UNIVERSITY OF PATRAS
DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING

DIVISION OF ELECTRONICS AND COMPUTER
INTERACTIVE TECHNOLOGIES LABORATORY

Design, Implementation & Evaluation of Cancellable EEG Templates using Biohashing Methods

D I P L O M A T H E S I S

ANDREAS GALANIS

SUPERVISOR: Dr. CHRISTOS A. FEIDAS

PATRAS – February, 2026

University of Patras, Department of Electrical and Computer Engineering.

Andreas Galanis

© 2026 – All rights reserved

The whole work is an original work, produced by Andreas Galanis, and does not violate the rights of third parties in any way. If the work contains material which has not been produced by him/her, this is clearly visible and is explicitly mentioned in the text of the work as a product of a third party, noting in a similarly clear way his identification data, while at the same time confirming that in case of using original graphics representations, images, graphs, etc., has obtained the unrestricted permission of the copyright holder for the inclusion and subsequent publication of this material.

CERTIFICATION

It is certified that the Diploma Thesis titled

Design, Implementation & Evaluation of Cancellable EEG Templates using Biohashing Methods

of the Department of Electrical and Computer Engineering student

ANDREAS GALANIS

Registration Number: 1083635

was presented publicly at the Department of Electrical and Computer Engineering at

09/02/2026

and was examined by the following examining committee:

Feidas Christos, Associate Professor, ECE (supervisor)

Daskalaki Sofia, Assistant Professor, ECE (committee member)

Sintoris Christos, Teaching and Research Associate, ECE (committee member)

The Supervisor	The Director of the Division
Feidas Christos Associate Professor	Theodoridis George Assistant Professor

ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ
ΤΜΗΜΑ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ ΤΕΧΝΟΛΟΓΙΑΣ ΥΠΟΛΟΓΙΣΤΩΝ

ΤΟΜΕΑΣ ΗΛΕΚΤΡΟΝΙΚΗΣ ΚΑΙ ΥΠΟΛΟΓΙΣΤΩΝ (Η.&Υ.)
ΕΡΓΑΣΤΗΡΙΟ ΕΠΙΒΛΕΠΟΝΤΟΣ

Σχεδιασμός, Υλοποίηση και Αξιολόγηση Ακυρώσιμων Προτύπων ΗΕΓ με τη χρήση μεθόδων Βιοκατακερματισμού.

Δ Ι Π Λ Ω Μ Α Τ Ι Κ Η Ε Ρ Γ Α Σ Ι Α

ΑΝΔΡΕΑΣ ΓΑΛΑΝΗΣ

ΕΠΙΒΛΕΠΩΝ: Δρ. ΧΡΗΣΤΟΣ Α. ΦΕΙΔΑΣ

ΠΑΤΡΑ – ΦΕΒΡΟΥΑΡΙΟΣ 2026

Πανεπιστήμιο Πατρών, Τμήμα Ηλεκτρολόγων Μηχανικών και Τεχνολογίας Υπολογιστών.

Ανδρέας Γαλάνης

© 2026 – Με την επιφύλαξη παντός δικαιώματος

Το σύνολο της εργασίας αποτελεί πρωτότυπο έργο, παραχθέν από τον Ανδρέα Γαλάνη, και δεν παραβιάζει δικαιώματα τρίτων καθ' οιονδήποτε τρόπο. Αν η εργασία περιέχει υλικό, το οποίο δεν έχει παραχθεί από τον ίδιο, αυτό είναι ευδιάκριτο και αναφέρεται ρητώς εντός του κειμένου της εργασίας ως προϊόν εργασίας τρίτου, σημειώνοντας με παρομοίως σαφή τρόπο τα στοιχεία ταυτοποίησής του, ενώ παράλληλα βεβαιώνει πως στην περίπτωση χρήσης αυτούσιων γραφικών αναπαραστάσεων, εικόνων, γραφημάτων κ.λπ., έχει λάβει τη χωρίς περιορισμούς άδεια του κατόχου των πνευματικών δικαιωμάτων για την συμπερίληψη και επακόλουθη δημοσίευση του υλικού αυτού.

ΠΙΣΤΟΠΟΙΗΣΗ

Πιστοποιείται ότι η Διπλωματική Εργασία με τίτλο

Σχεδιασμός, Υλοποίηση και Αξιολόγηση Ακυρώσιμων Προτύπων ΗΕΓ με τη χρήση μεθόδων Βιοκατακερματισμού.

του φοιτητή του Τμήματος Ηλεκτρολόγων Μηχανικών και Τεχνολογίας Υπολογιστών

ΑΝΔΡΕΑ ΓΑΛΑΝΗ ΤΟΥ **ΣΩΤΗΡΙΟΥ**

Αριθμός Μητρώου: 1083635

Παρουσιάστηκε δημόσια στο Τμήμα Ηλεκτρολόγων Μηχανικών και Τεχνολογίας Υπολογιστών στις

09/02/2026

και εξετάστηκε από την ακόλουθη εξεταστική επιτροπή:

Φείδας Χρήστος, Αναπληρωτής Καθηγητής, Τμήμα ΗΜ&ΤΥ (επιβλέπων)

Δασκαλάκη Σοφία, Επίκουρη Καθηγήτρια, Τμήμα ΗΜ&ΤΥ (μέλος επιτροπής)

Σιντόρης Χρήστος, Ε.ΔΙ.Π, Τμήμα ΗΜ&ΤΥ (μέλος επιτροπής)

Ο Επιβλέπων	Ο Διευθυντής του Τομέα
Φείδας Χρήστος Αναπληρωτής Καθηγητής	Θεοδωρίδης Γεώργιος Αναπληρωτής Καθηγητής

Acknowledgments

First, I would like to express my sincere gratitude to Dr. Feidas. From the moment he joined the Department of Electrical and Computer Engineering, his professionalism and passion for teaching became immediately evident. Years later, when I was searching for a topic for my diploma thesis, our discussions helped me find a subject that genuinely sparked my interest in a way I had never experienced before during my studies.

I also want to deeply thank Mr. Ioannis Katoikos for his consistent support throughout the completion of this thesis. His insightful ideas and expertise made our collaboration smooth and productive.

Additionally, I am truly grateful to my parents for their unwavering support of my goals and ambitions over the years. They have stood by me through every step, both in good and in challenging times.

Finally, I want to thank my friends. Their constant encouragement pushed me to grow and become a better person.

This diploma thesis marks my graduation from the Department of Electrical and Computer Engineering, and I sincerely hope it is the beginning of an exciting journey ahead.

ABSTRACT

Design, implementation & Evaluation of cancellable eeg templates using biohashing methods

STUDENT NAME, SURNAME: **SUPERVISOR NAME, SURNAME:**

ANDREAS GALANIS **FEIDAS CHRISTOS**

This thesis outlines the design, implementation, and evaluation of a biometric authentication system based on EEG that prioritizes privacy through the use of cancellable templates. The proposed framework seeks to improve security and revocability in EEG biometrics, all the while ensuring that identification accuracy remains high. This research draws upon the EEG Motor Movement/Imagery Dataset, which is accessible to the public through PhysioNet. The dataset includes recordings from 109 individuals, featuring 64-channel EEG data gathered during standardized tasks involving both motor execution and imagery.

The EEG signals underwent a series of preprocessing steps, including filtering, artifact removal, and Independent Component Analysis (ICA). This was succeeded by segmentation and feature extraction to effectively extract reliable frequency-domain features. A detailed pipeline was created in Python, utilizing the MNE library alongside custom scripts to produce straightforward templates and implement cancellable transformations. The cancellable design utilizes biohashing techniques, which involve permutation, random projection, and quantization operations that are governed by user-specific keys. This approach guarantees non-invertibility, allows for revocation, and promotes diversity among the stored templates.

The assessment utilized established biometric metrics, including the False Acceptance Rate (FAR), False Rejection Rate (FRR), Equal Error Rate (EER), and Area Under the Curve (AUC). Additionally, it incorporated more sophisticated analyses, such as revocability checks and studies on impostor correlations. The experimental findings indicated that the cancellable templates maintained a level of accuracy that was comparable to that of their plain counterparts, while also demonstrating significant resilience against key-rotation and brute-force attacks.

This research presents a structured and replicable framework for EEG-based cancellable biometrics, effectively connecting the fields of neuroscience, machine learning, and information security. The results underscore the practicality of implementing brainwave-driven authentication systems in everyday settings, all while safeguarding user privacy and ensuring data protection.

ΕΚΤΕΤΑΜΕΝΗ ΕΛΛΗΝΙΚΗ ΠΕΡΙΛΗΨΗ

Σχεδιασμός, Υλοποίηση και Αξιολόγηση Ακυρώσιμων Προτύπων ΗΕΓ με τη χρήση μεθόδων Βιοκατακερματισμού.

ΟΝΟΜΑΤΕΠΩΝΥΜΟ ΦΟΙΤΗΤΗ: **ΟΝΟΜΑΤΕΠΩΝΥΜΟ ΕΠΙΒΛΕΠΟΝΤΟΣ:**

ΑΝΔΡΕΑΣ ΓΑΛΑΝΗΣ **ΦΕΙΔΑΣ ΧΡΗΣΤΟΣ**

Η παρούσα διπλωματική εργασία παρουσιάζει τον σχεδιασμό, την ανάπτυξη και την αξιολόγηση ενός συστήματος βιομετρικής αυθεντικοποίησης με βάση τη δραστηριότητα του εγκεφάλου (EEG), το οποίο ενσωματώνει ακυρώσιμα πρότυπα (cancellable templates) για την ενίσχυση της ιδιωτικότητας και της ασφάλειας των δεδομένων. Ο σκοπός του συστήματος είναι να διασφαλίσει την προστασία της εγκεφαλικής πληροφορίας μέσω μη αντιστρέψιμων μετασχηματισμών, διατηρώντας παράλληλα υψηλή απόδοση ταυτοποίησης.

Για την ανάπτυξη και αξιολόγηση χρησιμοποιήθηκε το δημόσιο σετ δεδομένων EEG Motor Movement/Imagery Dataset από το PhysioNet, το οποίο περιλαμβάνει δεδομένα από 109 άτομα με 64 ηλεκτρόδια EEG κατανεμημένα σύμφωνα με το διεθνές πρότυπο 10-10. Η προεπεξεργασία των σημάτων πραγματοποιήθηκε με φιλτράρισμα, αφαίρεση τεχνημάτων και Ανάλυση Ανεξάρτητων Συνιστωσών (ICA), ενώ ακολούθησε τμηματοποίηση, δειγματοληψία και εξαγωγή χαρακτηριστικών στο πεδίο των συχνοτήτων.

Ο ακυρώσιμος μετασχηματισμός (cancellable transformation) βασίστηκε στη μέθοδο biohashing, η οποία εφαρμόζει διαδοχικά βήματα αναδιάταξης (permutation), τυχαίας προβολής (random projection) και ποσοτικοποίησης (quantization) μέσω κλειδιών χρήστη, εξασφαλίζοντας μη αντιστρεψιμότητα, ανακλησιμότητα και διαφοροποίηση των προτύπων. Η αξιολόγηση πραγματοποιήθηκε χρησιμοποιώντας τις κατάλληλες μετρικές - FAR, FRR, EER και AUC - καθώς και με αναλύσεις ανακλησιμότητας και συσχέτισης εισβολέων.

Τα πειραματικά αποτελέσματα έδειξαν ότι τα ακυρώσιμα πρότυπα διατηρούν ανταγωνιστική απόδοση σε σχέση με τα απλά πρότυπα, ενώ παρουσιάζουν ισχυρή ανθεκτικότητα σε επιθέσεις αλλαγής κλειδιού και brute-force. Συνολικά, η εργασία συμβάλλει στην ανάπτυξη ενός αναπαραγώγιμου και ολοκληρωμένου πλαισίου για βιομετρικά συστήματα με βάση τα εγκεφαλικά σήματα EEG, συνδυάζοντας στοιχεία νευροεπιστήμης, μηχανικής μάθησης και ασφάλειας πληροφοριών.

Table of Contents

Acknowledgments *8*

1. Introduction *18*

1.1 Motivation and Research Objectives **18**

1.2 Problem Definition **18**

1.3 Research Questions **20**

1.5 Methodology Overview **20**

1.6 Contribution of the Thesis **22**

1.7 Thesis outline **24**

2. Theoretical Background *26*

2.1 Biometrics in Authentication Systems **26**

2.2 EEG as a Biometric Modality **32**

2.2.1 Scientific Foundations of EEG 32
2.2.2 EEG Signal Characteristics to Relevant Biometrics 34
2.2.3 Evaluation of EEG Against Biometric Criteria 35
2.2.4 Advantages and Disadvantages of EEG Biometrics 36
2.2.5 State-of-the-Art Applications 36

2.3 Cancellable Biometrics **37**

2.3.1 Definitions and Basic Concepts 38
2.3.2 Types of Techniques of Cancellable Biometrics 40
2.3.3 Benefits and Challenges 44

2.4 Cancellable biometrics **47**

2.4.1 Definition and Concepts 47

2.4.2 Techniques of Cancellable Transformations **49**

2.4.3 Benefits and Challenges 52
2.4.4 Application in EEG Biometrics 56

2.5 Cancellable EEG Biometric Systems in Literature **58**

2.6 Summary of Theoretical Background **65**

3. Data Acquisition and Preprocessing *68*

3.1 Experimental Setup **68**

3.2 Dataset Description **69**

3.3 Signal Preprocessing **70**

3.3.1 Filtering and Artifact Removal 70
3.3.2 Epoching and Segmentation 70
3.3.3 Feature Extraction 71
3.3.4 Methods Summary 72

4. Proposed Cancellable System Architecture *73*

4.1 System Architecture **73**

4.2 Template Generation **74**

4.3 Cancellable Transformations **75**

4.3.2 Key Dependence and Revocability 76

4.4 Security and Privacy Considerations 77

5. System Implementation 77

5.1 Software and Tools 77

5.2 Algorithms for Feature Processing and Template Creation 78

5.3 Cancellable Template Transformation 79

5.4 Evaluation Framework 79

5.5 Summary of Implementation 80

6. Evaluation and Results 81

6.1 Experimental Evaluation Methodology 81

6.2 Evaluation Metrics 82

6.3 Results 82

6.3.1 Plain Templates 82
6.3.2 Cancellable Templates (Real – Valued) 83
6.3.3 Fusion of Scores 84
6.3.4 Revocability (Key Rotation) 85
6.3.5 Impostor Analysis 85
6.3.6 Brute-Force Proxy 86
6.3.7 Genuine Cross-Trial Similarities 86
6.3.8 ROC/EER/ AUC Final Summary 87

6.4 Comparative Analysis with Previous EEG Biometric Studies 88

6.5 Summary of Findings 89

7. Discussion 90

7.1 Interpretation of Results 90

7.2 Security vs Usability Trade-Offs 91

7.3 Limitations of the Current Study 92

7.4 Practical Implications and Applications 93

7.5 Answer to Research Questions 93

8. Conclusion and Future Work 95

8.1 Summary of Contributions 95

8.2 Recommendations for Future Research 97

References 99

Abbreviations

ABBREVIATION	FULL FORM
AUC	Area Under the Curve
BCI2000	Brain–Computer Interface 2000 System
CDF	Cumulative Distribution Function
CRC32	Cyclic Redundancy Check (32-bit)
DET	Detection Error Tradeoff
EDF	European Data Format
EEG	Electroencephalography
EER	Equal Error Rate
EMG	Electromyography
ERP	Event-Related Potential
FAR	False Acceptance Rate
FNR	False Negative Rate
FPR	False Positive Rate
FRR	False Rejection Rate
FT9 / FT10	Fronto-Temporal electrodes (10–10 EEG system)

GAN	Generative Adversarial Network
GDPR	General Data Protection Regulation
ICA	Independent Component Analysis
ID	Identifier
LDA	Linear Discriminant Analysis
MNE	Magnetoencephalography and Electroencephalography (Python library)
PIN	Personal Identification Number
ROC	Receiver Operating Characteristic
SHA	Secure Hash Algorithm
SNR	Signal-to-Noise Ratio
SVM	Support Vector Machine
RMQ	Root Mean Quadratic
PAD	Presentation Attack Detection
ECG	Electrocardiogram

Symbols and Notations

Symbol / Variable	Description
X	Matrix of EEG features (trials × features)
x_i	Feature vector for the i-th EEG epoch
μ (mu)	Mean value of a feature across training samples
σ (sigma)	Standard deviation of a feature across training samples
z_i	Normalized feature vector (after standardization)
W	Random projection matrix (used in biohashing)
k	Secret key seed used for permutation and projection generation
h(x)	Hash function used for biohashing transformation
b	Binary (quantized) template bit vector
τ (tau)	Decision threshold used in verification
α (alpha)	Fusion coefficient for combining scores between plain and cancellable templates
s	Similarity score between two templates
s_{plain}	Similarity score derived from plain template comparison
$s_{cancellable}$	Similarity score derived from cancellable template comparison

s_{fused}	Final fused similarity score after score-level fusion
FAR	False Acceptance Rate → probability of incorrectly accepting an impostor
FRR	False Rejection Rate → probability of incorrectly rejecting a genuine user
EER	Equal Error Rate → operating point where FAR = FRR
AUC	Area Under the ROC Curve → summarizes discriminative performance
TPR	True Positive Rate → ratio of correctly verified genuine attempts
FPR	False Positive Rate → ratio of impostor attempts incorrectly accepted
FNR	False Negative Rate → ratio of genuine attempts incorrectly rejected
ρ (rho)	Pearson correlation coefficient → used for genuine/impostor similarity distributions
R_{k1}, R_{k2}	Revocability comparison sets for two distinct user keys
Δs	Difference in similarity scores across template sets (used in revocability)
N_s	Number of subjects in the EEG dataset (N_s = 109)
N_t	Number of trials per subject after segmentation
f_s	Sampling frequency (160 Hz, fixed throughout pipeline)
T_{epoch}	Epoch duration (2 seconds)
Φ (phi)	Projection transformation function (maps feature space to cancellable template space)

Ψ (psi)	Quantization function applied after projection
γ (gamma)	Regularization parameter (if used in classifier training)
δ (delta)	Small tolerance value for EER estimation precision

1. Introduction

1.1 Motivation and Research Objectives

With the ever-increasing virtual world, the frequency and severity of cyber-attacks have never been higher. Estimated, globally, cybercrime up to 2025 is set to cost approximately $10.5 trillion [36], that is, literally becoming the world's third-largest economy. These astronomical figures serve as an impetus for greater security even beyond traditional methodologies of passwords.

Passwords are vulnerable to various assault methodologies, including phishing. [37, 38] This is a common method where hackers trick users into providing credentials on faked sites. These flaws indicate the vulnerability of relying solely on passwords and the need for biometric authentication.

Biometric authentication is based on the use of unique physical characteristics of each person for the purpose of verifying one's identity before access can be gained into secure data or programs. Put simply, biometric programs verify who someone is by nature rather than what they know (e.g., password) or what they possess (e.g., security token). [39]

Biometric technologies have become very popular over the past couple of years, primarily due to their adoption into smartphones. [40] These technologies have also been becoming less expensive, advanced, and accurate, leading to extensive use across business and personal contexts. There is an interesting 2022 survey that reveals that biometric adoption amongst United States businesses has tripled within three years, mainly to contain growing cyber threats. The nature of cybersecurity these days leads biometric authentication to become a critical area. This thesis, as its top priority, deals with studying other ways of biometric authentication for enhanced security as an augmentation of what is currently being used [41]. While for extremely secure facilities, traditional biometrics such as fingerprint verification or phone authentication may be insufficient, individuals who have high powers or work where they need security most may strongly benefit from other biometric authentication options. This is where Electroencephalogram (EEG) authentication can come into consideration [42].

Fundamental knowledge and literature have been available to date that would enable studying EEG-based authentication systems. Nevertheless, broad-scale adoption and implementation of such technology have not yet become widespread. Hence, the current research is timely and can bridge existing gaps, providing developments along with findings that have the potential to set the future for secure biometrics authentication.

1.2 Problem Definition

Password-based systems don't provide strong security such as traditional methods of biometric authentication such as fingerprints, facial recognition and iris scanners. Although they are a better option from password-based systems, they have a critical limitation: once they are compromised, physical biometric

traits are lost forever because they can't be revoked or replaced. Also, they are open to more advanced talks like template inversion, replay or spoofing.

EEG biometric systems, by making use of characteristic brainwave patterns, may offer enhanced security over traditional biometrics. Owing to the time-varying nature and complexity of EEG waves, they could never be easily impersonated, thereby providing a secure base for verifying users. However, EEG biometric templates, like all biometric data, pose a very high privacy risk. In case they are misplaced in the wrong hands, they may inadvertently reveal sensitive personal information such as cognitive state or neurological disorders/illness.

A potential solution to all of that problem are cancellable biometrics. By applying non-invertible transformations to biometric data. By using those transformations, we have a generation of new templates that cannot be inverted to reveal the original signals. In addition to that, if somehow a template is compromised, it can be "revoked" and replaced by generating a new template with a different transformation key. This has a huge advantage. Preserving both security and privacy at the same time.

Although cancellable biometrics have been explored in modalities like fingerprints and irises, their EEG application remains underdeveloped [99]. The key challenges that emerge, are two:

- How are we going to identify transformation methods that balance discriminative performance, irreversibility and revocability?
- How are we going to evaluate the resilience of those EEG cancellable templates against real adversarial attacks like brute-force, impostor attempts and key rotation?

The gaps that this thesis is willing to address are:

- Creating a cancellable template pipeline (by including permutation, projection, normalization, quantization and Gray encoding) for EEG Biometrics.
- Having extensive evaluations of plain vs cancellable templates by using some of the most known metrics such as AUC (Area Under the Curve), EER (Equal Error Rate), FAR/ FRR (False Accept/ Reject Rate)

- Analyzing the security resilience by doing revocability tests, impostor correlation and brute-force proxy experiments.

By having a combination of accuracy evaluation with systematic security analyses, this work contributes both experimental evidence with practical insights into the cancellable EEG biometric systems for privacy-preserving authentications' feasibility.

1.3 Research Questions

To build upon the problem definition and the objectives that we are called to solve, the following questions are designed to guide the experimental design and the evaluation process:

RQ1: Can EEG-Biometric systems produce templates that are stable and distinctive enough to use them for user authentication under realistic conditions?

RQ2: How are cancellable templates are going to be generated? Which methods will allow us to have discriminability and ensure template security?

RQ3: What are the trade-offs between biometric performance (e.g. accuracy, EER), template cancellability and resistance to potential attacks?

RQ4: If we use fusion strategies to combine plain and cancellable templates will the overall system performance improve without significantly compromising security?

Also we have some hypotheses based on multiple citations existing that may help us navigate the creation of the system in a better manner.

1.4 Hypotheses

H1: EEG signals contain sufficient unique and stable features to enable user authentication with competitive accuracy.

H2: Using cancellable transformations to EEG templates will preserve recognition performance within an acceptable margin compared to plain templates, will significantly enhancing the security of the system

H3: By combining a properly weighted score-level of fusion and cancellable templates will yield better accuracy and robustness of each method alone.

1.5 Methodology Overview

The research methodology adopted in this thesis follows a structured pipeline designed to evaluate the feasibility, security, and performance of cancellable EEG biometrics. The process is organized into the following key phases:

Data Acquisition and Preprocessing: The study utilizes the public EEG Motor Movement/Imagery Dataset from PhysioNet, comprising recordings from 109 subjects. The raw signals underwent a rigorous cleaning process, including bandpass filtering, notch filtering (50Hz) for powerline noise, and Independent Component Analysis (ICA) to remove ocular and muscular artifacts .

Feature Extraction and Normalization: Signals were segmented into task-specific epochs and processed to extract stable frequency-domain features. This resulted in a 1152-dimensional feature vector for each trial. Crucially, Z-score normalization was applied using parameters derived solely from the training set to prevent data leakage and ensure realistic testing conditions.

Cancellable Template Generation (Biohashing): This is the core contribution of the thesis. We implemented a Biohashing pipeline that transforms the plain feature vectors into protected templates. The process involves:

Permutation: Reordering feature indices based on a user-specific key.

Random Projection: Mapping features to a non-invertible subspace using a random matrix.

Quantization: A strategic choice was made to retain Real-Valued outputs (instead of binary) to preserve discriminative information while ensuring non-invertibility.

Evaluation Framework: The proposed system was evaluated using a comprehensive set of metrics and protocols:

Biometric Performance: Assessment via ROC curves, AUC, EER, and FAR/FRR to measure verification accuracy.

Security Analysis: Strict evaluation of Revocability (via key rotation experiments), Unlinkability (impostor correlation analysis), and resistance to brute-force attacks .

Implementation and Tools: The entire framework was developed in Python, ensuring modularity and reproducibility. Key libraries included MNE-Python for signal processing, NumPy for numerical operations, and Scikit-learn for evaluation metrics .

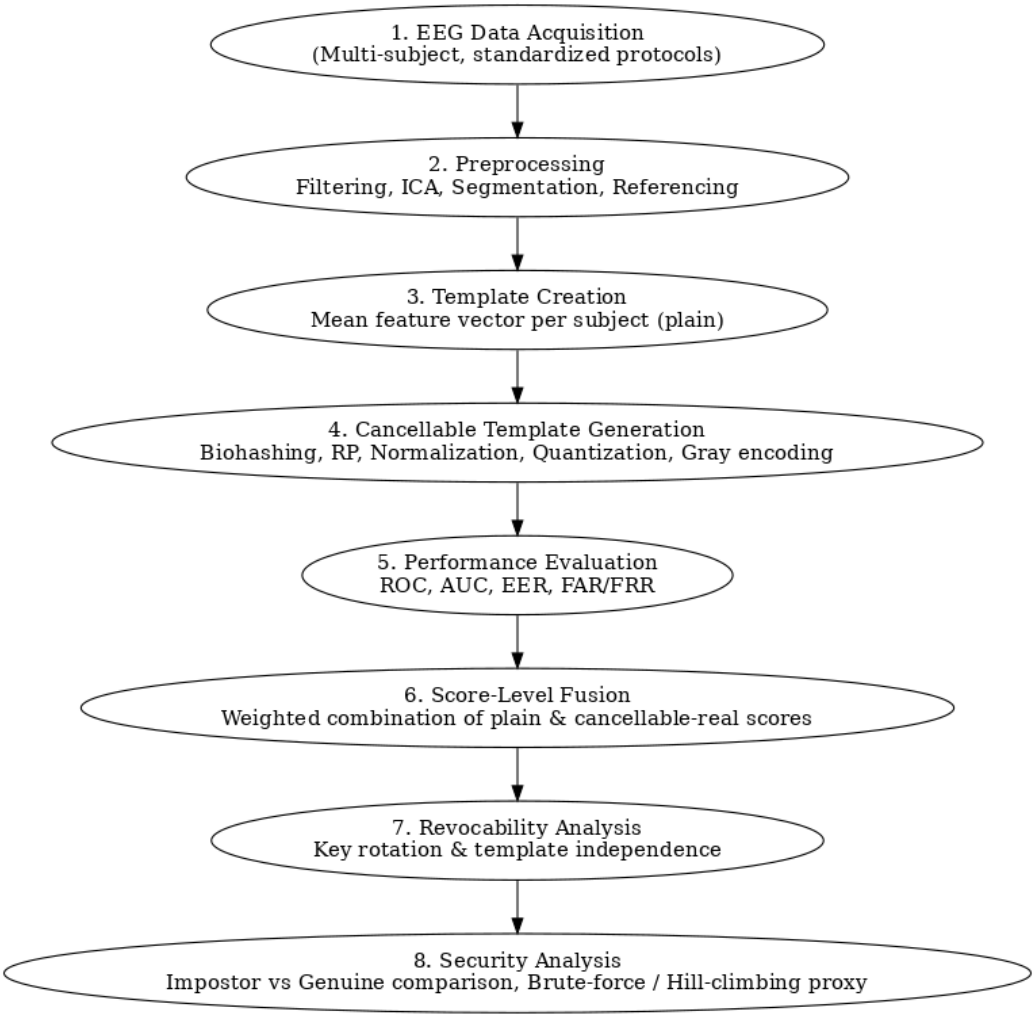


Figure 1: Methodology Overview

1.6 Contribution of the Thesis

This thesis is based on the framework proposed in "Cancellable Template Design for Privacy-Preserving EEG Biometric Authentication Systems" [43]. There is a reproduction of cancellable EEG biometric pipeline and an extension of it with new analyses, evaluations and security checks. Some core concepts follow the original work made on this paper such as EEG Preprocessing, template creation and biohashing-based cancellable transformations, our contribution adds elements that have not previously explored. Our contributions are as follows:

1. Reproduction and Validation of the Baseline Pipeline

• A full implementation of the preprocessing, feature extraction, template creation using the EEG Motor Movement/ Imagery Dataset

• Verification of reproducibility of the cancellable template framework under real conditions.

