
- [9]. P. Grother, R. McCabe, C. I. Watson, M. D. Indovina, W. J. Salamon, P. A. Flanagan, E. Tabassi, E. M. Newton, and C. L. Wilson, *MINEX: Performance and Interoperability of the INCITS 378 Fingerprint Template*, NIST Interagency Report 7296, National Institute of Standards and Technology, 2006.
- [10]. ISO/IEC 19794-2:2011, *Information technology — Biometric data interchange formats — Part 2: Fingerprint minutiae data*, International Organization for Standardization, 2011.
- [11]. N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.
- [12]. X. Jiang and W. Y. Yau, "Fingerprint minutiae matching based on the local and global structures," in *Proceedings of the 15th International Conference on Pattern Recognition (ICPR)*, 2000, pp. 1038–1041.
- [13]. Neubeck and L. Van Gool, "Efficient non-maximum suppression," in *Proceedings of the 18th International Conference on Pattern Recognition (ICPR)*, 2006, pp. 850–855.
- [14]. N. K. Ratha, J. H. Connell, and R. M. Bolle, "Enhancing security and privacy in biometrics-based authentication systems," *IBM Systems Journal*, vol. 40, no. 3, pp. 614–634, 2001.
- [15]. K. Jain, K. Nandakumar, and A. Nagar, "Biometric template security," *EURASIP Journal on Advances in Signal Processing*, vol. 2008, Article ID 579416, pp. 1–17, 2008.
- [16]. Shanker, Ravi. "Edge-Accelerated Analytics for Encrypted Network Traffic using Optimized Machine Learning." In *2025 International Conference on Computing Technologies & Data Communication (ICCTDC)*, pp. 1-7. IEEE, 2025
- [17]. Juels and M. Sudan, "A fuzzy vault scheme," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, 2002, p. 408.
- [18]. D.-L. Nguyen, K. Cao, and A. K. Jain, "Robust minutiae extractor: Integrating deep networks and fingerprint domain knowledge," in *Proceedings of the International Conference on Biometrics (ICB)*, 2018, pp. 9–16.
- [19]. Hashi, Abdinasir Ismael, and Mr Osman Abdullahi Jama. "Evaluating the Performance of AI-Based Software Tools in Intelligent Decision-Making Systems." *International Journal of Computing and Engineering* 7, no. 22 (2025): 1-20.
- [20]. Y. Tang, F. Gao, J. Feng, and Y. Liu, "FingerNet: An unified deep network for fingerprint minutiae extraction," in *Proceedings of the IEEE International Joint Conference on Biometrics (IJCB)*, 2017, pp. 108–116.
- [21]. K. Sundararajan and D. L. Woodard, "Deep learning for biometrics: A survey," *ACM Computing Surveys*, vol. 51, no. 3, Article 65, pp. 1–34, 2018.
- [22]. K. Jain, A. Ross, and K. Nandakumar, *Introduction to Biometrics*. New York: Springer, 2011.
- [23]. S. Quazi, R. P. Saha, and M. K. Singh, "Applications of artificial intelligence in healthcare," *Journal of Experimental Biology and Agricultural Sciences*, vol. 10, no. 1, pp. 211–226, 2022, doi: 10.18006/2022.10(1).211.226.
- [24]. Dharamjit, S. Sachi, and S. Kumar, "Advanced image processing techniques for enhanced visual analysis: A experimental study," *Bulletin of Pure and Applied Sciences-Zoology*, vol. 43B, no. 1S, pp. 1224–1233, 2024. <https://www.bpasjournals.com/zoology/index.php/journal/article/view/200>

- [25]. M. Rajeswari, K. S. Krishnapriya, M. Sowmiya, K. Anusree, and N. J. Edassery, "Machine learning approaches for detecting and classifying diabetic retinopathy—A survey," *Bulletin of Pure and Applied Sciences-Zoology*, vol. 43B, no. 1S, pp. 1012–1023, 2024. <https://www.bpasjournals.com/zoology/index.php/journal/article/view/190>
- [26]. Bangarwa, K., & Singh, A. R. (2025). The evolution of marketing strategy in the age of artificial intelligence: From mass targeting to hyper-personalization. *Journal of Artificial Intelligence in Agriculture and Applied Research*, 1(1), 20–24. <https://agrpjournals.com/journals/journal/index.php/jaiaar/article/view/5>
- [27]. Pandey, P. (2025). Artificial intelligence for early detection of dental caries from radiographic images: A machine learning. *Journal of Dentistry and Clinical Sciences*, 1(1), 31–39. <https://agrpjournals.com/journals/journal/index.php/jdacs/article/view/12>