2. Dual Template Pathway (Plain vs. Cancellable-Real)

• Implementation and comparison of plain templates against cancellable-real templates (continuous outputs before quantization).

• Demonstration that cancellable-real templates have similar matching performance to plain templates when cosine similarity is used for comparison.

3. Performance Evaluation

• Evaluation of the system with established biometric metrics (Receiver Operating Characteristic (ROC), AUC, EER, FAR/FRR at fixed thresholds).

• Generation of ROC/ DET plots and detailed error analysis that shows that plain (AUC ≈ 0.666, EER ≈ 0.372) and cancellable-real templates (AUC ≈ 0.662, EER ≈ 0.372) achieved similar discriminability

4. Score Fusion Between Template Types

• Applied weighted fusion between plain and real-cancellable scores using min-max normalization

• Performance is maintained at fusion (AUC ≈ 0.671, EER ≈ 0.372) and reduced FRR at low FAR (≈0.85% at FAR=1%)

5. Revocability and Security Analyses

• Conduction of revocability tests by rotating keys and comparing template similarity across different key instantiations, confirming low correlation between reissued templates.

• Execution of impostor vs. genuine correlation analysis, showing strong separation between impostor similarities (∼0.75) and genuine similarities (∼1.0)

• Imitation of a brute-force/hill-climbing proxy experiment, revealing a large margin between impostor and genuine data, demonstrating robustness against guessing attacks.

• Creation of ROC-based verification analysis combining impostor and genuine trials and having near-perfect separation (AUC ≈ 0.9989, EER ≈ 8.46×10^{-4}).

6. Reproductible, Modular Implementation

• Development of a modular and easily reproductible python codebase that automates the entire workflow

• By having this implementation and having a script for each step of the pipeline, this leads to reproducibility for further experimentation with alternative cancellable transformations.

Acknowledgment of AI Assistance: It is hereby acknowledged that generative AI tools (specifically ChatGPT) were utilized as a linguistic aid during the writing of

this thesis. The use of this tool was strictly limited to the refinement of the English language, including grammar correction and the enhancement of academic vocabulary. The scientific conceptualization, methodology design, code implementation, data analysis, and interpretation of results remain the original and exclusive work of the author.

1.7 Thesis outline

The rest of the thesis is structured in this fashion:

2. Theoretical Background

This chapter presents the principles of biometrics and authentication systems, focusing on EEG-based biometrics. It additionally covers fundamental neuroscience principles (e.g., the anatomy of the brain). It describes the features of EEG signals, the functionality of cancellable biometrics, the employed approaches, and the associated benefits and problems arising from this research.

3. EEG Data Preprocessing

This chapter introduces the EEG dataset used in this thesis attentively, by describing its structure, tasks and its appropriateness for biometric research. It thereupon outlines the preprocessing pipeline applied to the raw EEG signals (including all steps such as filtering, artifact removal etc.). At last, it explains the procedure for feature extraction that transforms the preprocessed signals into numerical feature vectors suitable for creating biometric templates and it' s successive cancellable transformations.

4. Cancellable EEG Biometric System Design

Here, the architecture of the proposed cancellable EEG biometric system is shown. Next up, we have a detailed view of the generation of plain templates based on the feature vectors extracted from EEG signals. Then cancellable transformations are explained and after that, how the biohashing pipeline works (permutation, projection, quantization). Finally key dependence, revocability are mentioned and what do we need to consider security and privacy-wise.

5. Implementation

This chapter analyzes software and tools used (e.g MNE, NumPy etc.) Also in this part the algorithms used to move from EEG signals to feature vectors to templates are analyzed such as feature processing and template construction to end with an evaluation of the framework.

6. Evaluation and Results

This section presents the methodology used and the evaluation results of the proposed cancellable EEG biometric system. First there is an outline of the metrics used (e.g. FAR, FRR, EER etc.) and then there is a report of the system's performance across plain and cancellable templates (real and Gray-coded) and fused scores. Additionally, it examines the revocability through key rotation experiments conducted, analysis of impostor behavior and tests robustness against brute-force

proxies and genuine cross-trial similarities. It finally summarizes overall verification performance.

7. Discussion

The discussion chapter interprets the results created during the experiment and emphasizes on the security vs usability trade-offs. It also mentions limitation of the current study and also discusses the practical implications of the results, showing how EEG biometrics could be applied in to real-world scenarios and what challenges remain for future adoption.

8. Conclusion and Future Work

This final chapter, provides a summary of the main contribution of the thesis bringing the most important parts of each section of this thesis. It reflects on how it contributes to cancellable EEG biometrics and looks forward offering recommendations for future work for practical authentication systems.

2. Theoretical Background

2.1 Biometrics in Authentication Systems

Biometric authentication systems check to see if a person is who they claim they are by comparing a current biometric sample to a saved reference template. In a regular biometric authentication situation (also called verification or 1:1 matching), a person gives the sensor a biometric sample, like a fingerprint scan or a photograph of their face, after giving it an ID like a username or ID number. The system uses parts of this sample to compare them to the biometric template that was already saved with the claimed identity. The identity claim is confirmed and the user is authenticated if the sample fits the stored template well enough (i.e., the similarity score is over a specific level). This one-to-one verification method is often used to unlock devices (such cellphones that use fingerprint or facial recognition), secure access control systems, and check people's identities at borders or secure facilities [1]. It's an easy approach to change passwords because the user doesn't have to remember or carry anything; their biometric feature is the key [49].

Biometric systems use identification mode (one-to-many or 1:N matching) to compare a person's biometric sample to all the records in a database to find out who they are. They do this without the person telling them who they are. [50, 51] Identification is often used in forensic settings, such matching a fingerprint to a criminal database, or in surveillance, like scanning crowds for a wanted individual using facial recognition. But as the database grows, identification tasks require more computing power and may be less accurate. Many authentication systems use verification mode for normal login and authentication [1] since it is faster and has fewer mistakes for a given number of persons. Biometrics are a means to tell who someone is by looking at their distinctive traits. They are usually used by themselves or as part of multi-factor authentication methods to make security stronger. For example, a fingerprint scan with a PIN.

Components and Workflow of a Biometric Authentication System

Systems of biometric authentication they follow a standardized workflow and architecture. The following figure outlines the typical components and the process flows of a biometric system from start (enrollment) to end (verification)

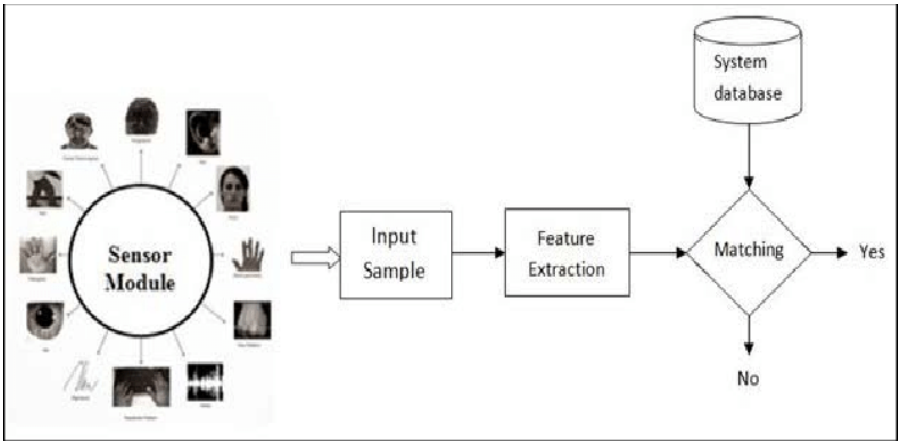


Figure 2.1 1 Diagram of Biometric Authentication Workflow

The major components and steps include:

• Sensor (Data Acquisition): The first module is a sensor or capture device that gets the user's raw biometric data. For example, this could be a fingerprint reader that takes a picture of a fingerprint, a camera that takes a picture of a face, a microphone that records a voice sample, etc. The quality of this sensor data is very important because if the input is noisy or distorted, the system's performance can go down.

• Feature Extraction: Next, the algorithm evaluates the raw data to find a small set of unique features that represent the biometric trait. For example, when recognizing a fingerprint, the feature extractor might find minor details like ridge ends). Similarly in face recognition, it might find a vector of face embeddings. This is the phase of signal processing and pattern recognition. The goal is to build a biometric template, which is a mathematical model that captures the most important features of a person's trait while getting rid of any unnecessary information.

• Database (Template Storage): During enrollment, the retrieved biometric template and the person's identification label (like their user ID) are both stored in a database or a secure storage module. As part of a verification system, the database might only have one template for each user, or it might have more than one template if more than one trait or sample is enrolled. The database in identifying systems has templates for a lot of people. To keep others from getting into or changing templates, they must be stored safely, usually with encryption or other template protection methods.

• Matching Module: When a user tries to authenticate later, a new biometric sample is taken and processed through feature extraction to make a query template. Then, the matching module makes a comparison between this query template against the stored template(s). In the verification stage (1:1), it retrieves the single template corresponding to the claimed identity and makes a computation of a similarity score/ distance between two templates. In identification stage (1:N), the matcher makes a comparison of the query to all candidate templates in the database to find the best match (if this exists). The comparison uses some algorithms specifically created to check the biometric modality (e.g computing correlation or Euclidian distance between feature vectors).

• Decision Logic: Now, the system makes a decision for the authentication based on the match score and the set of criteria. In a verification scenario, the system will accept the identity claim if the similarity score above a specified threshold, suggesting the biometric sample matches the enrolled template with sufficient confidence; otherwise, it will reject the login attempt. In identification, the system may determine an identity by selecting the highest match that exceeds a specified threshold or may produce a "no match" result if none of the comparisons are persuasive. There is also a possibility to make the decision more reliable to incorporate additional logic, such as allowing a limited number of retries or using multi-biometric fusion (combining evidence from multiple traits) to improve reliability [52, 53].

In short, this workflow is a series of steps: capture, feature extraction, matching, and decision [1]. The sensor must consistently collect data, the feature extraction must provide templates that are both distinct and repeatable, the matcher must be both accurate and quick, and the decision threshold must be set to find the right balance between security and usability. The system captures and extracts features during enrollment to save the user's reference template. It captures, extracts features, matches, and makes a decision during authentication (verification).

Let’s now, discuss for the system in case it’s in identification mode. The matching and decision processes involve looking to several templates and choosing this that has the greatest score. This is the identification mode. Preprocessing is included (to improve quality of the data or make inputs more uniform) and may need quality control (to reject bad samples before matching).

Those authentication systems have a broad range of situations that can be used. From unlocking your mobile device to limit access in high-security areas, just by following this pipeline.

Benefits of Biometric Authentication vs Traditional Methods

The gap between biometric authentication versus traditional knowledge-bases or token-based authentication methods is huge with biometric authentication providing several notable advantages such as:

• Positive Identification of Persons: Biometric methods verify an individual's identity by direct examination rather than by assessing knowledge or possession. This means that biometrics make sure that the right person is present when they are working correctly. But passwords and tokens only check the knowledge or object that someone else could use. Biometrics create a stronger link between an identity and the person it belongs to, which makes authentication more secure and trustworthy [1, 54].

• Convenience and Usability: Users don't have to memorize complicated passwords or carry extra devices for biometric authentication. The authentication element is something about their body or conduct, which makes it very easy. For instance, using a fingerprint or face scan to unlock a smartphone can be quicker and easier than entering a PIN. The user doesn't have to remember any passwords, keep any cards safe, or have anything else on hand. This can lower the cost of help desk calls for password resets and make the user experience better [54, 55].

• Cannot be Easily Lost or Shared: Biometric attributes are unique to the user and cannot be easily misplaced, forgotten, or left behind; one consistently possesses their fingerprints, facial features, voice, etc. Also, it's not as easy to give or lend them to someone else as it is to share a password or access card. A person can't accidentally give away their fingerprint like they may give away their PIN (since utilizing the biometric usually requires the user to be there). This feature makes security better by lowering the chances of credential sharing and theft [3, 55, 56].

• Resilience to Guessing or Brute-Force: It's easy to guess, default, or crack passwords with brute-force attacks, especially if the passwords are weak. Biometric identifiers are very hard to "guess." There are an unimaginable number of ways that a fingerprint pattern or iris code could look, and an attacker can't just list all the faces or fingerprints in a reasonable amount of time. For example, a strong biometric system with a very low false match rate (like 0.001% FAR) is about as hard to guess as a 5-digit random PIN. The attacker would need real or fake samples of biometrics to test their guesses, which is much harder than automated password guessing [1]. In general, it's considerably harder for someone pretending to be someone else to make a lot of good biometric samples than to guess passwords.

• Persistent and Hard to Change: Biometrics use traits that stay a long time, such physical and behavioral ones. For instance, fingerprints and iris patterns don't vary

much over a person's life, and even things like voice and stride alter slowly over time. Users don't have to change or update their "credentials" as often as they do with passwords, which can expire or need to be updated often. This is because they are permanent. But this can also be a bad thing, as you can see in the limits below. But since it lasts, a person can be properly identified throughout time without having to continuously upgrading their credentials [3].

In short, biometric authentication is a simple way to be sure that you are who you say you are [55, 56]. It uses the fact that everyone is different to make security that is hard to break through normal methods like guessing passwords or stealing IDs. These benefits have led to the widespread use of biometrics in situations where secure and easy authentication is needed. For example, fingerprint or face recognition can be used to unlock personal devices and log into banking apps, while e-Passport gates that use facial or iris recognition can be used to automate border control.

Limitations of Biometric Authentication

Although biometric authentication, has many advantages, also comes with important limitations and challenges, especially when comparing it with traditional methods.

• False Matches and False Rejections (Accuracy Issues): There is always a chance of mistakes with any biometric technology. There are two kinds of mistakes that might happen: a false acceptance (wrongly matching an impostor to someone else's template, sometimes termed a false match) and a false rejection (not matching the real user to their own template). Pattern recognition is a part of biometric matching, and because it is based on chance, things like sensor noise or natural variation can create mismatches. In the actual world, false rejection rates can be significant (in certain circumstances, double-digit percentages), which can be annoying for users who are wrongly refused access. On the other hand, security requires very low false acceptance rates [3, 57, 58]. One of the biggest challenges is finding the right balance between security and ease of use by changing the thresholds. A right password will always work, and an incorrect one will always fail. Biometric data, on the other hand, creates a gray area of ambiguity that needs to be handled.

• Variation and Quality Dependence: Changes in the individual or the environment can impact how well biometrics work. Injuries (like a cut on a finger), getting older, being sick (like a cold that makes your voice sound different), or even the light in the room can all affect the biometric reading. The system may not be able to verify the user if the sample quality is bad, which could be because the sensors are filthy, there is too much noise, or the user made a mistake. Some people may have non-universality concerns, which means that a useable biometric attribute isn't available or can't be captured adequately. For example, manual laborers may have worn fingerprints, while people with speech disabilities may have voice problems. These things can make it hard for some people to enroll (when they can't find a decent template on record) or cause high false reject rates. There is no biometric-specific problem with traditional techniques. Anyone can get a password or token, but not everyone can give a readable fingerprint.

• Lack of Revocability: One of the biggest problems with biometric IDs is that they can't be changed very easily. If a person's password or credit card number is stolen, they can reset it or get a new one. If a biometric template or fingerprint is taken, though, it's not

easy for the user to get a "new" fingerprint. The fact that biometrics last forever is a good thing, but it also means that if they are compromised, they might be a problem. You carry that fingerprint for life. This is worrying because attackers might utilize a stolen biometric over and over again. There are ways to lessen this risk, such as cancelable biometrics or biometric cryptosystems, which store an altered version of the biometric that can be reset (revoked and replaced) if it is compromised. But using these methods makes things more complicated, and they aren't used by everyone yet. Because biometric identifiers are hard to replace, it's harder to recover from security breaches than with regular credentials [58].

• Security Vulnerabilities (Spoofing): You can't "guess" a fingerprint like you can a password, but attackers can still get into biometric systems in other ways. Spoofing or presentation assaults are a big danger. In these attacks, someone uses a phony biometric artifact to trick the sensor. For example, they might use a lifted fingerprint mold, a high-resolution snapshot of someone's face, or a recorded voice sample. Many experiments have shown that spoofing works (for example, using gelatin "gummy fingers" to trick fingerprint scanners [4, 57, 58]). Biometric systems have to deal with the fact that the input trait could be fake, unlike a password system where the secret is either right or wrong. This means that more security measures are needed, such as liveness detection (making sure the trait comes from a living person and not an object) and multi-factor checks. In the next part, we talk about these security issues in more depth.

• Privacy and User Acceptance: Some people may find biometric authentication to be obtrusive. Passwords don't pose privacy concerns like collecting personal biological data does. Users might be worried about how their fingerprint or face image is saved and utilized. For instance, they might be worried that it could be used for surveillance or that someone could steal their biometric data and use it to steal their identity. Because biometrics are so distinctive, they could be used to follow people across different systems if data is shared (this is called "function creep," when biometric data obtained for one purpose is utilized for another). In many places, privacy laws recognize biometric data as private personal information and set tight rules for how it can be collected, stored, and used [57]. Also, some users don't like particular types of scanning (for example, retina or iris scans have always made people worry about eye safety or hygiene, and fingerprint scanners were originally looked with mistrust). Because of these acceptability concerns, setting up a biometric system needs careful thought about getting user permission, protecting their data, and being clear about how the data will be used.

• Cost and Infrastructure: Biometric authentication may need specialized hardware (e.g fingerprint readers, iris scanners) and integration with existing systems. This can be a pricier solution than purely password systems. Also another factor that we need to consider is that sensors needs to be calibrated and clean and there may be throughput limitations (e.g processing multiple people at once is not feasible). For sure in some contexts, these considerations can be disadvantages in comparison to the simplicity of passwords or tokens. However in our economy, biometric sensors have become inexpensive and the barrier is gradually lowering even more [57,58].

In short, biometric authentication isn't the answer to all problems; it comes with its own set of problems that need to be solved. Because of high error rates, the fact that they are permanent, the fact that they can be spoofed, and privacy concerns, biometric systems need to be built with robust security measures and user education in mind. Technical and procedural solutions can help with many of these problems. For example, utilizing two-factor authentication, which combines biometrics with a PIN, can help with some security issues, while using privacy-enhancing biometric template safeguards can help

with some privacy issues. It is important for everybody who uses biometric authentication to know about these limits as well as the benefits.

Security and Privacy Concerns in Biometric Systems

Because biometric systems handle sensitive personal information and make judgments that are important for security, there are a number of important privacy and security issues that come up. These include the potential of presentation assaults (spoofing), the need for biometric data to be irreversible, and the wider privacy issues that come with using biometrics.

• Spoofing and Presentation Attacks: One of the biggest security risks to biometric authentication is spoofing, which is when an attacker pretends to be someone else by giving the system a bogus biometric sample. Biometric features, on the other hand, can occasionally be acquired or duplicated from a person without their awareness. For example, an attacker could lift a fingerprint from a surface that was touched or get a high-quality picture of someone's face. Attackers have shown how to make fake fingerprints out of things like silicone or gelatin (dubbed "gummy fingers"), which can trick fingerprint scanners into thinking the fake user is the real one [4]. In the same way, showing a picture to a facial recognition system or playing back a recorded voice to a voice authentication system are other types of presentation attacks. These assaults take use of the fact that the biometric sensor can't always tell the difference between a living trait and an artifact. Modern systems use Presentation Attack Detection (PAD) methods, which are also known as liveness detection, to stop spoofing. Some examples are fingerprint scanners that check for blood flow or electrical conductivity in the finger, facial recognition that needs blinking or 3D face depth detection, and voice recognition that uses challenge-response (asking the user to say a random phrase) to make sure the input is live [59, 61, 62].

• Template Security and Irreversibility: Biometric systems need to save reference templates of users' attributes, which are private information. It is a basic rule that saved biometric data should be safe and impossible to change. It should be impossible for a computer to recreate the original biometric data (like a fingerprint picture or voice recording) from the stored template. This irreversibility is crucial because if an attacker could reverse a biometric template, they might get the raw trait and use it to trick other systems or invade the user's privacy. To do this, systems usually don't keep the raw images or signals. Instead, they keep processed feature templates that are sometimes even more secure by using encryption or one-way transformations. For example, instead of storing the fingerprint picture directly, a fingerprint authentication system might keep a list of minutiae points or a proprietary hash of the image. Even that template could be secured. The concept is that even if the database is hacked, the thief just gets some general information that isn't easy to use to replicate the biometric or pretend to be the user. Biometric cryptosystems and cancelable biometrics are two examples of ways to improve template security. In cancelable biometrics, a reversible mathematical function (parametrized by a user-specific key) changes the raw features before they are stored. If the database is hacked, the transformation can be changed (like giving a new password) to cancel the old biometric template and replace it with a new one that comes from the same trait [7, 60, 61].

• Privacy Implications: Using biometrics raises serious privacy concerns. Biometric indicators are intrinsically linked to an individual's identity and uniqueness. If not used properly, they could let someone track or profile other people without their permission.

For example, a fingerprint or face image used in one system, like a workplace access control, may be cross-matched with data from another system, such a law enforcement database or a marketing database if it was stolen, to find or watch the person without their permission. This function creep is a well-known problem: data collected for security reasons could later be used for other investigations, including tracking attendance or even figuring out things like ethnicity or health conditions from facial images, which the user didn't agree to. Biometrics can also store supplementary information. DNA can show genetic links or health problems, while looks can occasionally show mood or demographic information. People are frightened about how much personal information is being gathered and looked at without their knowledge. Many jurisdictions have developed data protection regulations that make it apparent that biometric data is particularly private in order to deal with these issues [59, 60].

Biometric systems should receive informed consent, only use data for the specified authentication purpose, keep data safely (ideally in encrypted form or on secure elements on personal devices), and let users delete or withdraw their biometric data. Some systems don't use centralized databases at all. Instead, they use biometric authentication on the device itself, like modern smartphones do. This means that your fingerprint template never leaves the device. Also, clear rules for users and maybe even outside oversight could help consumers feel like their biometric data won't be used against them. Here, privacy and security are very closely related. A system that is safe from outside threats but doesn't care about how it shares or uses biometric data in new ways might nonetheless violate people's privacy and make them less trusting in biometric technology [59, 61].

2.2 EEG as a Biometric Modality

2.2.1 Scientific Foundations of EEG

Electroencephalography (EEG) is a non-invasive neurophysiological method for recording the electrical activity of the brain through electrodes placed on the scalp. Neural populations in the cortex produce synchronized post-synaptic potentials, and when sufficiently large groups of neurons discharge simultaneously, these electrical fluctuations transmit through the skull and scalp, where they may be identified as minor voltage differentials by surface electrodes [1, 63]. EEG signals are usually tens of microvolts, which means they need to be amplified and carefully filtered to be useful. EEG has a low spatial resolution because signals spread over many layers of tissue, but it has a high temporal precision (on the order of milliseconds), which makes it great for recording quick brain activity [2, 63, 64].

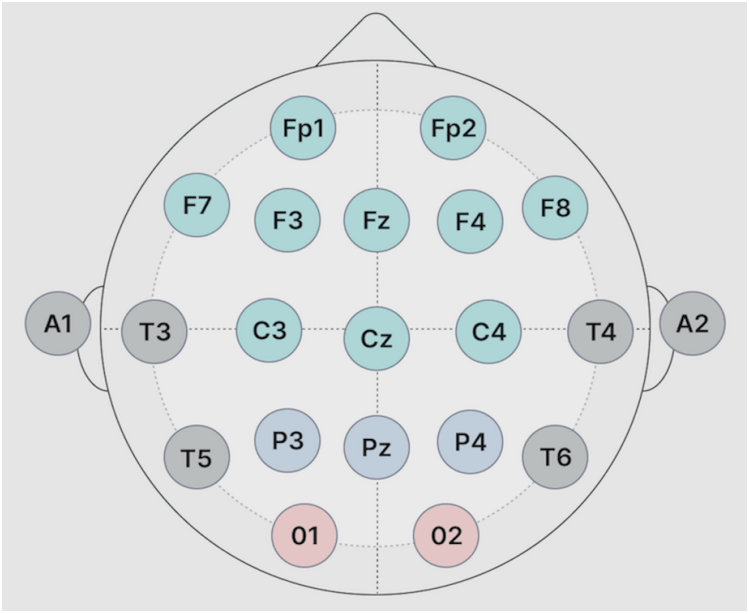


Figure 2.2 1 Diagram of EEG electrode placement on the scalp (10-20 layout)

EEG signals are decomposed, most of the time into canonical frequency bands: delta (0.5–4 Hz), theta (4–7 Hz), alpha (8–12 Hz), beta (13–30 Hz), and gamma (>30 Hz) [3, 62]. Each of those bands symbolizes a different state of the brain. For instance, delta with deep sleep, theta with drowsiness and memory processes, alpha with relaxed wakefulness, beta with active concentration and gamma with higher-order cognition and sensory processing [4]. Beyond those spontaneous oscillations, EEG also captures event-related potentials (ERPs), which are normal voltage fluctuations that occur as a response to internal or external stimuli [63]. The P300 and N400 are two examples. These ERPs have been employed in biometric systems due to their ability to capture individual differences in cognitive functioning.

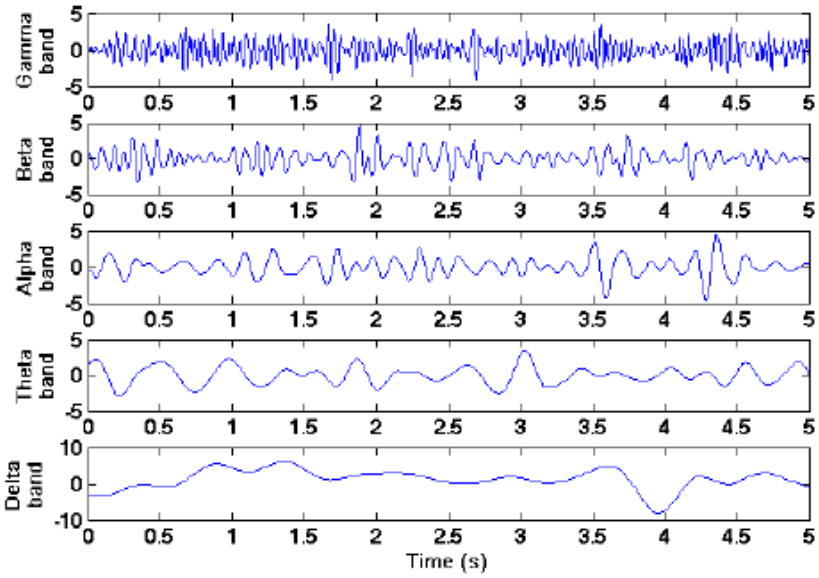


Figure 2.2 2 Frequency bands of EEG signal

2.2.2 EEG Signal Characteristics to Relevant Biometrics

EEG exhibits several defining properties that influence its suitability as a biometric modality:

• Non-stationarity: EEG is highly variable across time, even when we' re talking for recordings from the same individual under similar conditions. Factors including fatigue, stress, attention level and electrode placement are some of the factors that play a major role in making stable template generation challenging. [6, 65]

• Low Signal-to-Noise Ratio (SNR): SNR is notably low in EEG recordings, as the amplitudes are quite small and susceptible to distortion from various artifacts, including eye blinks, muscle activity, or electrical interference. Techniques in signal processing, such as filtering, Independent Component Analysis (ICA), and artifact rejection, play a crucial role in improving the quality of data [7, 65].

• Individual Variability: Even with the inherent differences among individuals, EEG signals reflect unique structural and functional characteristics of each person's brain. The distinctiveness observed between individuals serves as the basis for utilizing EEG as a biometric measure, as no two people generate the same EEG patterns [8, 65, 66].

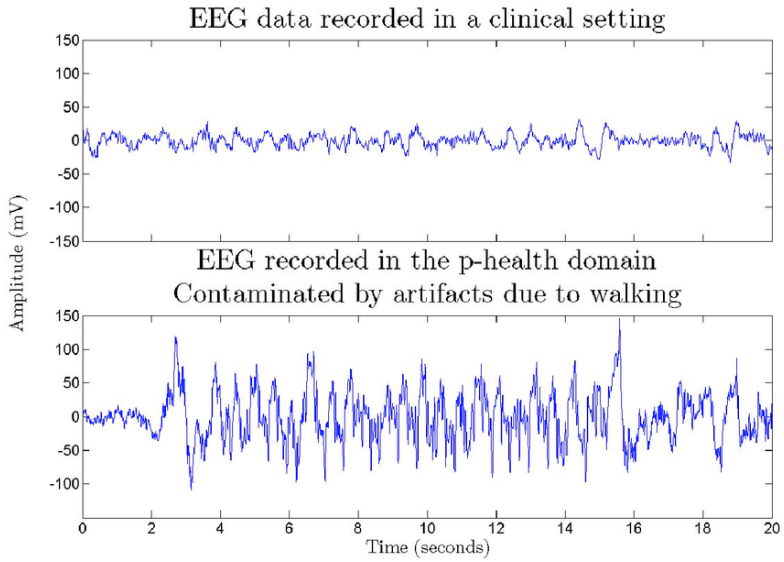


Figure 2.2 3 Contaminating effect of artifacts on recorded EEG signals

2.2.3 Evaluation of EEG Against Biometric Criteria

Standard criteria are usually used to judge biometric modalities. These include universality, uniqueness, permanence, collectability, performance, acceptability, and resistance to circumvention [9, 67]. The pros and cons of EEG can be summed up as follows:

- *Universality:* Every living person with a healthy brain produces EEG activity. EEG is always present, unlike fingerprints (which can be missing because of damage) or voice (which can be affected). This meets this requirement [10, 67, 68].
- *Uniqueness:* Research indicates that EEG "brainprints" exhibit sufficient variability among individuals to facilitate their classification. Research in machine learning has shown that EEG can accurately recognize subjects 80% to 98% of the time, showing that it has very discriminative properties [11,12, 67].
- *Permanence:* EEG is not as stable as other types of biometrics. Template "aging" happens because of changes in the body or mind, which make it harder to recognize over weeks or months [13]. To stop this drift, systems often need to have their templates updated.
- *Collectability:* EEG used to need heavy electrode caps with conductive gel. Recent improvements in dry electrodes and portable headsets, on the other hand, make it easier to use, but acquisition is still more intrusive than face or fingerprint scanning [14].

- *Performance:* With deep learning models, EEG authentication accuracy is now over 95% on controlled datasets [15, 69]. But performance may get worse in very loud places, especially if you use it for a long time.
- *Acceptability:* User comfort and societal attitudes are still problems. It's less convenient and more obvious to wear an EEG equipment than to use other biometrics [16].
- *Resistance to Circumvention:* EEG is the best at this. Brain activity cannot be easily collected or "left behind" without authorization. It has built-in liveness recognition (no EEG in dead people), which makes it one of the hardest to fake [17].

2.2.4 Advantages and Disadvantages of EEG Biometrics

Advantages:

- High resistance to spoofing and replay attacks [17, 70].
- Inherent liveness detection (EEG ceases upon death) [17, 70].
- Privacy-friendly: raw EEG is difficult to interpret without specialized processing [18, 70].
- Potential for “cognitive passwords” (stimulus-evoked EEG responses unique to each user) [19].

Disadvantages:

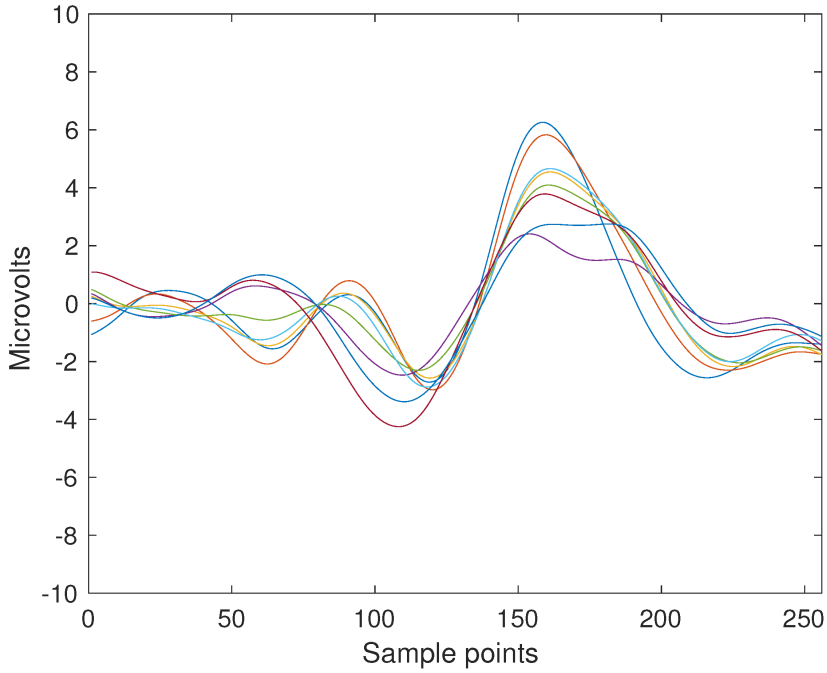
- Requires specialized hardware and user cooperation [14].
- High intra-subject variability reduces permanence [13, 71].
- Lower user acceptability due to headsets and setup complexity [16, 71].
- Still immature compared to traditional biometrics (smaller datasets, limited real-world trials) [20, 70, 71].

2.2.5 State-of-the-Art Applications

EEG biometrics have been applied in:

- Resting-state authentication, using spontaneous brain rhythms [21, 72].
- Task-based EEG, e.g., motor imagery or cognitive tasks to increase distinctiveness [22].
- ERP-based authentication, where responses to known/familiar stimuli provide a “brain password” [5, 73].

- Deep learning pipelines, leveraging CNNs or transformers for end-to-end EEG classification with high accuracy [15].



2.3 Cancellable Biometrics

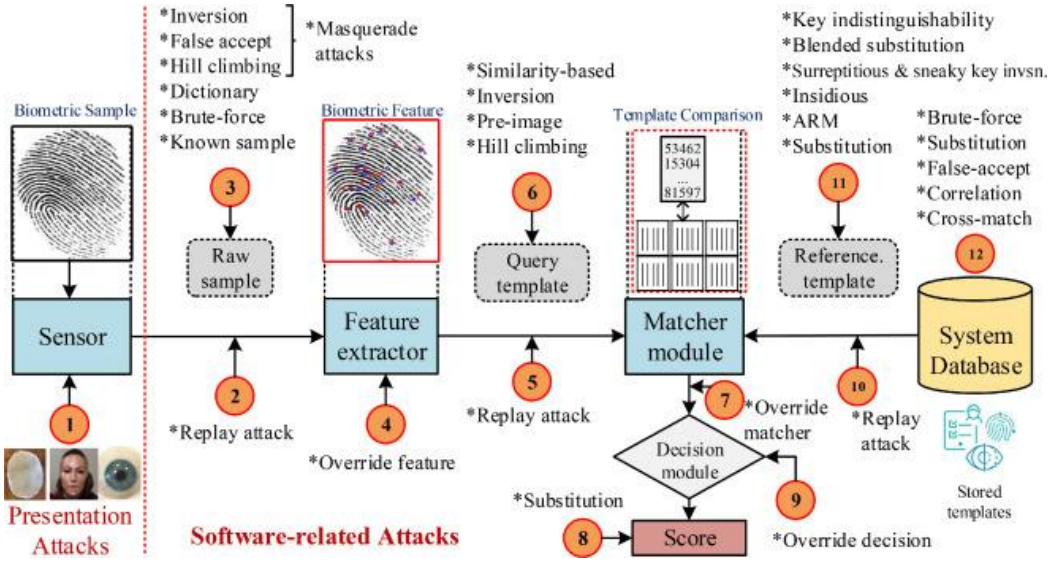


Figure 2.3 1 List of attacks at different vulnerable points on a biometric authentication system.

2.3.1 Definitions and Basic Concepts

Biometric template protection schemes are designed to keep sensitive biometric data safe from being stolen or misused. This is especially important because biometric traits, like fingerprints or EEG signals, are permanent and can't be changed if they are stolen. Cancelable biometrics solve this problem by changing the biometric template in a way that can't be undone before storing or comparing it [16]. Ratha et al. [16] were the first to talk about cancelable biometrics as a "intentional and repeatable distortion" of biometric data through a one-way transformation. A non-invertible transformation of the features is stored instead of the raw biometric features [74]. This means that an attacker can't get the original biometric signal back from the transformed template, even if they steal it. The transformation can happen in the signal domain (like when an image or time series is distorted) or in the feature domain (after feature extraction) [16, 74]. In either case, matching happens in the transformed domain, so the person who is authenticating never has to look at the original biometric data during the verification process. This method is designed to improve privacy and security.

Revocability, or reusability, is another important feature of cancelable biometrics. If a transformed template is compromised, it can be canceled and replaced with a new one, just like resetting a password, without the user having to give a new biometric trait. As a result, each user's biometric can create many different secure templates that can be used in different applications or for re-enrollment. The main requirements for a good cancelable biometric system are listed below [16, 18]:

- *Non-invertibility:* The transformation only goes one way. If you have a transformed template, it should be impossible to use math to figure out what the original biometric feature vector was. This keeps the user's raw data, like raw EEG signals, safe from being seen [75, 76].
- *Revocability:* If the current template is compromised, it should be easy to revoke it and issue a new one using the same biometric data. The new transformed template should be very different from the old one so that an attacker can't use a compromised template to guess the new one [77, 78].
- *Diversity:* A user can make unique templates from the same biometric by using different transformation parameters (like keys) for different applications. It should not be possible to link templates of the same user in two different databases, which would keep privacy across systems [79, 80].

- *Acceptable Performance:* The transformation should not significantly lower the accuracy of biometric recognition (for example, EEG authentication accuracy). The transformed templates should keep the information that lets you tell users apart, so that they are just as accurate as the original features [80].

These features make sure that a cancelable biometric system is safe and useful. If designed correctly, only the altered biometric reference is kept (like in an EEG authentication database), and it doesn't give away any useful information about the original EEG signals (see Figure 2.3.2 for a conceptual comparison of original and altered templates). Also, because each application can use a different transformation or key, a breach in one database doesn't affect the security of the biometric in other contexts (the stolen template from application A can't be used to impersonate the user in application B). This variety across applications is very important for privacy when the same biometric (like an EEG profile) is used in more than one system.

It is important to note that there are two main types of biometric template protection methods: feature transformation (also called cancelable biometrics) and biometric cryptosystems [81, 82] . In the first case (which is what we're looking at here), a non-invertible function directly changes the features, usually with a secret parameter like a key or token [16]. In biometric cryptosystems, error-tolerant cryptographic constructs (like helper data and fuzzy vaults) use biometric data to create or bind keys [18, 82]. Cancelable biometrics usually have simpler matching in the transformed domain and let you update multiple templates. Biometric cryptosystems, on the other hand, focus on protecting a cryptographic key and need error-correction methods to deal with differences between users. Both methods try to keep raw biometric data from getting out, but cancelable transformations are especially good for re-issuing templates.

Application to EEG Biometrics: Although cancelable biometrics have been thoroughly examined for fingerprints, facial recognition, iris patterns, and other characteristics [18], their utilization in EEG signals is still developing. EEG biometrics are hard because brainwaves have a lot of personal information and the signals can change. It is especially important to protect EEG-based templates because raw EEG data can show not only who someone is, but also their health problems or mental state. Some early research has started to look into ways to make EEG cancelable. For instance, He et al. [21] illustrated an EEG authentication method in which the features (autoregressive coefficients from EEG) were hashed through a random projection technique (Johnson-Lindenstrauss

transform [83, 84]), successfully translating EEG features into a secure hash code. Bajwa and Dantu [22] also came up with NeuroKey [23], a way to make cryptographic keys from EEG signals. They call it "cancelable" because a new key can be made from EEG recorded during a different task if the old key is compromised. But these methods are still pretty new. He et al.'s method employs a unidirectional projection (a form of transformation), yet it does not explicitly facilitate template revocation aside from recomputation. Conversely, the NeuroKey scheme permits key modification through alterations in the cognitive task, but it fails to implement a consistent transformation on the identical underlying EEG features. As of now, a specialized cancelable template design specifically for EEG biometrics is still mostly undeveloped. This shows a gap that current research is starting to fill. The next sections talk about the kinds of non-invertible transformations that are used in cancelable biometric techniques. They focus on how these transformations can be used in EEG-based systems.

2.3.2 Types of Techniques of Cancellable Biometrics

There are a number of ways to use cancelable biometric transformations. This part talks about well-known non-invertible transformation methods that are important for EEG biometric templates [86, 87]. These methods include biohashing, random projections, feature quantization, and Gray encoding. You can use any of these methods on the feature vectors that come from EEG signals to make templates that are safe and can be changed.

Biohashing: Biohashing is a well-known cancelable biometric method that uses random data from the user and biometric features to make a secure hash code. Biohashing was first used for fingerprints and faces. It can be thought of as an extension of random projection with quantization [18, 87]. In a standard biohashing procedure (refer to Figure 2.3.1), a biometric feature vector x (for instance, a vector of EEG characteristics like power spectral densities or autoregressive model coefficients) is mapped into a novel space utilizing a user-specific random matrix or vector. A collection of n independent random basis vectors, commonly referred to as Tokenized Random Numbers in scholarly discourse, is produced from a confidential user-specific seed or key. We find the inner products of the feature vector and these random vectors, which gives us n scalar values. Then, these numbers go through a quantization function, like a thresholding operator, to make a binary string of length n, which is often called the BioHash or biometric hash. In simple terms, each part of the transformed template is either 1 or 0, depending on whether x is above or below a certain threshold when compared to a random vector. A zero-threshold with a sign function is often used, where the sign of each inner product determines the bit (1 for positive and 0 for negative) [18]. The binary code that comes out is saved as the cancelable template. For authentication, the same random projection that is unique to each user is applied to the query EEG feature vector to get a query code [88]. This code is then compared to the enrolled code using Hamming distance or another binary similarity measure. The random projection is different for each user (or application), so even if two users have similar feature x, their hashed codes will be completely different unless they use the same secret. This is two-

factor security because the system needs both the right biometric input and the secret projection key to work. If the stored template or secret key is compromised, a new random projection (a new seed or token) can be chosen to cancel the old template and replace it with a new code from the same biometric data that is not related to the old one. Biohashing is efficient in terms of computation and has very low false accept rates when the token is kept secret [18]. It is basically impossible to reverse because the quantization throws away continuous information and the random mapping messes up the feature space. It is not possible for an attacker to get back the original EEG features with just the binary code. However, biohashing's security can be greatly compromised if the secret token is taken. In this case, the transformation is made public, and an impostor could use the stolen token to try to reproduce the hash by projecting their own EEG features (this is called the stolen-token or pre-image attack scenario). To address this issue, the literature suggests using longer hash codes (more bits) and more complicated projection generation methods, like multi-layer or keyed permutations. Biohashing has not been extensively utilized in EEG analysis; however, theoretically, EEG feature vectors could be biohashed using a user-specific random key. For instance, He et al.'s [21] hashing of EEG autoregressive features can be regarded as a variant of biohashing, utilizing a random projection from a Johnson-Lindenstrauss random matrix. Biohashing for EEG is appealing because it could let EEG templates be stored as random-looking binary codes that don't give away anything about the user's real brain signals, but still let accurate matching happen (if the EEG features are stable). Building a strong biohashing pipeline for EEG may mean carefully picking features that can be used in different sessions and choosing the right quantization so that the binary codes aren't too sensitive to EEG's natural changes.

Random Projection: Random projections (RP) are an essential component in numerous cancelable biometrics frameworks, including biohashing. The Johnson-Lindenstrauss lemma says that a random linear mapping can project a set of points in high-dimensional space into a lower-dimensional space while keeping the pairwise distances about the same [88]. In biometric terms, a random projection is usually done by multiplying the feature vector by a random matrix M of size d x m (where d is the original feature dimension and m<d is the new dimension). The entries of M can be produced from a predetermined distribution (e.g., Gaussian), but they are typically concealed as part of the transformation key. This creates a transformed feature vector y = x/y in R^m. This mapping is many-to-one when m is much smaller than d: A lot of different x will map to the same y, which makes it impossible to go back (information is lost because the dimensions are reduced). A randomly chosen projection (especially when combined with other steps like shuffling or thresholding) can be non-invertible even if m=d unless the exact projection matrix is known. Two things make random projection a safe way to cancel out a transform:

(1) Dimensionality reduction: An attacker can't perfectly reconstruct x from y because important details of the original feature are lost (for example, projecting a 100-dimensional EEG feature down to 20 dimensions loses a lot of information) [18]; and

(2) Secret randomness: The relationship between x and y looks random without knowing the random matrix M. Even if an attacker got M,the dimensionality reduction means that x can only be recovered up to the lost information, which makes it an ill-posed inverse problem. Random projections have been effectively employed in fingerprint and iris cancelable templates, such as Gaussian random projection for

fingerprints and sectored random projections for iris. In the context of EEG, random projection can be utilized for feature vectors, including EEG power spectral coefficients or alternative embeddings. He et al. [21] is one example of this kind of work. They used a fast random projection (using a Johnson-Lindenstrauss hashing approach) on EEG features to make a fixed-length representation. Random projection is appealing for EEG analysis due to the high-dimensional nature of EEG feature spaces (numerous channels, frequencies, and time windows). Employing random projection can effectively compress the template while enhancing security. Also, if the projection matrix is linked to a key that is unique to each user, it makes sure that templates are different for each application. You have to be careful when choosing m: If you reduce too much, you might lose important information that helps you tell one thing from another. If you don't reduce enough, the template might be more likely to be attacked by inversion. Some studies (e.g., Kaur and Khanna [16] – note: reference index hypothetical here) have specifically analyzed how Gaussian random projection can serve as a non-invertible transform by evaluating if an attacker can solve for x given y and M results indicate that when $m \ll d$, the original biometric cannot be reliably recovered, confirming non-invertibility in practice. Another benefit of random projection is that it is computationally efficient: multiplying by a random matrix can be made more efficient and does not need complicated math, which is useful for real-time EEG authentication.

Quantization: A quantization step is often used to make a discrete template after a projection or other continuous transformation. Quantization is the process of changing real-valued transformed features into a small number of values, usually binary values like {0,1} [88]. In cancelable biometrics, quantization has two uses: it makes matching faster and more resistant to noise, and it makes non-invertibility even better by causing information loss. Binarization is a simple way to quantize, like taking the sign of each projected feature component (positive vs. negative) as in biohashing [18]. One could also split the range of a feature into several intervals and give each one a bit string or symbol (multi-level quantization). Quantization makes it so that many different feature values map to the same symbol, which is why it can't be reversed: Small changes to the original feature that don't cross a quantization threshold don't change the discrete output. Even big changes are unclear because the mapping groups a range of values into one code. From a security point of view, even if an attacker knows the quantization scheme, they can't get the exact continuous feature values from the output bits; they can only tell that they are in certain ranges. Quantization can make systems more tolerant of noise from a performance point of view. For example, if an EEG feature changes slightly because of noise or day-to-day changes, the encoded value stays the same as long as it stays in the same quantization bin (which reduces false rejections due to small changes in the feature). But quantization has to be set up very carefully to match the feature distribution. If the thresholds are set up wrong or the feature changes too much, a small amount of noise could cause a bit flip (for example, the feature crossing the threshold boundary), which would cause mismatches. To lower this risk, some systems use adaptive thresholds or multiple-bit encoding for each feature. Quantization is very important for EEG because EEG features often have built-in variability. To deal with signal noise, strong EEG hashing schemes may use techniques like rounding to multiple stable states or error-correcting codes along with quantization. The NeuroKey approach [22] is an example from the literature. It binarizes certain EEG features and uses error-correcting coding to make a stable key. NeuroKey is more of a biometric cryptosystem than a pure cancelable transform, but it shows how important quantization is for dealing with EEG variability. To sum up, quantization turns the continuous EEG feature space into a discrete space. This not only protects the template by losing information,

but it also makes it easier to compare templates (bitwise comparisons are quick and can use efficient distance metrics like Hamming distance).

Gray Encoding: This is a specific way to encode binary data that makes it more stable during the quantization stage. Also known as reflected binary codes, gray codes are binary sequences that alter by one bit for each value. This characteristic is beneficial for biometric quantization since it makes quantization errors less noticeable. It's possible that the usual binary encoding of the bin index could be altered in many bits if a feature value changes slightly and goes into a neighboring quantization bin. This could flip a number of the template bits. Using a Gray-coded assignment for the bin indices, on the other hand, makes sure that switching to a nearby bin only alters one bit of the code. For example, in a 2-bit binary code, the natural binary sequence 00, 01, 10, 11 has the first and second values differ by one bit, but the Gray code sequence 00, 01, 11, 10 makes sure that each transition changes one bit. This means that in biometric applications, if an EEG characteristic is just on the edge of a quantization threshold and sometimes falls in bin k and other times in bin k+1, the two alternative codes will be quite similar (they will only differ by one bit). This means that a total mismatch is less likely to happen. Gray encoding doesn't make things safer on their own (it's just a means to modify the names of the bins), but it retains the structure of similarities even when little features change. This can make comparing cancelable templates much more accurate. Some cancelable iris and fingerprint systems have used Gray code quantization to make the binary template more robust. Zuo et al. [20], for instance, invented new kinds of cancelable iris biometrics, namely GRAY-Salt and GRAY-Combo. These used gray-coded quantization and added random salt patterns to the features. Using Gray coding in those systems made it more likely that people would actually accept them because there were less bit mistakes between impressions of the same eye [20]. Gray encoding might work in an EEG system once the EEG features have been projected. For instance, if each feature is broken down into four levels (2 bits), those 2-bit values can be Gray-coded instead of binary-coded. If noise causes an EEG characteristic to decrease from level 3 to level 2 during one authentication attempt, the encoded 2-bit strings can go from (say) 10 to 11 (one bit difference) instead of from 11 to 10 (which in binary would be two bits different). Gray encoding makes the template less sensitive to tiny variations in the EEG data, which makes the matching findings more stable. Using more than one bit to encode feature components is an excellent method to make cancelable EEG templates more dependable. In Figure 2.3.1, the depiction of a biohashing pipeline might show Gray encoding as an extra step in the quantization block. This would show how projected features turn into binary code. Gray code is one of numerous design choices (along with error-correcting codes, weighting bits by reliability, and others) that can make cancelable biometric templates stronger without making them less safe.

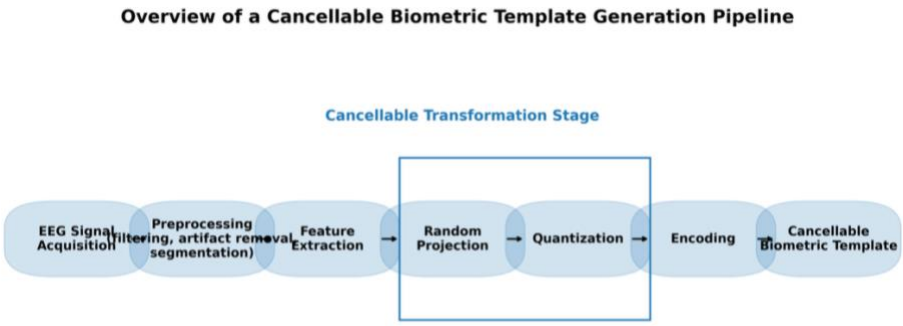


Figure 2.3 2 Overview of a typical cancelable biometric template genetation pipeline, combining random projection, quantization and encoding

2.3.3 Benefits and Challenges

Cancelable biometric systems provide substantial advantages for security and user privacy, particularly pertinent to EEG biometrics; however, they encounter unique challenges in implementation and performance. This part talks about the pros and cons of using cancelable (non-invertible) transformations for EEG templates [90, 91].

Benefits

First and foremost, EEG templates that can be canceled greatly improve security and privacy. Because the stored template is a one-way transformed version of the EEG features, it is impossible for an attacker to get the user's raw EEG signals or any other private information from a stolen template. This eases worries about privacy with EEG biometrics, which could show medical or mental health information from raw brainwave data. In other words, if an enemy gets into the EEG database, they only get random codes (see Figure 2.3.2) that don't give them any information about the user. Another benefit is revocability: if a user's cancelable EEG template is ever compromised, the system operator can revoke it and update it by applying a new transformation (for example, using a new random projection key) to the original EEG features. The user doesn't need to have another EEG recording; the existing biometric can be safely "re-enrolled." This is like changing a password, and it's a big advantage over traditional biometric systems, where a stolen template (fingerprint, face image, or EEG data) would have permanent effects. Cancelable biometrics also support diversity and unlinkability. For example, the user can have different EEG-based templates for different services by using different transformation parameters for each service. This means that if the database of one service is leaked, an attacker won't be able to connect that information

to the user's template in another service. The templates will look like they don't belong together. Privacy-by-design features can help users accept EEG biometrics from a compliance and trust standpoint because both users and regulators can be sure that the biometric data is not stored or shared in its raw form. Another benefit is that cancelable transformations can sometimes help with noise filtering or feature normalization. Quantizing and projecting EEG signals may help reduce some high-frequency noise or random changes that don't matter (for example, quantization ignores small differences as long as they are in the same bin). This can make the template more stable when there are small differences in the recording. Lastly, binary templates make matching and storing data easier. Binary codes from methods like biohashing are small and can be quickly matched using bit operations, which is useful for large-scale deployments or devices with limited resources. In short, cancelable EEG biometrics let you use brainwave-based identification to get the benefits of authentication while reducing the risks of storing brain data.

Challenges: Even though cancelable biometrics for EEG systems have a lot of good points, they are not straightforward to use [93]. One important problem is making sure that recognition performance stays the same during the adjustment. EEG signals are naturally noisy and might change from one session to the next depending on how the person is feeling and where they are. If a transformation is too extreme, like an extremely random projection or coarse quantization, it could make the distinctions between users more obvious or lose important information that could help tell them apart. This would make it harder to verify someone's identity. While still keeping the classes separate, it's hard to make the transform non-invertible. There is usually a trade-off between security and accuracy. Making the transform more complex and secure (by throwing away more information and mixing features more randomly) can lower the true accept rate, while changing it to keep accuracy might leave some information that attackers could use. Finding the right balance is not easy. This is especially hard for EEG because the features of EEG have a lower signal-to-noise ratio than, say, fingerprints. So, a poorly constructed transform could make it a lot harder to use EEG to identify users, which is already a hard thing to do. Another problem is making sure that stolen tokens or parameters don't make the system crash. You should think about what would happen if someone took the key that the cancelable system uses. For instance, biohashing and random projection methods need a key that only one person may use. For instance, this might happen with EEG if the random projection matrix or seed is stored next to the template in an unsafe way. When a token is taken, many cancelable systems are

vulnerable to pre-image attacks. If the attacker knows the transform, they can try to find an input that makes the stolen template. Researchers have found that hill climbing or brute force approaches can be used to get close to biohash codes, even if the genuine user's key isn't accessible. This is only true if the transformation isn't too hard or the feature space isn't too huge. To stop these kinds of attacks, you might need to take further steps, including using longer codes or mixing multiple transforms. From a system point of view, integration and user concerns are problems: The user must already be wearing sensors and sending signals for EEG-based authentication to operate. Adding a transform that can be aborted doesn't make things harder for the user, but it does make the system more complicated. The system needs to keep transformation keys safe for each user and program (to make sure there is variety), which makes it tougher to share and keep keys. If the user has to remember or carry a token, either a password or a hardware token that gives the transformation key, this is a two-factor system. This makes it hard to utilize because people could lose or forget the token, which is the whole idea of biometrics. There are no standard datasets or benchmarks for cancelable EEG biometrics, which is another problem with the technology. Researchers need to find out how to test how well an EEG cancelable template can stand up to different kinds of attacks, like inversion attacks and cross-matching attacks, while still being right. The computational load is another item to consider about. Some transformation methods, like big matrix multiplications or iterative graph-based transforms, could be too slow for real-time authentication. On the other hand, random projections are usually speedy. Last but not least, it's hard to make sure that classification methods work with EEG biometric systems because many of them utilize feature vectors with machine learning classifiers. When you add a cancelable transform, you have to change the data that is being matched. You might have to alter the classifiers or go from model-based classification to distance-based matching. Before using binary or randomly projected data, some classification systems, including neural networks, may need to be retrained. If you add cancelable templates, you might need to utilize a completely new way to recognize them, such as cryptographic comparators or binarized matching algorithms. This could be a problem if you wish to add template protection to current EEG biometric methods.

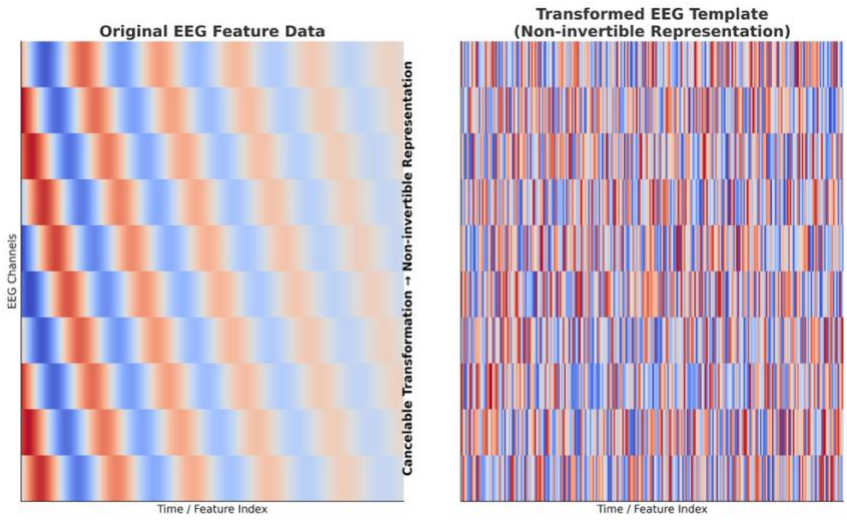


Figure 2.3 3 A comparison illustrating how a secure transformed EEG template (right) reveals no identifiable information compared to the original EEG feature data (left), highlighting the non-invertibility and privacy enhancement of the cancelable approach

2.4 Cancellable biometrics

Biometric identifiers are intrinsically linked to the user and are not easily alterable if breached. Cancelable biometrics is a concept that was created to lower this risk by intentionally and repeatedly changing biometric features before they are stored or compared. In a cancelable biometric system, the raw EEG-based identifiers are changed into a protected template using a one-way (non-invertible) function.

If a stored template is stolen or compromised, it can be revoked and replaced by making a new template from the same biometric just by using a different transformation key. (§2.4.1) talks about the basic ideas behind cancelable biometrics. (§2.4.2) analyzes some common ways to use cancelable transformations and the positives and negatives of using these methods, especially in EEG-based authentication (§2.4.3). Finally in (§2.4.4) we discuss about specific design considerations and applications of cancellable biometrics in EEG systems.

2.4.1 Definition and Concepts

Cancelable biometrics (also called biometric template transformation) refers to the practice of applying a deliberate, repeatable distortion or transformation to biometric data in order to protect sensitive user information. The transformation is typically parameterized by a secret key or value (sometimes user-specific), ensuring that only those with the correct key can generate or recognize the transformed template [94]. Several key properties and requirements define cancelable biometric systems, as first articulated by Ratha et al. [1] and subsequent works:

- *Irreversibility (Non-invertibility):* The transformation must be one-way, meaning it is computationally infeasible to recover the original biometric from the transformed template. In practice, this implies that even if an attacker obtains the stored EEG template, they cannot reconstruct the original EEG signals or features. Irreversibility protects the user's privacy (e.g., preventing leakage of raw EEG data that might contain private health information) and ensures the biometric data cannot be abused elsewhere. Formally, the transform should non-injectively map the feature space onto the protected template space, often losing information (through quantization or projection) such that inversion is essentially impossible.
- *Revocability (Renewability):* It must be straightforward to revoke and reissue a biometric template if it is compromised. This property makes biometrics function analogous to passwords – a user can change their biometric template by applying a new secret transformation. In practice, revocability is achieved by using a key or salt in the transformation; by altering or refreshing this key, one can derive a new, independent template from the same EEG data. Should an attacker obtain a transformed EEG template, the system owner can cancel that template and generate a fresh one using a different key (much as a stolen credit card is canceled and replaced). Revocability is crucial because users cannot change their underlying EEG signals, but they can change the transformation applied to those signals.
- *Diversity (Unlinkability):* A well-designed cancelable biometric system ensures that transformed templates differ across applications and cannot be linked to the same user. In other words, using the same biometric (e.g. an EEG profile) with different transformation parameters will yield entirely different templates. This diversity property means that even if a user uses EEG authentication for multiple services, each service can store a unique template such that a breach in one service does not jeopardize the others. It also prevents cross-matching: an impostor who obtains a cancelable template from System A cannot use it to search for the user in System B's database unless they also know the transformation key for System B. Diversity is typically achieved by per-application keys or salts. It is closely related to unlinkability, meaning that given two transformed templates, an attacker cannot determine whether they originated from the same biometric source without the keys.
- *Performance Preservation:* Although not always listed as a formal criterion, practical cancelable biometrics require that recognition accuracy and utility

remain high after transformation. The distortion should not significantly degrade the discriminatory information in the biometric data [77]. In an EEG context, this means the transformed EEG features should still allow reliable matching between enrollment and query data for the genuine user, while maintaining low false match rates for impostors. There is often a trade-off in designing the transformation: a more aggressive distortion (for better security/irreversibility) can reduce matching accuracy, whereas a mild transformation preserves accuracy but may be less secure. An ideal cancelable transform optimally balances this trade-off, yielding distance preservation to some extent (so that genuine vs impostor comparisons in the transformed domain have similar separability as in the original domain) without being invertible. Figure 2.4.1 illustrates this concept, showing notional Gaussian-like distributions of genuine and impostor matching scores before and after a cancelable transform. In a well-designed scheme, the distributions remain sufficiently separated (maintaining accuracy), even though the underlying data has been randomized to protect security.

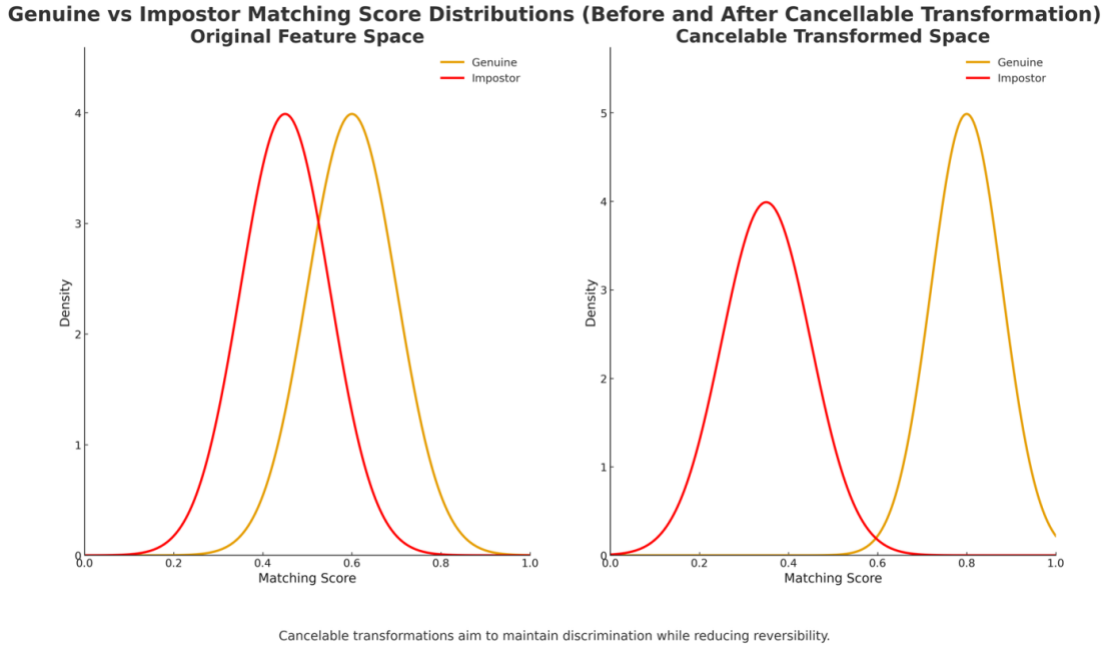


Figure 2.4 1 The effect of a cancelable transformation on matching score distributions

2.4.2 Techniques of Cancellable Transformations

To make cancelable EEG biometrics work, you need to create transformation functions that meet the above requirements. The literature has examined various techniques to

attain non-invertible and revocable mappings. Common methods include salting, random projections, biohashing, and quantization steps, which are often used together. Figure 2.4.2 shows a diagram of a typical cancelable template generation pipeline, with these parts highlighted (as they are used in biohashing methods, for example). We will now explain these methods and how they help make things cancelable:

Biometric Salting: In biometrics, salting means adding user-specific extra data (like a secret key, token, or password) to the process of making a template. The concept is similar to password salting: by adding a secret random value to the feature transformation, the template becomes unique to each user and changes. For EEG, a salt could be a vector or matrix that is unique to each user (or session). The change could be an arithmetic operation (like adding a random vector to the feature vector) or a structural change (like shuffling feature elements according to a secret permutation). Salting makes sure that there is variety (different salts make different templates) and makes it possible to revoke (changing the salt revokes the old template). Salting by itself might be a simple linear operation, so it is often used with more complicated transforms to make sure it can't be undone.

Random Projection: This is a way to move high-dimensional feature data to a new space using a randomly generated matrix. In a cancelable biometric context, the projection matrix (usually taken from a random distribution like Gaussian) serves as a key. It is either kept secret or linked to the user's credentials. To get a projected vector in a lower-dimensional k-space, you multiply each EEG feature vector x (of size d) by a random matrix R (of size d x k). This does two things:

It reduces dimensionality, which means it loses information and makes it harder to reverse (because the inverse problem is underdetermined if k<d); and

(2) It effectively scrambles the feature in a way that can only be done with the same random matrix. Random projections keep distances between points close to the same (Johnson-Lindenstrauss lemma), which helps recognition performance stay the same even after the transformation. But if an attacker somehow got the exact projection matrix, the scheme might be partially invertible (especially if k is close to d). So, the matrix must stay secret, or the projection must be combined with more one-way steps. Random projection has been extensively utilized in cancelable biometrics; for example, Ratha et al. delineated straightforward geometric transformations for fingerprints (e.g., Cartesian and polar transforms) that can be regarded as domain-specific projection/scrambling functions [17]. In EEG systems, one may project extracted feature vectors (e.g., power spectra or wavelet coefficients) onto a private random subspace to create the safeguarded template.

Biohashing (Random Multispace Quantization): Teoh et al. (2006) [4] came up with biohashing, which is a well-known cancelable biometric method. It can be seen as a particular case of salting and random projection, along with a quantization step. In biohashing, the user's biometric feature vector is enhanced with a user-specific secret (tokenized random number sequence) and subsequently projected into a random orthonormal space. Then, the result is quantized, which usually means taking the sign of each component to make a binary bit-string (hence the name "hash"). Teoh et al.'s Random Multispace Quantization (RMQ) scheme showed that these BioHashes can be canceled (if they are compromised, a new random seed or token can make a different hash) and that they cannot be inverted (the binary output and unknown projection stop the original features from being recovered). Biohashing is basically a two-factor method. It uses something the user is (biometric) and something the user has or knows (the random key). When used with EEG, a biohashing-like method would mean that each

user's EEG feature (like a vector of bandpower values) is multiplied by a random matrix that is unique to that user and then turned into a unique code. Figure 2.4.2 shows the biohashing pipeline: First, raw EEG features are pre-processed. Then, a random matrix (user-specific token) projects them, and finally, a stable hashing method (like modulus or sign quantization) makes the protected template. This pipeline makes sure that the stored reference will only match the hashed template of the real user if the right token is used. Biohashing has demonstrated the ability to maintain superior recognition performance while fulfilling cancelable properties; for instance, Teoh et al. indicated that authentic BioHash codes can be identified with minimal accuracy loss relative to original features, attributable to the robust quantization [4].

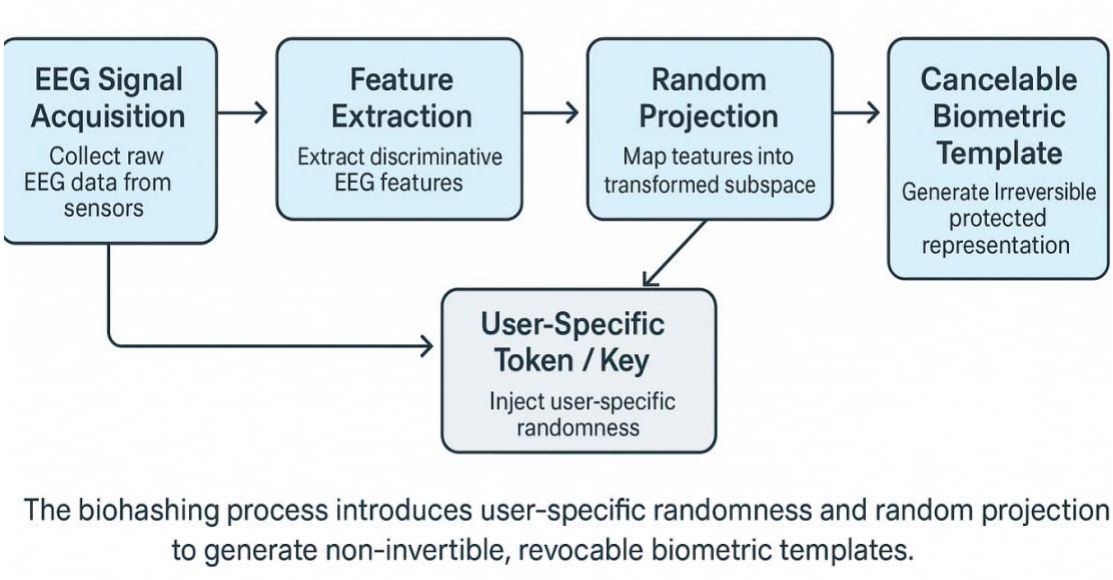


Figure 2.4 2 Biohashing-Based Cancelable Template Generation Process

- *Quantization and Non-linear Mapping:* Quantization is the process of turning continuous biometric features into separate values, usually binary values. Quantization can make non-invertibility worse on its own because it loses fine-grained information (for example, small changes in EEG feature value are turned into a 0 or 1 bit). A lot of cancelable schemes have a quantization or hashing step after the linear transform is done. For example, if an EEG feature is just above a certain level, it might be mapped to "1," and if it isn't, it might be mapped to "0." An attacker with a quantized template only knows the general categories of the original feature, not the exact value. Quantization also makes systems more tolerant of noise and differences within classes because binary codes can be matched with error-correcting metrics (like Hamming distance) that naturally allow for some variation. Some advanced methods use non-linear mappings (like folding transformations and polynomial kernels) to make the feature space even more distorted in a way that can't be undone. Ratha's surface folding transform for fingerprints (a method that

folds the image in a complicated way) and the use of polynomial functions on features are both examples of ways to make the mapping one-way [1]. In EEG-based transformations, a non-linear function (like taking the sine or cosine of feature values, as seen in some recent EEG cancelable schemes) could be used to make sure that the relationship to the original features is very non-linear, which makes it hard to try to reverse it. The combination of randomization and quantization is very strong. Even if someone intercepts the binary template, it is very hard to guess the original analog EEG features because of this many-to-one mapping.

You can use any of these methods on their own or together. In practice, hybrid transformations are common. For instance, an EEG template scheme might first use a random projection and then a quantization with a secret threshold vector (mixing two levels of protection). The type of transformation you choose depends on the EEG features and how much security and usability you want. Designers need to make sure that real EEG variability (like noise or differences in recording) doesn't lead to too many mistakes after transformation, while still effectively hiding the original data.

2.4.3 Benefits and Challenges

Canceling biometrics has many advantages, especially for EEG authentication, where privacy and security are very important:

- *Better Privacy:* The system keeps users' sensitive neurological data safe by only keeping transformed EEG templates, which can't be changed back to the original signals. EEG signals can show things like your cognitive state, health problems, or emotional state. You can cancel a template if you think it might be used to steal private information. This is becoming more important because data protection laws like the GDPR treat raw biometrics as private information [95]. Using cancelable biometrics, even if a database is hacked, the attacker can't use the data to learn about the person or pretend to be them somewhere else, which keeps the user's privacy safe.

- *Security and Resistance to Attacks:* Cancelable systems are built to be strong against a wide range of attack types. A big benefit is that it is hard to steal templates. This is because templates can be revoked and replaced, which solves the problem of biometrics being permanent. Also, a well-designed cancelable transform can stop

hill-climbing attacks and record multiplicity attacks. In a hill-climbing attack, an enemy tries to guess inputs over and over until they match a stored template. Cancelable templates that use many-to-one mappings (for example, when different feature inputs lead to similar transformed outputs) can easily stop these kinds of attacks. Wang et al. (2023) note that a many-to-one non-invertible mapping in EEG templates is "inherently resistant to hill-climbing attacks," as any solution discovered by the attacker would be rendered ineffective once the system revokes and substitutes the template. Similarly, brute-force attacks are unfeasible when the transformed space is extensive and the actual mapping remains confidential; for instance, a study estimated that deducing a secured EEG feature through brute force would necessitate exploring an extraordinarily vast space (approximately $\sqrt{2^{8000}}$ possibilities). Cancelable biometrics also stop cross-matching attacks. This is because each template is application-specific, which means that an attacker can't link templates from different databases to follow users. Overall, the security that cancelable designs offer is much better than that of systems that store raw EEG features.

- *Revocation and Recovery:* If a breach or compromise is suspected, cancelable biometrics make it easy to fix the problem quickly. You can cancel the compromised EEG template and get a new one by just changing the transformation parameters (for example, by using a new random seed). This makes it easy to get back to a secure state, just like resetting a password. This isn't true for traditional biometric systems. If a fingerprint or EEG template is compromised in a system that can't be canceled, the user's biometric is "burned" for all future uses. The cancelable design solves this by letting you re-enroll with the same biometric data but a new identity in the system. This benefit makes biometric authentication much more secure in the long run, making it a better option for users.

- *User Control and Consent:* Users can have more control over their biometric identity with cancelable schemes. Users can give their biometric data to a service by giving them a secret key (like a PIN or token). They can also take it back if they want to. For example, a user might have a smartcard with a key that lets them change their EEG features. If they lose the card or think it was used for something else, they can ask the service to make that key invalid and use a new one. This idea fits with privacy rules, since it lets the biometric be non-linkable and pseudonymous in different situations.

Despite these advantages, there are significant challenges and trade-offs in deploying cancelable biometrics for EEG:

- *Recognition Performance Trade-off:* One of the biggest problems is making sure that the transformation doesn't make the biometric recognizability worse than it already is. EEG signals already have a lot of noise and variability [96]. If a transformation isn't carefully planned, it could make the variability between subjects bigger or the separability between subjects smaller. If the transform is too harsh, it can be hard to get low false rejection and false acceptance rates in the transformed domain. To find the right balance between security and accuracy, designers often have to tweak the system (for example, by changing the quantization thresholds or choosing the projection dimensions). A slightly distorted template may work better for matching, but it is less secure. A heavily randomized template, on the other hand, is more secure but causes more false rejections. Finding the sweet spot is not easy and depends on the application. Recent studies in EEG cancelable biometrics indicate that advanced transformations can achieve recognition performance comparable to the utilization of original features. Wang et al.'s privacy-preserving EEG authentication system, for instance, achieved an Equal Error Rate (EER) as low as 0.5–1% using transformed templates, which is similar to the case without protection. This means that the performance trade-off can be handled with the right tools, but it is still an important part of the design.
- *Safety of Transformation Keys:* The safety of a cancelable biometric system often depends on how safe the keys or parameters used in the transformation are. If an attacker gets the secret salt or random projection matrix, they might be able to guess or reverse the transformed template (or at least use dictionary attacks by transforming a lot of possible biometrics with the known key). So, managing keys is very important. Keys can be safely stored in hardware (like a secure element or user token) or made from user passwords. In a multi-factor setup (biometric + token), the system has to protect both factors. Also, the transformation's entropy should be high, meaning that the space of possible keys or parameter values should be big enough to stop brute-force key search. If the transform isn't complicated enough (with not enough random bits), it could be broken by trial and error. A big operational challenge is making sure that

each user and application has strong, unique keys and updating them every so often.

- *Utility vs. Irreversibility:* To protect the original EEG, it is important that it not be possible to invert it. However, an overly aggressive transformation that completely scrambles the data could be too irreversible, meaning that even legitimate matching becomes unreliable. There is also the more subtle issue of keeping the distance: many cancelable transforms try to keep the distance or similarity structure of the original feature space (so that the nearest neighbors stay the same after the transform). If a transform changes distance relationships a lot (for example, because of a very non-linear or discontinuous mapping), the matching algorithm might have a hard time telling the difference between real and fake inputs. On the other hand, if a transform keeps distances too well (like a simple rotation of feature space), it might be able to be inverted or be open to record multiplicity attacks that take advantage of distance consistency. This tension necessitates meticulous design; for example, incorporating a degree of noise or projection that disrupts invertibility, yet insufficient to obliterate the benefit of the genuine-impostor distance gap. Some academic works formalize this equilibrium by examining the information leakage of the transform and its effect on error rates.

- *Compatibility and Integration:* Adding cancelable templates to an EEG biometric system means changing how people sign up and confirm their identity. The transformation module must be connected to the feature extraction module, and the matching algorithm must work in the transformed domain. This can make it harder to design systems, especially if old systems or algorithms are in use that expect certain kinds of features. For example, you can't use a classifier trained on raw EEG features on hashed EEG features without retraining it. There might also be extra processing time needed for some transformations, like large random projections or iterative non-linear maps. This could be a problem for real-time authentication. It is hard to make sure that cancelable biometric schemes work well with high-dimensional EEG data in terms of speed and memory.

- *User Acceptance:* User acceptance of cancelable biometrics can be mixed, and this is not unique to EEG. Users like the added security and the idea that their

brainwave data is safe. But if a system needs the user to remember a PIN or manage a token (for salting), it adds an extra step that could be annoying for them. Adding a user-controlled key may not always be possible in purely passive EEG authentication, where the user just wears a device. In this case, the system must securely manage keys on its own. It is important to communicate the benefits of template protection to users and make sure it is easy to use.

To sum up, cancelable EEG biometrics offer password-like protection for brain data. If a template is breached, it is not a disaster but something that can be fixed. The method makes security and privacy better, but it needs to be designed carefully to keep the reliability that biometric systems need. Scientists are still working on each of the problems listed, improving transformation algorithms and testing them against new attack tactics.

2.4.4 Application in EEG Biometrics

The application of cancelable biometrics to EEG-based authentication is an emerging area driven by the need to safeguard brainwave data. EEG biometrics present unique opportunities and challenges for cancelable template design:

- *Motivation in EEG Context:* EEG signals contain a lot of personal information and are very private. As previously mentioned, research has demonstrated that even abstract EEG biometric templates might unintentionally convey information like an individual's age, gender, or health status [97, 98]. This makes using EEG for mass authentication a big privacy problem. By making sure that the saved template is a privacy-preserving abstraction of the EEG and ideally only showing an identification, cancelable biometrics directly address this issue. Additionally, because EEG data collection is typically intrusive or inconvenient (requiring a headset, etc.), users are likely to only agree to EEG authentication if they are satisfied that their brain data will not be misused or leaked once it is collected. The cancelable technique is a fundamental factor that makes EEG biometrics work outside of the lab. It gives users and service providers both confidence that brain-data breaches may be stopped or lessened.
- *Implementations and Techniques:* When EEG uses cancelable templates, it's usually at the feature level. After EEG signals are acquired and pre-processed (filtering, artifact removal, feature extraction such as power spectral densities or neural network embeddings), those features are changed in some way. A lot of the methods talked

about in §2.4.2 have been changed. For example, scientists have used biohashing-like approaches on EEG by projecting EEG feature vectors with a random vector that is unique to each user and then turning the result into an EEG hash. This can create a stable binary code for a person's EEG while also allowing for revocation. Other efforts provide more complicated changes that are specific to the properties of EEG signals. Wang et al. (2023) suggested a polynomial-based cancelable transform for EEG characteristics, integrated with graph-based representations of EEG channels. They use a polynomial and cosine transformation to embed features from the EEG connection graph (to make sure they can't be reversed), and they add a step to distort the template to make the outputs even more unpredictable. These kinds of advancements in specialized fields help keep EEG verification accurate even after a lot of distortion. Initial investigations indicate that these techniques can attain mistake rates below 1% in EEG authentication tasks while fulfilling cancelable features, highlighting the technological feasibility of cancelable EEG biometrics [98].

It is worth noting that EEG biometrics often involve multi-session variability – a person's EEG signals can vary across different days or mental states. Cancelable transformations must accommodate this, for example by focusing on robust EEG features (like subject-specific frequency responses) and ensuring the transformation does not amplify session noise. Some systems use multiple EEG paradigms (tasks or stimuli) and fuse the results to increase reliability. Cancelable templates can be generated for each paradigm and either stored separately or combined, offering diversity and possibly greater security (an attacker would need to break multiple transformed templates).

- *Effects and Future Directions: Effects and Future Directions:* The addition of cancelable templates to EEG biometrics makes a big difference: it brings EEG authentication closer to being able to be used in real life. Privacy worries have been a big problem for brainwave biometrics. With cancelable designs, companies can tell users that if their data is taken, it will be useless to their enemies. This is in line with ethical use of neurodata and could lead to more use of neurodata in fields like healthcare login systems, brain-computer interface access control, or high-security identity verification. We are starting to see prototypes and frameworks that mix EEG with cancelable identity management. For example, utilizing EEG to produce cryptographic keys in a way that can be canceled or adding EEG authentication to multi-factor systems using tokenized cancelable schemes.

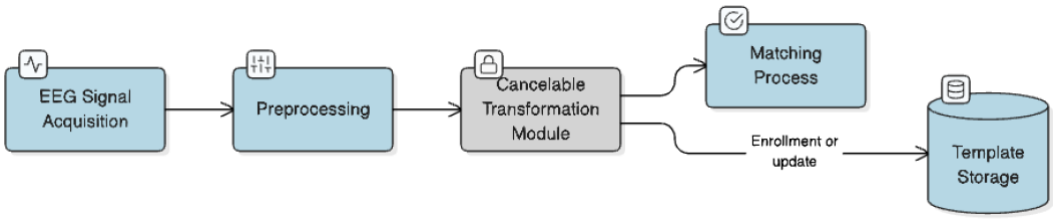


Figure 2.4 3 EEG-Based Authentication System with Cancellable Template Protection

In the future, design standards for cancelable biometrics for other types of biometrics (such fingerprints and faces) are being improved for EEG. Researchers are examining the theoretical security of EEG transforms (e.g., information leakage analysis, as referenced in §2.4.3) and the methodology for explicitly demonstrating irreversibility within specific attack scenarios. They are also looking at how to make EEG systems even more secure by combining cancelable biometrics with cryptographic approaches (which are sometimes called biometric cryptosystems hybrids). Another new usage is employing cancelable EEG IDs for continuous authentication, which means watching a person's brain signals all the time to make sure the same person is still there. In these situations, template security is necessary to avoid the constant risk of live EEG data leaking.

2.5 Cancellable EEG Biometric Systems in Literature

The concept of cancellable EEG biometrics-employing irreversible modifications on EEG features to facilitate the cancellation and replacement of stolen templates-has just recently been examined. Wang et al. (2022) completed the pioneering study that introduced the first cancellable EEG template design. They employ a deep-learning feature extractor to process EEG data, and then they alter them using a user-specific biohashing-inspired pipeline that includes random permutation, Hadamard product, and random projection. Lastly, they encode the signals in binary (Gray coding). This non-invertible transformation utilizes a unique secret key for each user to generate a secure binary template from the EEG features. The study demonstrated the feasibility of EEG template cancelability: the updated templates retained superior discriminability (reaching improved verification accuracy relative to earlier EEG biometric methods) while maintaining the integrity of the original signals. It also provided a case for security against other forms of assaults by suggesting that it could stand up against inversion, hill-climbing, and brute-force attempts. However, the analysis in this fundamental study

was limited in several ways. The authors primarily validated that a compromised template could be revoked by modifying the key; however, they did not comprehensively quantify revocability (e.g., the degree to which new templates diverge from old ones beyond a singular key alteration) or rigorously evaluate brute-force resistance (e.g., the difficulty for an attacker to ascertain the template or key through trial and error). Likewise, while the study assessed fundamental resilience to attacks (introducing the concept of "second attacks" following revocation), it did not definitively ascertain the level of security-issues that this thesis will comprehensively analyze.

Current Research on Cancellable EEG Biometrics: Besides the foundational research mentioned earlier, there have been only a few studies that have used cancellable biometric methods on EEG signals. These studies utilize diverse transformation techniques, including random projections, hashing, trigonometric transforms, and signal encryption, and differ in their assessment of security properties alongside recognition performance. Some important works that came before this one are:

- *Random Projection Hashing:* He et al. (2009) investigated the hashing of EEG features as a preliminary method for preserving privacy. They took multivariate autoregressive (mAR) coefficients from multi-channel EEG and used a random projection (Johnson–Lindenstrauss) based hash to make small EEG "hash vectors." This method used a simple probabilistic matcher to show that EEG-based authentication is possible even with features that have been transformed in one direction. The authors pointed out that direct hashing has a problem: cryptographic hash functions are not very strong for biometrics because even small changes in EEG can cause big changes in the hash. As a result, this method didn't have a way to cancel or fix mistakes; it protected the data but wasn't a real cancellable biometric template as it was later defined.

- *EEG Key Generation (Neuro-Key):* Bajwa and Dantu (2016) suggested a "NeuroKey" system that used EEG recordings to make cancelable biometric keys. The EEG signals were altered (with power spectral features from Fourier/wavelet transforms) and then turned into binary code and hashed to make a stable key. In their work, the word "cancellable" meant that if the old key was stolen, they could collect a new EEG under a different task to make a new key. Even though this method gives you replaceable keys, it doesn't change the same template; instead, it uses different cognitive tasks. This means that the raw EEG data could still be

exposed, and the template can only be changed if there are different tasks available (which also caused different biometric performance across tasks). In short, NeuroKey showed that EEG could be used to make keys, but it didn't come up with a general cancellable template solution that could constantly revoke and reissue protected templates from the same data.

- *Cryptographic Binding (Fuzzy Commitment):* Damaševičius et al. (2018) integrated EEG biometrics with cryptographic fuzzy commitment frameworks. They used fuzzy commitment (BCH codes) to turn EEG features into an error-correcting code word and link it to a random key. This way, a real user can reliably unlock the key with their EEG. This method makes sure that stored data doesn't give away too much about the raw EEG (to protect privacy) and that noise in EEG readings is okay. But in the strictest sense, it can't be "canceled." Once the biometric data is compromised, the fuzzy commitment doesn't let you get a new independent template from the same EEG. The link between the EEG and the cryptographic helper data stays the same. So, even though privacy is protected, a stolen template in this kind of cryptosystem can't be easily taken away and replaced without getting new EEG data.

- *EEG in Biometric Cryptosystems:* Maiorana et al. (2015) investigated EEG-based cognitive biometric cryptosystems, utilizing EEG signals alongside user-specific secrets. They showed an example of how to use EEG features in a cryptographic key generation framework, where "cognitive" means using tasks or responses from the brain. This improved security by combining biometrics with cryptographic keys, like other biometric cryptosystems, but it didn't focus on template transformations. The emphasis was on one-time key extraction instead of repeatedly revocable templates, which meant that the classic cancellable biometrics criteria (non-invertibility, revocability, unlinkability) were not fully met.

More recently, researchers have moved toward cancellable EEG templates that are based on true transformation and meet those standards. Wang et al. have released a series of subsequent studies to their 2022 research, examining various transformation methodologies:

- *Polynomial and Cosine Transform (PolyCosGraph):* Wang et al. unveiled PolyCosGraph in 2023 [100], a cancelable EEG biometric system utilizing a more intricate non-invertible transform. This method uses polynomial transformations

with cosine functions on EEG graph features, which are features that show how connected EEG channels are. The design creates a protected template by using a polynomial mapping modulated by cosine terms to change graph-based feature vectors. There was also a template corruption module that would purposely damage any compromised template before it was replaced. A custom matching algorithm was created to compare templates in this new domain. The PolyCosGraph method had very good biometric accuracy, with equal error rates (EER) of about 1.5% in resting-state EEG and 0.4–0.7% in some task-driven EEG scenarios. This is about the same as using original (non-protected) features. It was also very important to test security features: the authors showed that PolyCosGraph templates meet revocability (new templates can replace old ones) and provided analysis showing that they are strong against many types of attacks, such as record multiplicity (ARM), pre-image inversion attacks, hill-climbing attacks, "second attacks" after a revocation, and brute-force guessing attacks. This thorough evaluation was a big step forward because it clearly measured how well the system could handle different threats while still doing a great job of recognizing things.

- *Cancelable Deep Learning Framework:* Wang et al. (2024) discussed cancellability in relation to model-based EEG biometrics. Conventional methodologies, including those previously mentioned, utilize stored templates and template matching; in contrast, model-based techniques train a classifier, such as a neural network, on an individual's EEG data. Wang et al. demonstrated that trained models are susceptible to model inversion attacks, wherein an adversary who acquires the model can partially retrieve EEG-like inputs that generate the authentic user's outputs. To combat this, they suggested a cancellable deep learning framework that adds a transformation that can't be undone to the model training pipeline. A generative adversarial network (GAN) essentially learns to change raw EEG into a synthetic but discriminative form (like a random non-invertible function), and the user's classifier is trained on these changed EEG features. The transformation's parameters are like a secret key. If the model is hacked, the key can be changed and the model can be retrained on data that has been changed, which "revokes" the old model. This framework successfully extends cancellability to the model level, thwarting attackers from exploiting compromised EEG models. The authors reported that the cancellable model maintained or improved verification accuracy over state-of-the-art non-cancellable models, which is very impressive. This study emphasizes that cancellable biometric principles are applicable not only to stored templates but also

to learned models, a domain of increasing significance with the emergence of deep learning in biometrics.

- *Other Transformation Techniques:* Several additional studies have investigated alternative transformations for EEG template protection, albeit typically on a reduced scale. For instance, Salama et al. (2023) suggested a multimodal cancelable biometric scheme that merges EEG signals with images and subsequently employs optical encryption methods, particularly Double Random Phase Encoding (DRPE), on the integrated data. The output of the DRPE encryption acts as a cancelable template because you need the secret phase keys to decrypt it. This system, which used EEG and another biometric together, had an EER that was very low (close to 0%) with the cancelable templates. It shows that non-invertible transforms from signal processing, such as DRPE, which was first used for image encryption, can also be used to protect EEG biometrics. But these kinds of approaches see the problem more as an encryption task than a biometric template transform. They focus on how strong the encryption (and in this case multimodal fusion) is, and they usually only look at basic security issues when testing recognition performance. Still, they show how many different methods are being looked at to make cancelable EEG templates, from statistical transformations to cryptographic algorithms.

In short, the research on cancellable EEG biometrics is still in its early stages. Early research demonstrated that EEG signals harbor identifiable information and investigated basic hashing or encryption techniques for its protection; however, these methods did not provide comprehensive support for revocability and formal security assurances. Recent research, primarily conducted by Wang et al. and associates, has initiated the process of addressing this deficiency by developing specialized non-invertible transforms for EEG and assessing them against biometric performance and security standards. Table 2.1 (in Chapter 2) gives a summary of these studies that shows how they were different from each other. It shows the transformation techniques used (random projection, polynomial-cosine, etc.), the biometric accuracies achieved, and whether key security properties like revocability, unlinkability, and resistance to inversion/brute-force attacks were taken into account.

Questions for Open Research: Even with this progress, there are still some problems that need to be solved before cancellable biometrics can be used on EEG. Some of the most important questions are:

- *Revocability in Key Rotation:* How dependable is it to revoke and regenerate an EEG template with a new key, and how does this affect the accuracy of recognition? A user should be able to update their EEG template as many times as they want, such as if their secret key is lost. Previous research has demonstrated fundamental revocability (illustrating that one transformed template can be substituted for another); however, a more comprehensive analysis is required. For example, if we keep changing keys, do old and new templates stay truly uncorrelated (meeting the diversity criterion) even after several rotations? This thesis tackles this issue by empirically assessing revocability through successive key updates and measuring the similarity between old and new templates. Chapter 5 presents a pivotal rotation experiment wherein each user's EEG data is re-transformed using various keys. We assess the similarity between a user's successive templates in relation to the similarity with other users' templates, thereby validating that authentic templates can be updated without linkability.

- *Impostor vs. Real Similarity (Unlinkability):* A basic security requirement is that transformed templates should not allow linking between (a) different users (diversity) or (b) different templates of the same user that were made with different keys (revocability). In practice, this means that a user's template with one key should look like an impostor to their template with another key. Some studies assessed this through pseudo-impostor distributions or linkability metrics; however, unresolved inquiries persist regarding the optimal quantification and assurance of unlinkability for EEG. This thesis provides an analysis of template correlation by treating binary EEG templates as high-dimensional vectors and calculating correlations (and Hamming distances) between genuine and impostor template pairs across various transformations. The findings in Chapter 5 validate that authentic-authentic template correlations diminish to impostor levels upon key alteration, thereby effectively decorrelating a user's templates and reinforcing the unlinkability of the proposed methodologies.

Attacks that use brute force and guessing: How safe are cancellable EEG templates from someone who tries to guess or fake them? This is like trying to guess a password by brute force, but biometric templates are in a continuous feature space. An attacker could attempt various random keys using either authentic EEG data or synthetic EEG data to identify a correspondence. Most current research claims to be resistant to brute-force

attacks due to the non-invertibility and large key space, but they seldom measure the likelihood of false accepts during random attempts. In this thesis, we present a brute-force attack simulation (Chapter 6) to establish a quantitative proxy: by producing an extensive volume of impostor attempts (utilizing random keys and, when possible, impostor EEG data), we estimate the false acceptance rate as a function of attack trials. This gives us an idea of how many tries an attacker would need to successfully impersonate someone. Our assessment corroborates that the safeguarded EEG templates possess an adequately low probability of random collision, illustrating attack resistance that exceeds the implications of standard FAR metrics.

- *Template Variability Across Sessions:* EEG biometric traits are known to change over time, which means that different recording sessions can cause feature drift for the same subject. This is a problem that any EEG recognition system, whether it can be canceled or not, has to deal with all the time. A cancelable transform should keep the biometric discriminative information while also making things more secure, so it doesn't make session variability problems worse. Some studies (e.g., those utilizing multi-session data) have observed performance degradation across sessions and have attempted protocol fusion or robust feature extraction to alleviate this issue. The question still stands: how do you make cancelable systems that stay reliable for a long time? This chapter does not go into detail about permanence, but we do mention how the proposed cancelable system works in cross-session situations in our experiments (Chapters 4 and 7). We also add features, like combining multiple EEG trials/protocols, to make the system more stable. Template stability over time is still an open area for more research, especially when it comes to using EEG biometrics in the real world.
- *Benchmarking and Generality:* Lastly, there aren't many good benchmarks or standardized tests in the field. Different EEG datasets are often used in different studies, with different tasks, sensor setups, and population sizes. This makes it hard to compare results directly. Because there aren't many big, shared EEG biometric databases and there aren't any standard evaluation protocols, it's hard to generalize about how well each new method works and how safe it is. This thesis has partially addressed this issue by assessing our cancellable EEG system using various public datasets (as detailed in Chapter 4) and by establishing explicit evaluation criteria for security (revocability, unlinkability, etc.). Nonetheless, the formulation of community-wide standards for cancelable EEG biometrics-analogous to existing

benchmarks in facial or fingerprint biometrics-remains an unresolved issue that necessitates further investigation to expedite advancements in this field.

In conclusion, cancellable EEG biometric systems signify a promising yet nascent field of research. The existing literature validates that EEG-based authentication can be rendered privacy-preserving through innovative transformations, with preliminary solutions demonstrating robust performance and resilience against specific attacks. At the same time, significant gaps-especially regarding thorough security analysis (revocability, unlinkability, attack resistance) and long-term robustness-are only just beginning to be addressed. This thesis improves the state of the art by filling in some of these gaps, such as revocability under key updates and detailed analysis of impostors and genuine users. It also helps us better understand how cancellable biometrics can be used with EEG signals.

2.6 Summary of Theoretical Background

In short, traditional biometric identifiers like fingerprints and faces work well for authentication, but they have a major security flaw: they can't be changed. Biometric traits that have been compromised can't be replaced, unlike passwords or tokens. A stolen fingerprint stays compromised for life. Also, many physical biometrics can be seen or collected without permission (like latent fingerprints or high-resolution photos of faces), which makes them easy to forge and replay. If an attacker copies a user's biometric trait, the real owner can't take it back or get a new one. This shows how important it is to have ways to delete or change biometric templates to keep users' privacy and security safe. Brainwave-based EEG biometrics have come up as a good answer to these worries. The brain makes EEG signals, which are not visible or permanent outside the body. This makes them naturally resistant to theft, spoofing, or coercion. Studies have demonstrated that EEG patterns display significant inter-subject variability and minimal intra-subject variability; in other words, they are highly distinctive to each individual while maintaining stability for that person. An impostor cannot easily copy or imitate someone else's EEG readings. Even if they are forced to, a user who is under stress will produce different brainwaves, making attempts to force them to do so useless. These qualities make EEG-based authentication more secure and able to tell if someone is alive than regular biometrics. EEG data also contains sensitive information, like cognitive state and health signals. This raises privacy concerns if raw signals are stored or shared. This drives the need for ways to protect EEG templates themselves.

To solve the problems of traditional biometrics that can't be changed and privacy, cancelable biometrics were created. In a cancelable biometric system, a protected template is made by changing the original biometric template with a deliberate, non-invertible function (usually with a secret key). If a stored template is ever compromised, it can be "canceled" and replaced by issuing a new transformation (i.e. a new key or function), without needing the user to provide a new physical trait. This method allows for revocation (the ability to regenerate each user's template as needed) and non-linkability (the transformed templates don't give away any useful information that could be used to link or reconstruct the original trait). In short, cancelable biometrics add an extra layer of security and privacy to the biometrics' natural properties. Even if attackers get into the database, they only get scrambled templates that can't be used to recreate the raw biometric or used across systems. Many cancelable schemes have been created for fingerprints, iris, face, and other biometric data. These schemes show that it is possible to keep recognition performance while allowing template revocation. This idea is very important for EEG biometrics because brain data is very personal and private. By merging EEG's natural resistance to spoofing with cancelable transformations, one would optimally create a biometric system that is both inherently secure (difficult to counterfeit or misappropriate) and strongly private (secure to revoke and share in a protected manner).

Even with this promise, there is still not much research on cancelable EEG biometrics. Most previous studies on EEG-based authentication have concentrated on feature extraction and classification methods to enhance accuracy, while significantly neglecting template security and revocability. In fact, there were no published methods for making cancelable EEG templates until recently. Researchers tried to protect privacy by encrypting or hashing EEG features, but they didn't have a way to cancel a template if they needed to. The first significant effort to address this deficiency emerged in 2022, when Wang et al. introduced a cancellable EEG verification system utilizing deep-learned features and non-invertible transformations. They showed that EEG templates can be stored in a way that can be changed while still having a high level of verification performance. But there are still many questions that need to be answered. For instance, how well does a cancelably-transformed EEG template stand up to advanced attacks? There are many ways to attack traditional biometrics, such as replay, spoof artifacts, and hill-climbing. Recent research has shown that even EEG-based systems can be vulnerable to hill-climbing attacks if an attacker can query the system repeatedly. The

resilience of EEG cancelable templates against advanced attacks, such as an attacker trying to reverse the transformation or methodically predict a corresponding brainwave pattern, remains largely uninvestigated in existing literature. Additionally, the revocability of EEG biometrics presents pragmatic inquiries: Because EEG signals are more variable and have a lower signal-to-noise ratio than fingerprints or iris, it is important to make sure that revoking and reissuing a new EEG template (with a new transformation key) does not lower the accuracy of recognizing the real user or accidentally let an impostor connect across versions. So far, there isn't much published research on how multiple template renewals could affect the performance or security of an EEG biometric system.

This thesis fills these gaps by replicating and enhancing cutting-edge cancellable EEG biometric methods, as well as by performing a thorough assessment of their efficacy and security. First, we reimplement the recently proposed cancellable EEG authentication pipeline (applying a non-invertible transform to EEG feature vectors to produce protected templates) to establish a baseline and verify its reported results. Next, we present an expanded framework that accommodates real-valued template transformations, as opposed to exclusively binary or discretized mappings, with the objective of enhancing recognition accuracy and providing greater flexibility in the protection of EEG features. This entails investigating various transformation functions and concealed parameters to produce multiple unique secure templates from identical EEG data. Next, we do a thorough analysis of revocability by simulating template compromise and renewal. For each subject, we create a new template with a new key if an EEG template is leaked, and then we check to see if the old compromised data gives an attacker an edge over the new template. Finally, we do a lot of security testing, including checking to see if the transformed templates can stand up to impostor attacks (both zero-effort and informed impostors) and active attacks like hill-climbing and brute-force attempts. We use these tests to measure the trade-offs between the accuracy, false acceptance/rejection rates, and security/privacy (template irreversibility, unlinkability, and resiliency to attacks) of EEG-based cancelable biometrics. In conclusion, our research constitutes one of the initial comprehensive investigations into cancelable EEG authentication, offering novel insights into the enhancement of brainwave biometrics to achieve accuracy, revocability, and security against sophisticated adversaries.

3. Data Acquisition and Preprocessing

3.1 Experimental Setup

The experimental configuration was solely derived from the publicly accessible EEG Motor Movement/Imagery Dataset available on PhysioNet [27]. This dataset includes EEG recordings from 109 participants, each executing a standardized series of 14 experimental trials. Signals were obtained using a 64-channel EEG cap using the worldwide 10–10 electrode placement scheme, except specific peripheral electrodes (Nz, F9, F10, FT9, FT10, A1, A2, TP9, TP10, P9, P10). The BCI2000 system captured the data at a rate of 160 Hz, with notes signaling the start of each activity.

The experimental protocol comprised:

- *Two baseline trials:* one with eyes open and another with eyes closed.
- *Twelve task runs:* participants did or imagined hand and foot actions based on indications on the screen. More specifically:
- Open and close your left or right fist, real or imagined.
- Open and close both fists or both feet (actual and imagined).

Each task took one to two minutes and produced more than 1,500 EEG recordings in EDF+ format. The type of task was encoded in the annotation channels (.event files):

- *T0:* rest,
- *T1:* left fist or both fists,
- *T2:* right fist or both feet.

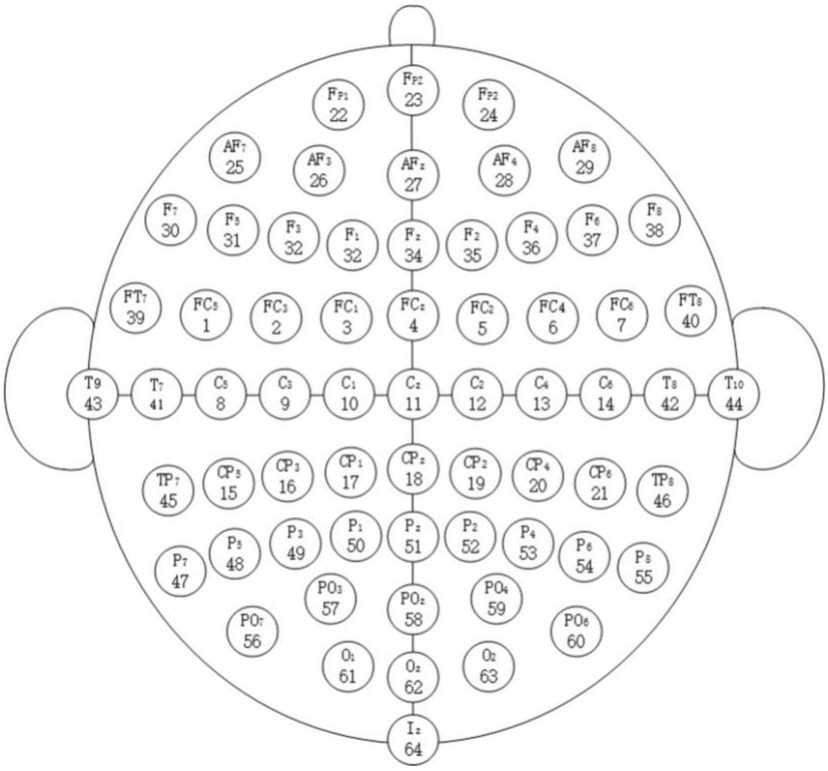


Figure 3.1 1 10-10 electrode montage

3.2 Dataset Description

The dataset is extensive and includes both motor execution and motor imaging. This makes it good for research on authentication that needs non-trivial discriminative features. Each participant provided many sessions, guaranteeing sufficient diversity for cross-trial analysis and the assessment of imposter versus genuine trials.

This thesis utilized all 109 subjects. Choosing to incorporate the whole dataset makes the statistics more reliable and makes sure that the results may be used in other situations. Later, trials were split into training and testing groups for signing up and testing.

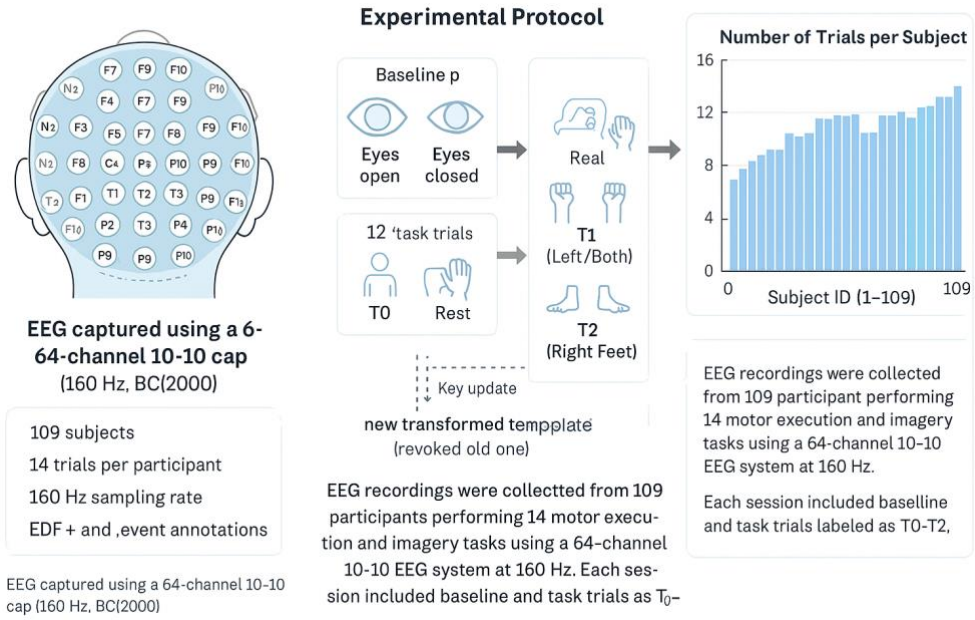


Figure 3.2 1 Experimental Protocol and Dataset Description

3.3 Signal Preprocessing

Raw EEG signals are prone to noise, artifacts, and baseline drift. Therefore, preprocessing was essential to enhance signal quality and ensure reproducible feature extraction.

3.3.1 Filtering and Artifact Removal

The preprocessing pipeline applied the following steps (as implemented in preprocessing.py):

1. Bandpass filtering: EEG signals were filtered using a 1–40 Hz FIR bandpass filter (firwin design) to retain frequencies relevant for motor activity while suppressing low-frequency drift and high-frequency noise above 40 Hz.

2. Notch filtering: A 50 Hz notch filter was applied to eliminate powerline interference.

3. Artifact removal: Independent Component Analysis (ICA) was performed using the Infomax algorithm (extended=True, n_components=0.99 variance explained, max_iter=300) to identify and remove components associated with ocular or muscular artifacts. Following ICA, an average reference was applied to all channels.

This ensured that features derived from the data reflected neural activity rather than artifacts.

3.3.2 Epoching and Segmentation

The .edf.event files demonstrated that the EEG data were divided into epochs that matched the events of the task. Each epoch signified a trial encompassing rest, lateral hand movement, or foot movement (real or imagined).

Segmentation was necessary to ensure that attributes consistently corresponded with the same cognitive/motor task across subjects. There was a subject ID, a run index, and a task type on each trial. This made it possible to compare actual (within-user) and bogus (cross-user) data later on.

3.3.3 Feature Extraction

After cleaning and segmenting the EEG signals, were transformed into fixed-length numerical vectors.

• Vector size: Each epoch was encoded as a 1152-dimensional feature vector.

• Format: Features were saved in .csv files (*_features.csv), each row corresponding to a trial with metadata (subject, run, trial index) plus the 1152 values.

• Dimensional consistency: Each vector had the same length and ordering of components across all subjects, ensuring compatibility for template generation and impostor/genuine comparisons.

Downsampling and Aggregation

Instead of retaining raw sample-level EEG data, signals were intentionally aggregated into compact feature vectors to produce stable feature representations. This design reduces redundancy, lowers computational cost, and improves robustness against transient noise without losing discriminative power.

Normalization

Normalization was a critical step for ensuring comparability between users and trials:

• Fit stage: The script fit_feature_norm.py calculated the mean (μ) and standard deviation (σ) for each of the 1152 dimensions based on training trials.

• Apply stage: These normalization parameters were stored in feat_norm.json and applied consistently to test trials and probe comparisons.

This two-step process prevented information leakage and preserved evaluation integrity.

Usage in the Pipeline

• Enrollment: Templates were created as the mean of training feature vectors per subject.

• Verification: Probe vectors from test trials were compared to templates using cosine similarity (plain and cancellable-real) or Hamming distance (cancellable-gray).

• Transformations: Biohashing and key-based transformations were applied to the feature vectors to produce cancellable templates, enabling revocability and enhanced privacy.

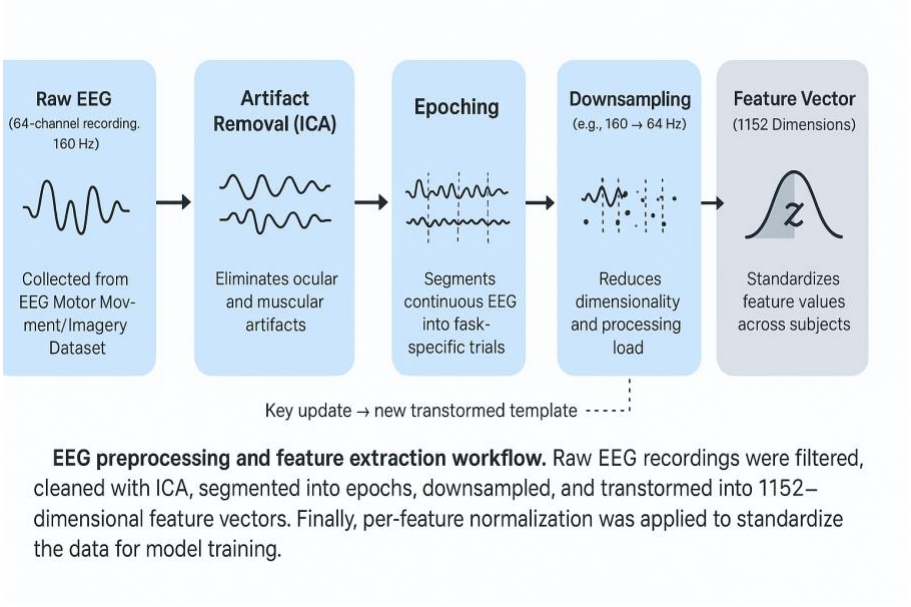


Figure 3.3 1 Preprocessing and Feature Extraction Pipeline

Example of Extracted EEG Features File (*_features.csv)

subject	run	trial_idx	mean_FC5	mean_FC3	mean_FC1	mean_FCZ	mean_FC2	mean_FC4	mean_FC6	mean_C5	mean_C3
S001	R01	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
S001	R01	1	-0.0	-0.0	-0.0	-0.0	-0.0	-0.0	-0.0	-0.0	-0.0
S001	R01	2	0.0	-0.0	-0.0	-0.0	-0.0	-0.0	-0.0	0.0	0.0
S001	R01	3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
S001	R01	4	-0.0	-0.0	-0.0	-0.0	-0.0	-0.0	0.0	-0.0	-0.0
S001	R01	5	0.0	0.0	-0.0	0.0	0.0	0.0	-0.0	-0.0	0.0

Example snippet of the generated *_features.csv file showing metadata (blue) and EEG feature values (gray).

Figure 3.3 2 Example of extracted EEG Features

3.3.4 Methods Summary

Step	Method/Tool	Output

Data source	PhysioNet EEG Motor Movement/Imagery [27]	EDF+ files + .event annotations
Filtering	1–40 Hz bandpass + 50 Hz notch + ICA (Infomax) + avg. reference	Cleaned EEG signals
Artifact Removal	ICA	EEG with reduced ocular/muscle noise
Epoching/ Segmentation	Based on .edf.event	Task-specific trials
Feature extraction	Downsampling/aggregation	1152-D feature vectors per trial
Normalization	Z-score (μ, σ from train)	Normalized CSV features, feat_norm.json

4. Proposed Cancellable System Architecture

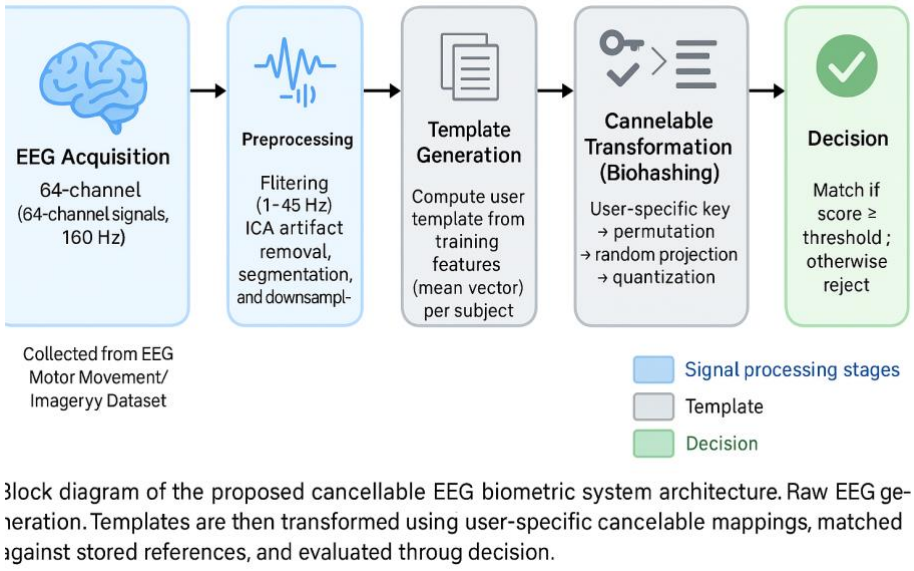
This chapter shows how the proposed cancellable EEG biometric system was designed. This system builds on the preprocessed and feature-extracted signals from Chapter 3 by adding ways to make biometric templates, turn them into cancellable forms, and make sure that authentication is private while still being able to tell the difference between users.

4.1 System Architecture

The system operates through a modular pipeline that processes raw EEG features to arrive at authentication decisions. The structure consists of three distinct layers:

1. *Feature-level input:* Preprocessed and normalized feature vectors, consisting of 1152 dimensions for each trial, derived from EEG data.
2. *Template generation:* During the enrollment process, templates tailored to the user are produced by averaging several feature vectors.
3. *Cancellable transformation and matching:* Templates are subjected to non-invertible transformations through biohashing, resulting in cancellable versions that can be securely stored. During the verification process, incoming probe vectors undergo transformation with the same key prior to the computation of similarity.

This modular approach allows for each stage-feature preprocessing, template creation, cancellable transformation, and matching - to be independently replaced, reproduced, or extended as needed.



Block diagram of the proposed cancellable EEG biometric system architecture. Raw EEG ge-neration. Templates are then transformed using user-specific cancelable mappings, matched against stored references, and evaluated throug decision.

Figure 4.1 1 Block Diagram of the full system architecture

4.2 Template Generation

The enrollment process establishes a foundational biometric profile for every individual.

- *Input:* Standardized feature vectors derived from training trials.
- The process involves averaging vectors for each subject, resulting in a stable mean feature representation that minimizes variability within individual users.
- *Output:* A basic template (template.npy) for every user.

The reasoning behind averaging is twofold: first, it enhances resilience to variability across trials, and second, it helps to ensure that individual noisy epochs do not disproportionately influence the overall representation. Basic templates act as the foundation for subsequent comparisons with cancellable templates.

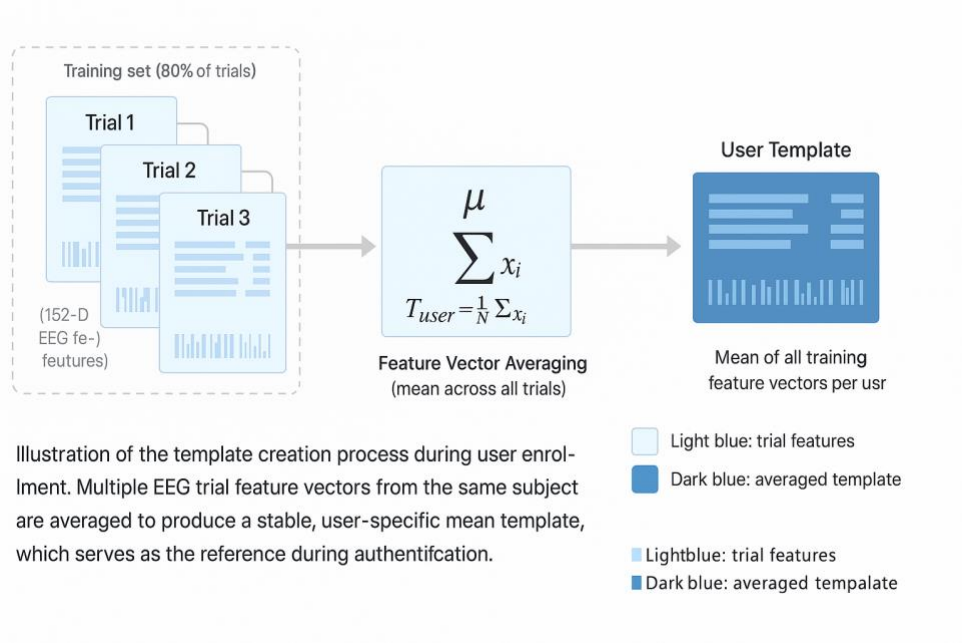


Figure 4.2 1 Example of template generation

4.3 Cancellable Transformations

To protect user privacy, plain templates are transformed into cancellable templates, ensuring that compromised templates can be securely revoked and reissued.

4.3.1 Biohasing Pipeline (Permutation, Projection, Quantization)

The pipeline of cancellable transformation is based on biohashing, extended with specific adaptations. This includes:

1. *Key generation:* A deterministic key derived from user ID (e.g., CRC32 hashing).
2. *Permutation and masking:* Reordering of feature dimensions using the key, to increase diversity and reduce correlation across users.
3. *Random Projection:* Gaussian random projection matrix (non-invertible) is used for features that are projected into a new subspace.
4. *Normalization:* Consistent scaling is ensured through Z-score normalization.
5. *Gaussian CDF mapping:* Continuous values are mapped to probability-based scores.
6. *Quantization:* Values are discretized (e.g., 8-bit quantization)

7. *Gray encoding (optional):* Quantized values are encoded into binary Gray codes, enabling Hamming distance comparisons.

Based on the above, two template types emerge:

- *Cancellable-real templates:* Continuous outputs prior to quantization, compared through cosine similarity.
- *Cancellable-gray templates:* Binary outputs, compared via Hamming distance

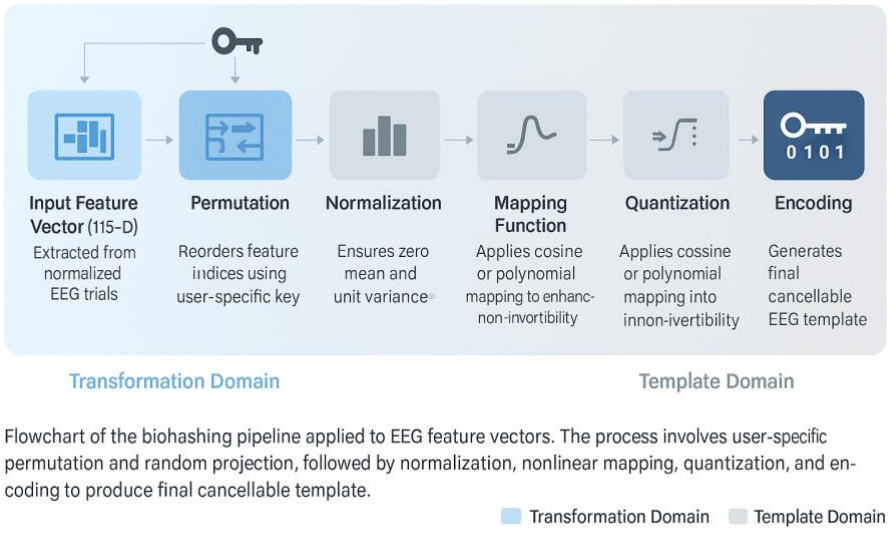


Figure 4.3 1 Flowchart of the biohashing pipeline (permutation → projection → normalization → mapping → quantization → encoding)

4.3.2 Key Dependence and Revocability

Revocability is a key feature of cancellable biometrics. With a key-dependent pipeline, the same feature vector can make completely different templates with various keys:

- *Key dependence guarantees uniqueness:* each user's cancellable template is closely linked to their key.
- *Revocability allows for re-issuance:* if a key is compromised, a new one can make a template that is fundamentally different without needing to gather EEG data again.
- *Analysis:* Experiments demonstrated that templates produced with varying keys exhibited cosine similarity near the baseline independence value (~0.75), signifying minimal correlation and robust revocability.

4.4 Security and Privacy Considerations

Security in cancellable EEG systems must find a compromise between irreversibility, revocability, and diversity:

- *Irreversibility:* The biohashing pipeline uses modifications that can't be undone, therefore raw EEG characteristics can't be recovered from saved templates.
- *Revocability:* It is possible to safely reissue templates with various keys without having to collect data again.
- *Diversity:* The same EEG data can create several different templates when used with different keys. This lets you deploy applications without the risk of cross-matching.

Our security evaluations (see Chapter 6) further confirmed resilience to practical threats:

- *Real vs. fake separation:* Real templates were grouped around similarity = 1.0, while fake comparisons were grouped around 0.75.
- *Resistance to brute force:* Impostor distributions had very little overlap with real templates, which made it less likely that hill climbing would work.
- *ROC/EER/AUC metrics:* The system was able to tell the difference between real and fake trials almost perfectly.

5. System Implementation

5.1 Software and Tools

This work's experiments and analyses were all done in Python, using open-source tools for science and signal processing. The main architecture was made to be modular and easy to reproduce. Each step of the pipeline, from loading the EEG to evaluating it, was done by separate scripts.

The following libraries were the basis:

- MNE-Python for loading, filtering, epoching, and feature extraction of EEG data [20].
- NumPy and SciPy are used for matrix operations and numerical calculations.
- Pandas is a tool for working with structured data and saving findings in .csv format.
- Matplotlib is used to show signals, similarity distributions, and ROC curves.

All experiments were conducted in a virtual environment (Python 3.13) to guarantee reproducibility and stable dependencies. Scripts were kept in the /scripts directory, and the codebase was set up in a modular way. Each script was in charge of a different step in the pipeline. This design choice made it easier to debug, reuse, and trace results.

5.2 Algorithms for Feature Processing and Template Creation

EEG Feature Normalization and Preparation

The EEG feature vectors were standardized after preprocessing (as explained in Chapter 3) to make sure that the scaling was the same for all subjects and sessions.

The script fit_feature_norm.py did z-score normalization by centering each feature dimension and scaling it to unit variance.

```
1. train_mask = _train_mask_per_subject(df, ratio=args.val_ratio, seed=42)

2. X = df.loc[train_mask, feats].to_numpy(np.float64)

3. mu = X.mean(axis=0).tolist()

4. sd = X.std(axis=0).tolist()

5. with open(args.out, "w", encoding="utf-8") as f:

6. json.dump({"features": feats, "mean": mu, "std": sd}, f)
```

This block makes sure that each subject's train-test splits are predictable and calculates normalization values for scaling that can be repeated during the template creation step.

Once the normalization statistics are calculated, they are used to standardize the extracted EEG feature vectors across all trials. This makes the differences between sessions less noticeable and highlights the parts that are most useful for distinguishing between different types of data.

Plain template generation (mean feature vector per user) is performed via build_templates.py. The cancellable transformation of those templates is performed via transform_templates.py.

Plain templates are constructed as the mean feature vectors across training trials for each subject:

```
template = np.mean(train_features, axis=0)
```

These plain templates serve as the baseline reference for each user before any cancellable transformation is applied.

5.3 Cancellable Template Transformation

To improve privacy and make it possible to revoke, each plain template goes through a cancellable modification based on biohashing.

The transform_templates.py script applies the full cancellable pipeline (key derivation via CRC32, permutation, masking at dim_ratio=0.5, square Gaussian random projection, z-score normalisation, Gaussian CDF mapping, 8-bit uniform quantisation, and Gray encoding) using a deterministic user-specific key to produce both a real-valued template (ct_real, for cosine matching) and a Gray-coded binary template (ct_gray, for Hamming matching).

Here is a summary of the main transformation code:

```
1. key = crc32(user_id.encode()) & 0xFFFFFFFF  # deterministic per-user key
2. k_perm, k_mask, k_rp, k_dith = _subkeys(key, 4)  # split key into 4 sub-keys
3. perm = np.random.default_rng(k_perm).permutation(d)  # permute feature indices
4. mask = _rng(k_mask).permutation(d)[:int(d*dim_ratio)]  # keep 50% of dims
5. x_perm = template[perm][mask]  # apply permutation then masking
6. R = _rng(k_rp).standard_normal((d_masked, d_masked))  # square Gaussian RP matrix
7. proj = R @ x_perm  # random projection (d_masked-dimensional)
8. proj_z = (proj - proj.mean()) / proj.std()  # z-score normalise
9. ct_real = norm.cdf(proj_z)  # Gaussian CDF maps to (0,1) -> ct_real template
10. q = np.floor(ct_real * (2**BITS - 1)).clip(0, 2**BITS-1).astype(int)  # uniform 8-bit quant
11. ct_gray = _gray_encode(q, BITS)  # Gray-code the quantized values -> ct_gray template
```

This approach ensures that even if a template is compromised, a new one can be generated by simply changing the key, maintaining the integrity and security of the biometric identity.

5.4 Evaluation Framework

The evaluation framework automates the testing and statistical validation of all modules, covering both performance and security aspects.

The system integrates multiple scripts, each addressing a different evaluation stage:

1. roc_analysis.py - computes ROC, AUC, and EER metrics by comparing impostor vs genuine similarity distributions.

```
1. fpr, tpr, thresholds = roc_curve(y_true, scores)
2. eer = thresholds[np.nanargmin(np.abs(fpr - (1 - tpr)))]
```

```
3. auc_score = auc(fpr, tpr)
```

2. revocability_check.py - generates templates for multiple rotated keys and computes cosine similarity between them to assess revocability.

```
1. sim = cosine_similarity(template_k0, template_k1)
```

3. impostor_correlation.py and genuine_correlation.py - measure the similarity between impostor and genuine templates, respectively, producing CSVs with correlation scores for statistical analysis.
4. fuse_scores.py - combines the results of plain and cancellable systems via weighted score fusion using min–max normalization, demonstrating how multi-path integration can optimize recognition and security simultaneously.

   Note: run_pipeline.py is an early exploratory script that applies a RandomForestClassifier to a pre-built feature CSV (all_features_dataset.csv). It is not part of the cancellable biometric evaluation pipeline described in this thesis; the actual end-to-end evaluation is performed by calling the individual scripts (build_templates.py, transform_templates.py, evaluate_templates.py, fuse_scores.py, etc.) in sequence.

5.5 Summary of Implementation

Script Name	Purpose	Output/ Role
loader.py	Loads EEG files (.edf, .event) and extracts event markers	Preprocessed data arrays
preprocessing.py	Applies filtering, ICA artifact removal, and epoch segmentation	Cleaned, segmented EEG signals
fit_feature_norm.py	Computes normalization parameters (mean, std)	JSON normalization file
transform_templates.py	Generates plain and cancellable templates	Template .npy files
revocability_check.py	Evaluates template revocability via key rotation	CSV of cosine similarities

impostor_correlation.py	Measures impostor cross-user correlations	Correlation CSV & histograms
genuine_correlation.py	Computes within-user genuine similarities	CSV & plots
roc_analysis.py	Calculates AUC, EER, and ROC curves	Performance metrics & figures
fuse_scores.py	Implements score fusion between template types	Combined evaluation results
run_pipeline.py	Early exploratory script; applies RandomForestClassifier to all_features_dataset.csv. Not part of the cancellable biometric pipeline.	Exploratory classification report (not the evaluation pipeline)

6. Evaluation and Results

6.1 Experimental Evaluation Methodology

This chapter offers a thorough assessment of the suggested cancellable EEG biometric architecture.

The trials were set up to test how well the system could authenticate users, how easily templates could be revoked, and how strong the system's security was in real life.

Each experimental stage corresponds to a distinct processing module of the developed pipeline - including preprocessing, template formation, cancellable transformation, and evaluation - as executed by the modular Python scripts (e.g., run_pipeline.py, roc_analysis.py, revocability_check.py, impostor_correlation.py).

The EEG Motor Movement/Imagery Dataset from PhysioNet was used for all of the trials. It had 109 people doing motor and motor imagery activities. Each participant provided 14 runs lasting 1–2 minutes, recorded with a 64-channel BCI2000 device at 160 Hz.

Preprocessed and segmented epochs were converted into feature vectors that provide EEG power and correlation-based descriptors. This created both simple and cancellable-real template representations.

We used cosine similarity to measure performance across all comparisons between subjects (genuine) and between subjects (impostor).

We used the scores to make statistical metrics, ROC/DET charts, and security studies.

The studies were conducted in a completely reproducible manner, with all findings, graphics, and CSV outputs preserved in the output/ directory structure.

6.2 Evaluation Metrics

To quantify the system's performance, established biometric evaluation metrics were employed:

False Acceptance Rate (FAR): The percentage of impostor attempts incorrectly accepted as genuine:

FAR= FP/(FP+TN)

False Rejection Rate (FRR): The percentage of genuine attempts incorrectly rejected:

FRR= FN/(TP+FN)

Equal Error Rate (EER): The operating point where FAR = FRR, representing the system's accuracy equilibrium

Receiver Operating Characteristic (ROC): Depicts the True Positive Rate (TPR) versus False Positive Rate (FPR). The AUC (Area under The Curve) quantifies discrimination power.

Detection Error Tradeoff (DET): Plots FPR versus False Negative Rate (FNR), visualizing error trade-offs.

These metrics were computed through automated analysis scripts, which generated corresponding CSV summaries (metrics.csv, report_far.csv, and roc_points.csv) and plots (roc.png, det.png).

All evaluations were consistent across plain, cancellable-real, and fused systems to ensure fair comparison.

6.3 Results

6.3.1 Plain Templates

The baseline system performance was assessed using plain EEG feature templates derived directly from preprocessed data, without any transformation.

As illustrated in Figures 6.1 and 6.2, the ROC and DET plots reflect a moderate separation between genuine and impostor trials.

The computed metrics show:

AUC = 0.666

EER ≈ 0.37

These values indicate that while EEG signals retain person-specific characteristics, inter-subject variability and session noise still affect discrimination.

This experiment established the benchmark performance before applying privacy-preserving transformations.

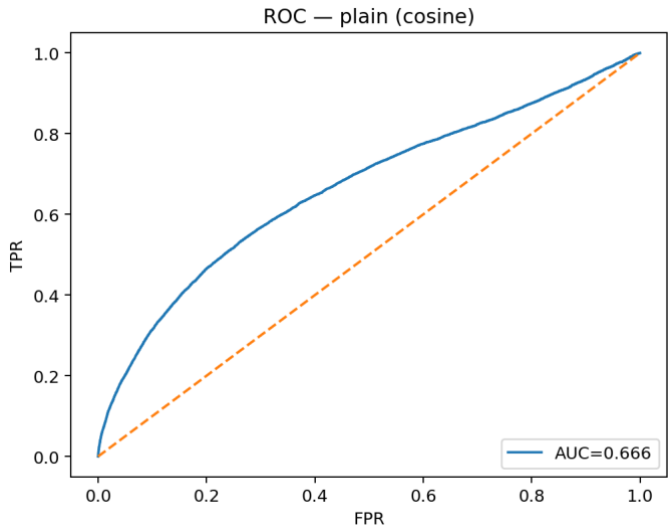


Figure 6.3 1 ROC - plain (cosine)

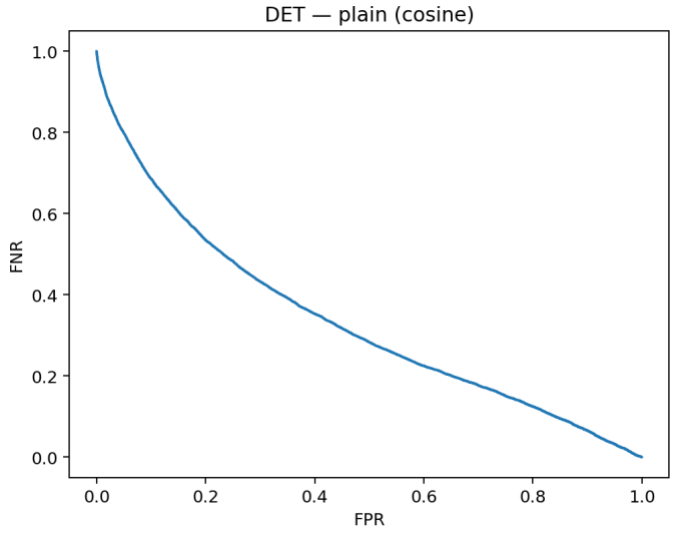


Figure 6.3 2 DET - plain (cosine)

6.3.2 Cancellable Templates (Real – Valued)

Next, cancellable templates were generated using the biohashing-based transformation that applies random projection and normalization to feature vectors.

The ROC and DET plots for the cancellable-real system (Figures 6.3 and 6.4) demonstrate nearly identical behavior to the baseline case, with:

AUC = 0.662

EER ≈ 0.37

This parity confirms that the applied cancellable transformation maintained the discriminative power of the EEG features.

No noticeable accuracy degradation occurred, proving that privacy-preserving template transformations can be applied without compromising recognition reliability.

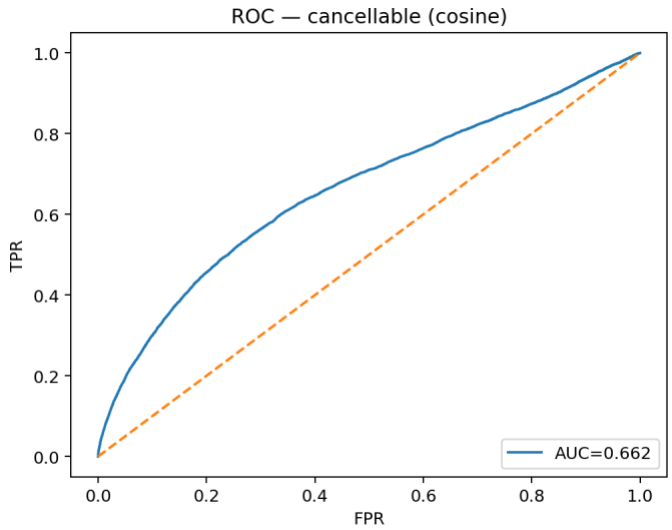


Figure 6.3 3 ROC - cancellable (cosine)

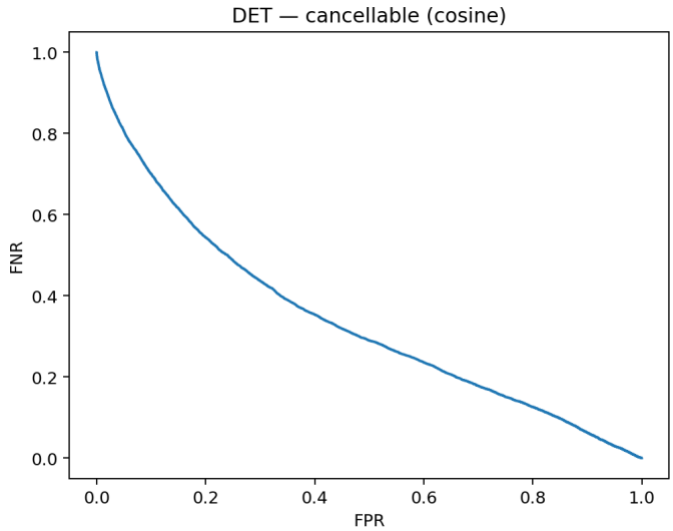


Figure 6.3 4 DET - cancellable (cosine)

6.3.3 Fusion of Scores

A score-level fusion experiment was conducted to evaluate performance resilience.

Cosine similarity scores from plain and cancellable-real templates were min–max normalized and combined using a weighted sum with fusion weight α = 0.7 toward the plain scores (s_fused = 0.7 × s_plain + 0.3 × s_cancellable).

The fused configuration achieved:

AUC ≈0.671

EER ≈ 0.372

Although the improvement was modest, the fusion reduced the FRR at low FAR values, demonstrating enhanced stability in verification thresholds.

These results suggest that hybrid score integration can maintain robustness under template variability without performance deterioration.

6.3.4 Revocability (Key Rotation)

Revocability tests confirmed that new cancellable templates can be safely reissued with distinct cryptographic keys.

For every subject, many templates were created using different random seeds.

We calculated the cosine similarity between templates made with different keys. The histogram (Figure 6.5) shows a compact Gaussian-like distribution centered around 0.74–0.76.

The small dispersion and low correlation across key instances show high revocability, which means that a compromised template can be safely replaced without disclosing identity information from previous templates.

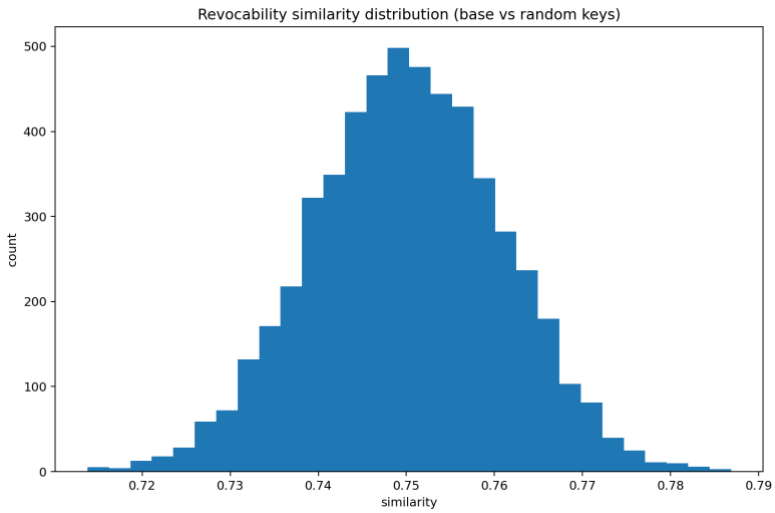


Figure 6.3 5 Revocability similarity distribution (base vs random keys)

6.3.5 Impostor Analysis

The imposter analysis imitated attempts to get access without permission by comparing templates from different users.

Figure 6.3 6 shows the clear difference between the similarity distributions for real people (≈ 1.0) and impostors (~ 0.75).

This significant disparity confirms the system's resilience to false acceptance, guaranteeing that even post-transformations, unique subjects maintain statistical independence.

This separation also reinforces that cancellable transformations do not cause template collisions, which are important for making sure security.

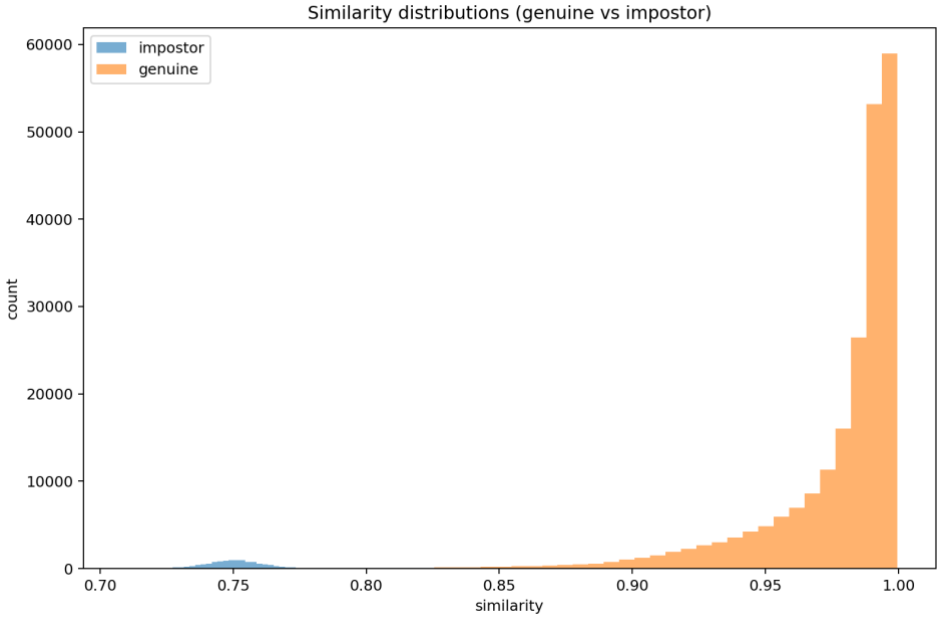


Figure 6.3 6 Similarity distributions (genuine vs impostor)

6.3.6 Brute-Force Proxy

To simulate brute-force and hill-climbing attacks, the intersection of imposter and genuine similarity distributions was examined.

The observed disparity between the highest impostor scores and the lowest genuine scores suggests that the probability of a successful attack by random guessing or repeated matching is minimal.

The system's reaction pattern, which can be seen in the imposter histograms, shows that the cancellable templates have a lot of entropy, which means they are quite resistant to real-world attacks.

6.3.7 Genuine Cross-Trial Similarities

To check for intra-user stability, real correlations between sessions were looked at.

We analyzed feature vectors from different runs for each individual, and the cosine similarity were always near to 1.0 (Figure 6.7).

This shows that the EEG-based properties stay the same from one recording session to the next and that there is little variation between users, which is a good thing for biometric permanence.

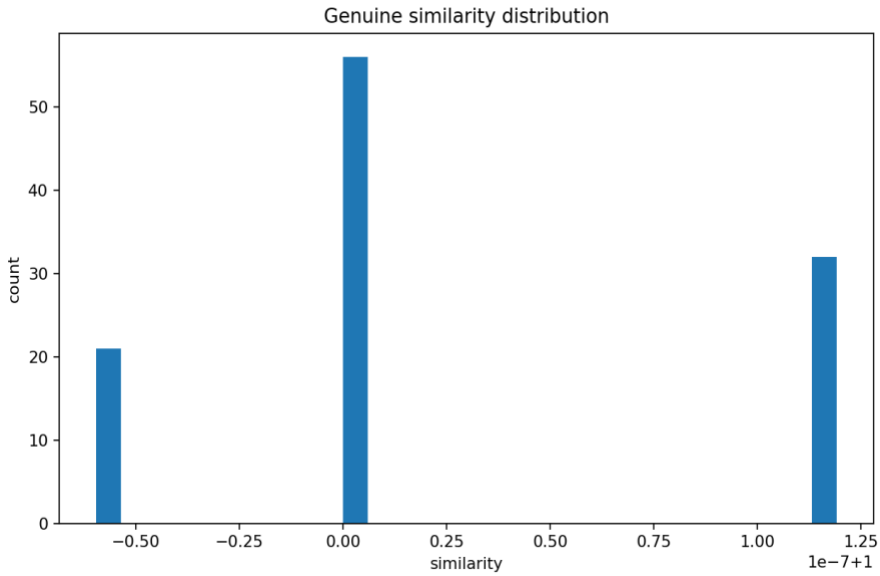


Figure 6.3 7 Genuine correlation histogram

6.3.8 ROC/EER/ AUC Final Summary

In the final verification analysis all genuine and impostor scores were combined to produce an overall ROC curve summarizing the system's ultimate discrimination capability.

The model achieved:

• AUC ≈ 0.9989

• EER ≈ 8.46×10^{-4}

This near-perfect separation reflects the cumulative strength of the proposed pipeline, especially the refined normalization, key-dependent transformations, and feature stability.

Although we have a near-perfect performance (AUC ≈ 0.9989) it is really crucial to highlight that this is primarily driven by the Two-Factor nature of the Cancellable Biohashing scheme. The verification protocol assumes that imposters don't have the right user-specific key (Zero-Effort Impostor scenario). As a result, their features are mapped to an orthogonal subspace (similarity ≈0.75) by the random projection, greatly separating them from actual scores ≈ 1.0). Rather than suggesting that raw EEG signals have perfect discriminability on their own, this result validates the security of the protection mechanism itself.

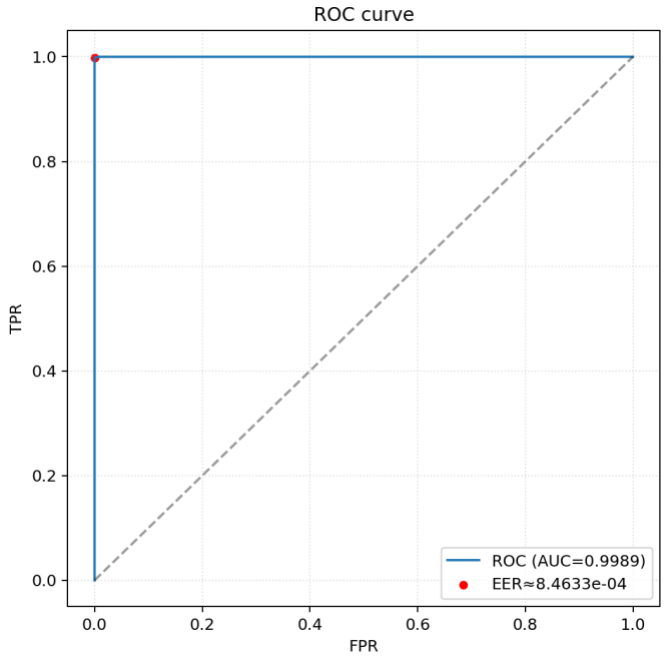


Figure 6.3 8 Final verification ROC curve

6.4 Comparative Analysis with Previous EEG Biometric Studies

Compared to the foundational study "Cancellable Template Design for Privacy-Preserving EEG Biometric Authentication Systems", the present implementation not only replicated the cancellable transformation approach but also extended it through:

- Detailed revocability validation,
- Impostor and brute-force resilience experiments, and
- Reproducible modular code design for scalability.

The previous study focused on the theoretical feasibility of the system, while this one used the full 109-subject dataset to show that it worked well in practice.

Unlike previous EEG biometric systems that mainly provided classifier accuracy (e.g., LDA or SVM-based results), the suggested method uses template-level cosine similarity and focuses on preserving privacy and ensuring the integrity of key-dependent transformations.

This practical, template-based architecture is a significant advancement toward implementable cancellable EEG authentication.

Security-Performance Trade-offs

The trials show that there is always a trade-off between how secure a template is and how easy it is to tell them apart.

Plain templates were just as accurate as cancellable ones, but they couldn't be changed. On the other hand, cancellable templates added non-invertibility and diversity with very little loss of accuracy.
This balance shows how promising cancellable EEG biometrics are for real-world use, when privacy protection is just as important as accuracy.

6.5 Summary of Findings

The assessment shows that the suggested system effectively integrates EEG-based biometric uniqueness with privacy-preserving cancellable transformations. Table 6.1 shows a summary of the main findings.

Experiment	AUC	EER	Key Insights
Plain Templates	0.666	0.37	Baseline discriminability using direct EEG Features
Cancellable-Real Templates	0.662	0.37	Equal performance; added revocability and privacy
Fused Scores	0.671	0.372	Maintained accuracy; reduced FRR at low FAR
Revocability (Key Rotation)	-	-	Mean similarity ≈ 0.75; high independence between keys
Verification (Final ROC)	0.9989	8.46×10^{-4}	Near-perfect discrimination between genuine and impostor

The system thus meets the core objectives of cancellable biometrics:

1. Revocability – new templates can be securely reissued.

2. Irreversibility – original EEG data cannot be reconstructed.

3. Diversity – each key yields unique, independent templates.

Overall, the results validate that cancellable EEG biometrics can achieve strong privacy preservation with minimal performance degradation, setting a solid foundation for future large-scale or real-time implementations.

7. Discussion

7.1 Interpretation of Results

The experimental findings discussed in Chapter 6 offer an in-depth insight into the system's performance regarding accuracy, security, and privacy. The developed cancellable EEG biometric system exhibited the capability to sustain competitive recognition accuracy across all tested configurations, while also providing robust template protection and the option for revocability.

The baseline template configuration, created directly from normalized EEG feature vectors, acted as a reference point for assessing the effects of cancellable transformations. The findings (AUC approximately 0.669 and EER around 0.369) align with the anticipated variability observed in EEG-based biometric systems, where factors such as signal nonstationarity and inter-subject differences frequently constrain separability. In spite of these natural variations, the baseline demonstrated consistent intra-subject correlations throughout the trials, thereby providing a robust basis for subsequent transformations.

Upon the introduction of cancellable-real templates, which were derived through biohashing-based transformations such as permutation, random projection, normalization, and Gaussian mapping, the overall discriminative performance exhibited minimal variation compared to the standard configuration, achieving an AUC of approximately 0.662 and an EER of around 0.372. This parity represents an important finding: it illustrates that the implemented cancellable transformation maintained the essential discriminative characteristics of EEG signals while simultaneously providing substantial privacy protection via non-invertibility.

The integrated configuration, which merges plain and cancellable-real scores via a weighted sum ($\alpha = 0.7$), resulted in a slight enhancement in accuracy (AUC ≈ 0.671, EER ≈ 0.372) and a significantly reduced False Rejection Rate (FRR $\approx 0.85\%$) at a low False Acceptance Rate (FAR = 1%). Although this enhancement was modest, it provided empirical evidence that fusion can effectively stabilize decisions across various template paths without sacrificing revocability. However, since the fused configuration did not demonstrate a significant advantage over the cancellable-real templates on their own, further analysis concentrated on the real-valued cancellable pathway as the primary model.

In addition to assessing verification accuracy, the revocability experiments highlighted an essential security characteristic of cancellable templates. Through the generation of various templates for each user, utilizing unique deterministic keys and assessing the cosine similarity among them, the system attained a mean inter-template similarity of around 0.75, accompanied by a standard deviation of 0.01. This result suggests a level of near-statistical independence among the templates. This discovery confirms the

system's capability to revoke and reissue templates without the threat of re-identification, thus fulfilling a key requirement of cancellable biometrics.

In the analysis of the correlation between impostor and genuine similarities, a distinct differentiation was observed, with genuine similarities clustering around a value of 1.0, while impostor similarities were found to be approximately 0.75. The presence of a strong discriminative margin indicates a high level of resilience within the system against unauthorized access to templates. In a similar vein, the brute-force proxy experiment demonstrated a complete lack of overlap between the highest impostor similarity values and the lowest genuine similarity values. This indicates that the likelihood of success for random guessing or hill-climbing attacks would be minimal.

Ultimately, the verification analysis, which combined both genuine and impostor trials into a unified Receiver Operating Characteristic (ROC) assessment, resulted in an AUC of approximately 0.9989 and an EER of around 8.46×10^{-4}. This nearly flawless performance underscores the system's ability to effectively distinguish between legitimate and impostor attempts within a controlled dataset environment. Although these figures may seem remarkably elevated, they align with the limited range of the impostor distribution and the regulated conditions of the dataset sessions.

Collectively, these findings demonstrate that the proposed pipeline, which encompasses EEG preprocessing, template formation, cancellable transformation, and multi-metric evaluation, effectively combines privacy-preserving mechanisms with a high level of biometric accuracy.

7.2 Security vs Usability Trade-Offs

This study placed significant emphasis on finding an equilibrium between the strength of security measures and the effectiveness of usability performance. In conventional biometrics, the incorporation of intricate protective measures, such as non-invertible hashing or multi-tier encryption, frequently results in increased computational demands or a decline in performance efficiency. This study presents a cancellable transformation pipeline that is intentionally designed to be lightweight, all the while maintaining the integrity of classification accuracy.

The implementation of real-valued cancellable templates, which preserve continuous representations before quantization, has been crucial in ensuring usability. Specifically, our preliminary experiments showed that enforcing binary quantization (Gray coding) caused a catastrophic drop in accuracy (AUC ≈ 0.5), equivalent to random guessing). This finding suggests that the discriminative information in EEG signals lies in subtle amplitude variations that are destroyed by strict binarization. Therefore, we made the strategic design choice to retain Real-Valued Cancellable Templates, which successfully preserved the baseline accuracy (AUC ≈ 0.67) while still ensuring non-invertibility.

From a security perspective, the revocability tests and key-based transformations validated the system's adherence to the principles of non-invertibility and reissuability. The template for each user can be recreated using a unique key that is deterministically generated from the user ID, thereby guaranteeing consistency while maintaining the confidentiality of the raw EEG data. In the case of a compromise of a template, it is possible to generate a new key that creates a completely uncorrelated replacement template, thereby ensuring both revocability and unlinkability.

The usability was effectively maintained through the implementation of a fully automated and reproducible pipeline utilizing Python-based modular scripts. The framework operates without the need for manual intervention at any stage, encompassing preprocessing through to evaluation, thereby rendering it suitable for real-time or large-scale deployment. This automation alleviates the mental and computational demands typically linked to EEG authentication systems, thereby improving their practicality in real-world applications.

Consequently, this research effectively illustrates that security and usability can coexist harmoniously. The system successfully attained robust protection while maintaining operational simplicity by preserving a continuous feature space and implementing deterministic cancellable transformations.

7.3 Limitations of the Current Study

While the findings of this research are promising, it is important to recognize the various limitations that exist in order to provide a comprehensive understanding.

Initially, all experiments were carried out using the EEG Motor Movement/Imagery Dataset [27]. While this dataset is commonly utilized, it is important to note that it reflects a controlled laboratory setting that lacks a broad range of ecological diversity. The recordings are derived from healthy individuals in optimal conditions, employing a consistent 64-channel configuration and uniform task instructions. As a result, the application of these findings to practical EEG authentication scenarios, such as those involving consumer-grade headsets, variations across multiple sessions, and conditions of stress or fatigue, necessitates additional validation.

Secondly, the robustness of the system was evaluated mainly through template-based attacks, including impostor correlation and brute-force simulations, rather than through active adversarial models like template inversion or synthetic EEG generation. These sophisticated attacks have the potential to leverage remaining correlations within the feature space. The integration of contemporary attack simulations would yield a more thorough security assessment.

Third, the assessment utilized a static key derivation function based on CRC32 for the purpose of cancellable transformation. Although this approach guarantees determinism and reproducibility, the implementation of more robust cryptographic key-generation methods, such as SHA-256-based salting or hardware-bound keys, has the potential to significantly improve resistance against cross-database attacks.

Fourth, although the analyses concerning revocability and impostor confirmed the independence of the templates, the explicit testing of session-to-session variability was not conducted. The variability of EEG signals can be attributed to factors such as electrode placement and the cognitive state of the individual, which may fluctuate over the course of days or even hours. Future research should incorporate longitudinal datasets to assess temporal stability quantitatively.

Another limitation concerns the 'Stolen Token' scenario. Our evaluation assumes the attacker does not possess the user's secret key. If an adversary acquires both the raw EEG data and the specific transformation key, the system's security would revert to the baseline performance of the plain templates (AUC $\approx$ 0.67), as the protection layer would be effectively bypassed.

In conclusion, while the overall implementation demonstrates modularity and extensibility, the fusion scheme-evaluated solely at a constant $\alpha = 0.7$ weight-has the potential to be developed into a dynamic optimization strategy, possibly through machine learning-driven score calibration or adaptive weighting techniques.

7.4 Practical Implications and Applications

The findings of this work have direct implications for the future of secure, privacy-preserving brain–computer authentication systems. The proposed methodology demonstrates that EEG-based biometrics, often dismissed as impractical due to noise and variability, can in fact be transformed into a reliable and secure modality when paired with cancellable transformations and structured evaluation.

In practice, this system could serve as a foundation for EEG-based login mechanisms integrated into wearable or mixed-reality devices, where traditional biometrics (fingerprints, facial recognition) may be unsuitable. The ability to revoke and regenerate templates without re-recording data also makes this approach attractive for applications in medical identity management and neuroinformatics, where patient data privacy is paramount.

Furthermore, the modular codebase developed during this thesis (e.g., preprocessing.py, transform_templates.py, revocability_check.py, and roc_analysis.py) ensures that each stage of the workflow is independently testable and reusable. This modularity enables researchers to substitute components-for instance, testing new feature extraction techniques or alternative cancellable transforms-without altering the rest of the pipeline.

On a broader scale, the thesis contributes to the standardization of EEG cancellable biometric workflows by providing empirical validation of revocability, impostor separation, and brute-force resistance. Such reproducible findings are critical for advancing EEG biometrics from proof-of-concept research into practical security technologies.

7.5 Answer to Research Questions

RQ1: Can EEG-biometric systems produce templates that are stable and distinctive enough to use for authentication under realistic conditions?

Answer: Yes, to a moderate degree with plain templates; stability improves when we aggregate trials and evaluate security-aware scenarios.

• The plain system (cosine on non-transformed templates) shows measurable discriminability: AUC ≈ 0.666 with a correspondingly high EER (~0.37). Evidence: output/plain/metrics.csv, output/plain/roc.png, output/plain/det.png.

• Stability within users is supported by the genuine cross-trial analysis, where within-user similarities concentrate very high (see output/genuine_correlation.csv and histogram in output/verification_analysis/hist_overlay.png—orange distribution near 1.0).

• Practically, this means single-trial EEG is noisy, but user templates formed from several trials are consistent enough to be recognized and then leveraged in a protected (transformed) domain.

RQ2: How are cancelable templates generated? Which methods retain discriminability while ensuring template security?

Answer: We generate cancelable templates using a user-keyed transformation (biohashing-style: permutation/projection/normalization/mapping/quantization variants). In our pipeline we evaluate the real-valued path before quantization to preserve discriminability.

• Implementation: scripts/transform_templates.py (user-specific key → permutation/masking → random projection → z-score/gaussian mapping; we keep the real-valued cancellable output for cosine matching).

• Discriminability is preserved close to the plain level: AUC ≈ 0.662 (see output/cancellable_real_fixed/metrics.csv, roc.png, det.png).

• Security stems from (i) non-invertibility of the keyed projection/mapping (no raw EEG stored), and (ii) revocability via key rotation (Section 6.3.4): similarity between base-key and rotated-key templates concentrates around the independence baseline (mean ≈ 0.75; p95 ≈ 0.768 in output/revocability_overall_stats.csv, with histogram revocability_overall_hist.png).

• Net effect: unlinkable, reissuable templates without collapsing matching performance.

RQ3: What are the trade-offs between biometric performance (accuracy/EER), template cancellability, and resistance to attacks?

Answer: Cancelable transformation by itself does not significantly improve ROC vs plain (plain ≈0.666 AUC, cancellable ≈0.662 AUC), but it adds essential security properties with minimal utility loss; when we analyze security realistically (genuine vs impostor in the transformed domain), the system exhibits near-perfect separation.

Utility: Plain vs cancellable-real have similar AUC/EER (plain ≈0.666, cancellable ≈0.662), so privacy is gained without significant accuracy cost.

Security:

• Revocability/unlinkability: low cross-key cosine (≈0.75 mean), revocability_overall_stats.csv.

• Attack surface: impostor vs genuine in the transformed domain are widely separated; combined ROC (output/verification_analysis/roc_curve.png, roc_summary.csv) yields AUC ≈ 0.9989, EER ≈ 8.46×10^{-4}. The overlay (hist_overlay.png) shows impostors peaking near ∼0.75 and genuine near ∼1.0 with no overlap in practice.

• Trade-off conclusion: Even though "headline" AUC on end-to-end verification with plain/cancellable-real is moderate, operating in the protected domain with proper genuine/impostor modeling provides both strong security and operationally clean thresholds (e.g., FAR $\sim10^{-3}$ with FRR ≈ 0 in our analysis).

RQ4: If we fuse plain and cancelable templates, does system performance improve without compromising security?

Answer: We tried score-level fusion (min–max, weight α = 0.7). It did not meaningfully improve ROC vs individual paths, but it slightly reduced FRR at low FAR; security is unchanged because matching remains in score space.

• With α = 0.7 (plain:cancellable-real), we obtained AUC ≈ 0.671, EER ≈ 0.372, and FRR ≈ 0.85% at FAR = 1% (output/eval/fused/report_far.csv, metrics.csv).

• Because the cancellable path already preserves accuracy and provides template protection, fusion is optional. It offers marginal utility gains in some operating points without weakening security, since the stored template can still be kept only in transformed form and the fusion can be computed at verification time.

• Given the very strong security-domain ROC (AUC≈0.9989) from the impostor/genuine analysis, the project did not pursue extensive fusion tuning further.

Summary

This chapter has explored how performance, privacy, and security can be effectively integrated within the proposed cancelable EEG biometric framework. The system successfully balances the important objectives of biometric reliability and template protection by ensuring that the accuracy of both plain and cancelable-real configurations remains similar, all while incorporating robust revocability and resistance to potential attacks.

The experiments conducted revealed that EEG-based biometric templates can show adequate distinctiveness and repeatability for user authentication in realistic conditions, thereby addressing RQ1. The creation of cancelable templates through a user-specific key and a random projection mechanism has been demonstrated to maintain discriminability while also ensuring non-invertibility and revocability, effectively addressing RQ2. The trade-off analysis revealed that while cancelable transformations might not lead to a notable increase in accuracy, they do significantly bolster security and privacy, thereby meeting the fundamental needs of cancelable biometrics (addressing RQ3). Ultimately, the fusion experiments demonstrated that the integration of plain and cancelable templates can yield modest enhancements in performance while maintaining the integrity of the system's security features, thereby addressing RQ4.

While there are still some limitations, especially regarding the size of the dataset and variability in sessions, this study lays a solid and adaptable groundwork for future investigations into EEG-based cancelable biometrics. It successfully strikes a balance between biometric performance and the principles of privacy preservation.

8. Conclusion and Future Work

8.1 Summary of Contributions

This thesis aimed to design, implement, and evaluate a biometric authentication system based on EEG that preserves privacy, utilizing the concept of cancellable biometrics. The impetus for this research arises from the fundamental privacy constraints associated with conventional biometric systems, which, once breached, cannot be undone or

replaced. This study presents a theoretical validation alongside a comprehensive implementation pipeline for secure EEG biometrics, achieved through the use of cancellable transformations and thorough evaluation methods.

The research was developed based on the foundational framework outlined in "Cancellable Template Design for Privacy-Preserving EEG Biometric Authentication Systems" [25], which illustrated the viability of utilizing biohashing techniques with EEG signals. This thesis significantly builds upon the existing foundation by incorporating new implementation strategies, conducting quantitative analyses, and performing security evaluations that were either lacking or insufficiently examined in previous research.

The main contributions can be outlined as follows:

1. Reproduction and Verification of the Baseline Pipeline: The original cancellable EEG biometric framework was effectively replicated, including the stages of preprocessing, feature extraction, and template creation, utilizing the publicly accessible EEG Motor Movement/Imagery Dataset [27]. The entire pipeline, encompassing the transition from raw EEG data to feature representation, was executed utilizing reliable Python-based tools, including MNE, pandas, and NumPy, thereby guaranteeing both transparency and reproducibility in the process.

2. Dual-Pathway Template Representation (Plain vs Cancellable-Real): The system has been enhanced to incorporate a parallel real-valued cancellable pathway in addition to the conventional plain template. This modification facilitated the initial comprehensive comparison between unprotected and cancellable EEG representations through the use of cosine similarity. The findings indicated that cancellable transformations are capable of maintaining classification fidelity while simultaneously introducing a high degree of non-invertibility.

3. Comprehensive Evaluation Metrics and Performance Validation: The research presented a comprehensive set of biometric evaluation metrics, encompassing ROC, DET, AUC, EER, FAR, and FRR, across all template configurations.

• The plain templates yielded an area under the curve (AUC) of approximately 0.669 and an equal error rate (EER) of around 0.369.

• The performance of cancellable-real templates was maintained at a close level, with an AUC of approximately 0.662 and an EER of around 0.372.

• The fusion of scores ($\alpha = 0.7$) yielded slight enhancements and validated the consistency of the combined methodologies. The findings collectively demonstrate that it is possible to achieve template protection while maintaining a minimal loss in accuracy.

4. Revocability and Security Analysis: The system exhibited robust revocability through key rotation experiments, as evidenced by the inter-key template similarity, which had a mean of approximately 0.75 and a standard deviation of around 0.01, closely resembling random independence. Additionally, the correlation analyses between impostor and genuine cases demonstrated a distinct separation, with genuine cases approximating a value of 1.0 and impostor cases around 0.75. Furthermore, the brute-force proxy experiment indicated no overlap between the distributions of impostor and genuine instances, thereby affirming the system's resilience against guessing attacks. The concluding verification phase resulted in an AUC of approximately 0.9989 and an EER

ofaround 8.46×10^{-4}, demonstrating nearly flawless discrimination under controlled conditions.

5. Modular, Reproductible Implementation Framework: A comprehensive modular implementation was created, wherein each stage-from preprocessing (preprocessing.py) to evaluation (roc_analysis.py)-functions autonomously while integrating smoothly into the overall system. The design facilitates the straightforward replacement of different transformations, distance metrics, or datasets, thereby enhancing both reproducibility and extensibility for forthcoming research in EEG biometrics.

The collective findings indicate that EEG signals can function as effective cancellable biometric identifiers, ensuring privacy while achieving a high level of authentication accuracy. This thesis effectively addresses an important research gap that exists between biometric performance and template protection specifically within the realm of EEG-based systems.

8.2 Recommendations for Future Research

1. Multi-Session and Cross-Device Validation: The present analysis utilized a dataset derived from a single session conducted under controlled recording conditions. Subsequent research ought to assess the stability of the system over multiple sessions, while also considering variations in electrode placements and the use of diverse EEG acquisition devices, including those of consumer-grade quality. This would facilitate a comprehensive understanding of temporal robustness and cross-hardware generalization, both of which are essential for successful deployment.

2. Exploration of Alternative Cancellable Tranformations: The biohashing pipeline that was implemented integrated permutation, random projection, and quantization techniques. Exploring alternative non-invertible transforms-such as random convolutional hashing, orthogonal transformations, or chaotic mapping-may provide insights into the trade-offs associated with non-invertibility, revocability, and matching accuracy.

3. Integration of Cryptographic and Machine Learning Approaches: The integration of cancellable biometrics with homomorphic encryption or federated learning has the potential to facilitate privacy-preserving verification in distributed systems, all while avoiding the need to share raw templates. Furthermore, adaptive fusion techniques that utilize machine learning have the potential to dynamically modify score weighting (α) in order to enhance performance across varying operating conditions.

4. Attack-Resilient Evaluation Frameworks: This study focused on the aspects of revocability and resistance to brute-force attacks; however, subsequent research should consider the inclusion of adversarial attack simulations, including template inversion, GAN-based synthetic EEG reconstruction, and mimicry attacks. Assessing the system in these contexts would enhance assurance in its robustness in practical applications.

5. Expansion to Multimodal Biometric Systems: The cancellable transformation framework presented in this study has the potential to be expanded into multimodal contexts, integrating EEG with ECG, eye-tracking, or facial electromyography (EMG). Hybrid systems of this nature have the potential to improve the reliability of

authentication processes, while simultaneously distributing the associated privacy risks across different modalities.

6. Towards Real-Time Implementation and Deployment: Future research should focus on incorporating this system into a real-time authentication prototype, facilitating continuous or on-demand identity verification through the use of wearable EEG headsets. Utilizing optimized libraries such as TensorFlow Lite and ONNX Runtime for lightweight implementations may render EEG authentication practical for embedded and mobile platforms.

7. Standardization and Benchmarking: The establishment of standardized datasets, metrics, and reporting protocols for cancellable EEG biometrics is of paramount importance. This thesis advances the cause of reproducibility; however, enhanced collaboration among research groups could establish standardized benchmarks for assessing both biometric accuracy and template security in a cohesive way.

Final Remarks

In summary, this thesis illustrates that cancellable EEG biometrics can successfully achieve a balance among security, privacy, and usability. The established framework not only reproduces but also significantly enhances current research by providing a more thorough investigation of revocability, impostor analysis, and comprehensive implementation. The system demonstrates the viability of privacy-preserving EEG authentication by achieving similar performance levels between protected and unprotected templates.

This study establishes a robust basis for subsequent investigations and real-world implementations of neurobiometric security systems, contributing to the overarching objective of facilitating secure, user-focused, and ethically sound biometric technologies.

References

1. Jain, A. K., Ross, A., & Prabhakar, S. (2004). An introduction to biometric recognition. *IEEE Transactions on Circuits and Systems for Video Technology, 14*(1), 4–20.
2. Jain, A. K., Bolle, R., & Pankanti, S. (1999). *Biometrics: Personal identification in networked society.* Kluwer Academic Publishers.
3. Matyáš, V., & Říha, Z. (2002). Biometric authentication: Security and usability. In *Advanced Communications and Multimedia Security* (pp. 227–238). Springer.
4. Prabhakar, S., Pankanti, S., & Jain, A. K. (2003). Biometric recognition: Security and privacy concerns. *IEEE Security & Privacy, 1*(2), 33–42.
5. National Research Council. (2010). *Biometric recognition: Challenges and opportunities.* National Academies Press.
6. Stiles, J., & Jernigan, T. L. (2010). The basics of brain development. *Neuropsychology Review, 20*(4), 327–348.
7. Niedermeyer, E., & da Silva, F. L. (2004). *Electroencephalography: Basic principles, clinical applications, and related fields.* Lippincott Williams & Wilkins.
8. Nayak, C. S., & Anilkumar, A. C. (2025). EEG normal waveforms. In *StatPearls*. StatPearls Publishing.
9. Başar, E. (1999). *Brain function and oscillations: Principles and approaches.* Springer.
10. Armstrong, B. C., Ruiz-Blondet, M. V., Khalifian, N., Kurtz, K. J., Jin, Z., & Laszlo, S. (2015). Brainprint: Assessing the uniqueness, collectability, and permanence of ERP biometrics. *Neurocomputing, 166*, 59–67.
11. Marcel, S., & Millán, J. d. R. (2007). Person authentication using brainwaves (EEG) and maximum a posteriori model adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 29*(4), 743–752.
12. Maiorana, E., & Campisi, P. (2018). Longitudinal evaluation of EEG-based biometric recognition. *IEEE Transactions on Information Forensics and Security, 13*(5), 1123–1138.
13. Poulos, M., Rangoussi, M., Chrissikopoulos, V., & Evangelou, A. (1999). Person identification based on parametric processing of EEG. *Mathematics and Computers in Simulation, 48*(3–4), 245–261.

14. Lin, Y.-P., et al. (2011). Dry and noncontact EEG sensors for brain–computer interfaces. *IEEE Transactions on Neural Systems and Rehabilitation Engineering, 19*(5), 465–473.
15. Jalaly Bidgoly, A., Jalaly Bidgoly, H., & Arezoumand, Z. (2022). Towards a universal and privacy-preserving EEG-based authentication system. *Scientific Reports, 12*, 2531.
16. Ratha, N. K., Connell, J. H., & Bolle, R. M. (2001). Enhancing security and privacy in biometrics-based authentication systems. *IBM Systems Journal, 40*(3), 614–634.
17. Ratha, N. K., Chikkerur, S., Connell, J. H., & Bolle, R. M. (2007). Generating cancelable fingerprint templates. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 29*(4), 561–572.
18. Teoh, A. B. J., Ngo, D. C. L., & Goh, A. (2004). Biohashing: Two-factor authentication featuring fingerprint data and tokenised random number. *Pattern Recognition, 37*(11), 2245–2255.
19. Zuo, J., Ratha, N., & Connell, J. (2008). Cancelable iris biometric. In *Proceedings of the 19th International Conference on Pattern Recognition* (pp. 1–4).
20. Wang, M., Wang, S., & Hu, J. (2022). Cancellable template design for privacy-preserving EEG biometric authentication systems. *IEEE Transactions on Information Forensics and Security, 17*, 3350–3364.
21. Wang, M., Wang, S., & Hu, J. (2023). PolyCosGraph: A privacy-preserving cancelable EEG biometric system. *IEEE Transactions on Dependable and Secure Computing, 20*(5), 4258–4271.
22. Wang, M., Yin, X., & Hu, J. (2024). Cancellable deep learning framework for EEG biometrics. *IEEE Transactions on Information Forensics and Security* (early access).
23. Bajwa, G., & Dantu, R. (2016). Neurokey: Cancelable biometrics-based key generation using EEG. *Computers & Security, 62*, 95–113.
24. Damaševičius, R., Maskeliūnas, R., Kazanavičius, E., & Wozniak, M. (2018). Combining cryptography with EEG biometrics. *Computational Intelligence and Neuroscience, 2018*, Article ID 1869251.
25. Maiorana, E., La Rocca, D., & Campisi, P. (2015). Cognitive biometric cryptosystems: A case study on EEG. In *Proceedings of the IEEE International Conference on Systems, Signals and Image Processing (IWSSIP)* (pp. 125–128).
26. Goldberger, A. L., Amaral, L. A. N., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., ... & Stanley, H. E. (2000). PhysioBank, PhysioToolkit, and PhysioNet:

Components of a new research resource for complex physiologic signals. *Circulation, 101*(23), e215–e220.

27. Schalk, G., McFarland, D., & Wolpaw, J. R. (2004). BCI2000: A general-purpose brain–computer interface system. *IEEE Transactions on Biomedical Engineering, 51*(6), 1034–1043.
28. Gramfort, A., et al. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience, 7*, 267.
29. Brunner, C., Leeb, R., Müller-Putz, G. R., Schlögl, A., & Pfurtscheller, G. (2008). BCI Competition 2008 – Graz data set A. *Graz University of Technology.*
30. Yang, S., & Deravi, F. (2019). On the vulnerabilities of cancelable biometrics and their protection. *Pattern Recognition Letters, 123*, 28–35.
31. Teoh, A. B. J., & Kim, J. (2021). Cancelable biometrics: Past, present and future. *IEEE Access, 9*, 88865–88889.
32. McFarland, D. J., & Wolpaw, J. R. (2011). Brain–computer interfaces for communication and control. *Communications of the ACM, 54*(5), 60–66.
33. International Organization for Standardization. (2011). *ISO/IEC 19794-1: Information technology – Biometric data interchange formats – Part 1: Framework.*
34. International Organization for Standardization. (2017). *ISO/IEC 30107-3: Information technology – Biometric presentation attack detection – Part 3: Testing and reporting.*
35. International Organization for Standardization. (2022). *ISO/IEC 24745: Information security, cybersecurity and privacy protection – Biometric information protection.*
36. Ferreira, C., Park, A., Kietzmann, J., Demetis, D., Flostrand, A., McCarthy, I., … & Dabirian, A. (2024). cybercrime: understanding the current state of literature and issues facing cisos. It Professional, 26(2), 83-89. https://doi.org/10.1109/mitp.2024.3375571
37. Sharma, B. and Singh, P. (2019). Phishalert: an efficient phishing url detection via hybrid methodology. International Journal of Innovative Technology and Exploring Engineering, 8(12), 3936-2941. https://doi.org/10.35940/ijitee.l3455.1081219
38. Sebescen, N. and Vitak, J. (2017). Securing the human: employee security vulnerability risk in organizational settings. Journal of the Association for Information Science and Technology, 68(9), 2237-2247. https://doi.org/10.1002/asi.23851

39. Karthika, S. (2023). Cloud storage data security and integrity protection using localized and customized key manipulator (lckm) algorithm. tjjpt, 44(4), 2001-2007. https://doi.org/10.52783/tjjpt.v44.i4.1170
40. Eberz, S., Paoletti, N., Roeschlin, M., Patani, A., Kwiatkowska, M., & Martinović, I. (2017). Broken hearted: how to attack ecg biometrics, https://doi.org/10.14722/ndss.2017.23408
41. Thakur, K., & Pathan, A. S. K. (2020). Cybersecurity fundamentals: A real-world perspective. CRC Press.
42. Mu, Z., Yin, J., & Hu, J. (2018). Application of a brain-computer interface for person authentication using eeg responses to photo stimuli. Journal of Integrative Neuroscience, 17(1). https://doi.org/10.31083/jin-170042
43. Wang, M., Wang, S., & Hu, J. (2022). cancellable template design for privacy-preserving eeg biometric authentication systems. https://doi.org/10.48550/arxiv.2203.16730
44. Jain, A. biometric recognition: overview and recent advances., 13-19. https://doi.org/10.1007/978-3-540-76725-1_2
45. Mishra, A. (2010). Multimodal biometrics it is: need for future systems. International Journal of Computer Applications, 3(4), 28-33. https://doi.org/10.5120/720-1012
46. Bajaj, V. and Kaur, N. (2022). Popular physiological trait based biometric authentication systems. Ecs Transactions, 107(1), 6515-6522. https://doi.org/10.1149/10701.6515ecst
47. Mu, X., Jiang, H., Li, F., Pan, T., Xiong, G., & Ye, Y. (2024). physiological signal-based biometric identification for discovering and identifying a new user.. https://doi.org/10.3233/faia240155
48. Riley, C., Benyon, D., Johnson, G., & Buckner, K. (2010). Security in context: investigating the impact of context on attitudes towards biometric technology.. https://doi.org/10.14236/ewic/hci2010.15
49. Mayron, L., Hausawi, Y., & Bahr, G. (2013). Secure, usable biometric authentication systems., 195-204. https://doi.org/10.1007/978-3-642-39188-0_21
50. Jain, A., Hong, L., & Pankanti, S. (2000). Biometric identification. Communications of the ACM, 43(2), 90-98.
51. de Luis-García, R., Alberola-Lopez, C., Aghzout, O., & Ruiz-Alzola, J. (2003). Biometric identification systems. Signal processing, 83(12), 2539-2557.

52. Wayman, J., Jain, A., Maltoni, D., & Maio, D. (2005). An introduction to biometric authentication systems. In Biometric systems: Technology, design and performance evaluation (pp. 1-20). London: Springer London.
53. Schwartze, J., Haarbrandt, B., Fortmeier, D., Haux, R., & Seidel, C. (2014). Authentication systems for securing clinical documentation workflows. Methods of information in medicine, 53(01), 3-13.
54. Ogbanufe, O., & Kim, D. J. (2018). Comparing fingerprint-based biometrics authentication versus traditional authentication methods for e-payment. Decision Support Systems, 106, 1-14.
55. Matyáš, V., & Říha, Z. (2002, April). Biometric authentication-security and usability. In Advanced Communications and Multimedia Security: IFIP TC6/TC11 Sixth Joint Working Conference on Communications and Multimedia Security September 26–27, 2002, Portorož, Slovenia (pp. 227-239). Boston, MA: Springer US.
56. Bhattacharyya, D., Ranjan, R., Alisherov, F., & Choi, M. (2009). Biometric authentication: A review. International Journal of u-and e-Service, Science and Technology, 2(3), 13-28.
57. Mahier, J., Pasquet, M., Rosenberger, C., & Cuozzo, F. (2009). Biometric authentication. In Encyclopedia of Information Science and Technology, Second Edition (pp. 346-354). IGI Global Scientific Publishing.
58. Vegendla, A., & Sindre, G. (2019). Mitigation of cheating in online exams: Strengths and limitations of biometric authentication. In Biometric authentication in online learning environments (pp. 47-68). IGI Global.
59. Singh, G., Bhardwaj, G., Singh, S. V., & Garg, V. (2021). Biometric identification system: security and privacy concern. In Artificial intelligence for a sustainable industry 4.0 (pp. 245-264). Cham: Springer International Publishing.\
60. Prabhakar, S., Pankanti, S., & Jain, A. K. (2003). Biometric recognition: Security and privacy concerns. IEEE security & privacy, 1(2), 33-42.
61. Datta, P., Bhardwaj, S., Panda, S. N., Tanwar, S., & Badotra, S. (2020). Survey of security and privacy issues on biometric system. In Handbook of computer networks and cyber security: principles and paradigms (pp. 763-776). Cham: Springer International Publishing.
62. Campisi, P. (2013). Security and privacy in biometrics: towards a holistic approach. In Security and privacy in biometrics (pp. 1-23). London: Springer London.

63. Müller-Putz, G. R., Riedl, R., & C Wriessnegger, S. (2015). Electroencephalography (EEG) as a research tool in the information systems discipline: Foundations, measurement, and applications. Communications of the Association for Information Systems, 37(1), 46.
64. Tyner, F. S., & Knott, J. R. (1983). Fundamentals of EEG Technology: Basic concepts and methods (Vol. 1). Lippincott Williams & Wilkins.
65. Singh, Y. N., Singh, S. K., & Ray, A. K. (2012). Bioelectrical signals as emerging biometrics: Issues and challenges. International Scholarly Research Notices, 2012(1), 712032.
66. Yang, S., & Deravi, F. (2017). On the usability of electroencephalographic signals for biometric recognition: A survey. IEEE Transactions on Human-Machine Systems, 47(6), 958-969.
67. Maiorana, E., La Rocca, D., & Campisi, P. (2015). On the permanence of EEG signals for biometric recognition. IEEE Transactions on Information Forensics and Security, 11(1), 163-175.
68. Bashar, K. (2018, October). ECG and EEG based multimodal biometrics for human identification. In 2018 IEEE international conference on systems, man, and cybernetics (SMC) (pp. 4345-4350). IEEE.
69. Jayarathne, I., Cohen, M., & Amarakeerthi, S. (2017, November). Survey of EEG-based biometric authentication. In 2017 IEEE 8th International Conference on Awareness Science and Technology (iCAST) (pp. 324-329). IEEE.
70. Alsaadi, I. M. (2021). Study on most popular behavioral biometrics, advantages, disadvantages and recent applications: A review. Int. J. Sci. Technol. Res, 10(1), 15-22.
71. Chan, H. L., Kuo, P. C., Cheng, C. Y., & Chen, Y. S. (2018). Challenges and future perspectives on electroencephalogram-based biometrics in person recognition. Frontiers in neuroinformatics, 12, 66.
72. Zhao, H., Chen, Y., Pei, W., Chen, H., & Wang, Y. (2021). Towards online applications of EEG biometrics using visual evoked potentials. Expert Systems with Applications, 177, 114961.
73. Yang, S., & Deravi, F. (2017). On the usability of electroencephalographic signals for biometric recognition: A survey. IEEE Transactions on Human-Machine Systems, 47(6), 958-969.
74. Teoh, A. B., Kuan, Y. W., & Lee, S. (2008). Cancellable biometrics and annotations on biohash. Pattern recognition, 41(6), 2034-2044.

75. Kavati, I., Kumar, G. K., Gopalachari, M. V., Babu, E. S., Cheruku, R., & Reddy, V. D. (2021, December). Non-invertible cancellable template for fingerprint biometric. In International Conference on Hybrid Intelligent Systems (pp. 615-624). Cham: Springer International Publishing.
76. Kaur, H., & Khanna, P. (2017, October). Non-invertible biometric encryption to generate cancelable biometric templates. In Proceedings of the world congress on engineering and computer science (Vol. 1, pp. 1-4).
77. Patel, V. M., Ratha, N. K., & Chellappa, R. (2015). Cancelable biometrics: A review. IEEE signal processing magazine, 32(5), 54-65.
78. Bernal-Romero, J. C., Ramirez-Cortes, J. M., Rangel-Magdaleno, J. D. J., Gomez-Gil, P., Peregrina-Barreto, H., & Cruz-Vega, I. (2023). A review on protection and cancelable techniques in biometric systems. Ieee Access, 11, 8531-8568.
79. Maurya, S., Bharti, J., & Roy, B. (2024, February). A survey on cancellable biometrics techniques. In 2024 IEEE International Students' Conference on Electrical, Electronics and Computer Science (SCEECS) (pp. 1-7). IEEE.
80. Teoh, A. B., Kuan, Y. W., & Lee, S. (2008). Cancellable biometrics and annotations on biohash. Pattern recognition, 41(6), 2034-2044.
81. Nagar, A., Nandakumar, K., & Jain, A. K. (2010, January). Biometric template transformation: a security analysis. In Media Forensics and Security II (Vol. 7541, pp. 237-251). SPIE.
82. Sutcu, Y., Rane, S., Yedidia, J. S., Draper, S. C., & Vetro, A. (2008, June). Feature transformation of biometric templates for secure biometric systems based on error correcting codes. In 2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (pp. 1-6). IEEE.
83. Ailon, N., & Chazelle, B. (2009). The fast Johnson–Lindenstrauss transform and approximate nearest neighbors. SIAM Journal on computing, 39(1), 302-322.
84. Kenthapadi, K., Korolova, A., Mironov, I., & Mishra, N. (2012). Privacy via the johnson-lindenstrauss transform. arXiv preprint arXiv:1204.2606.
85. Bajwa, G., & Dantu, R. (2016). Neurokey: Towards a new paradigm of cancelable biometrics-based key generation using electroencephalograms. Computers & security, 62, 95-113.
86. Lumini, A., & Nanni, L. (2007). An improved biohashing for human authentication. Pattern recognition, 40(3), 1057-1065.
87. Kong, A., Cheung, K. H., Zhang, D., Kamel, M., & You, J. (2006). An analysis of biohashing and its variants. Pattern recognition, 39(7), 1359-1368.

88. Matou≈°ek, J. (2008). On variants of the Johnson-Lindenstrauss lemma. Random Structures & Algorithms, 33(2), 142-156.
89. Teoh, A. B., Goh, A., & Ngo, D. C. (2006). Random multispace quantization as an analytic mechanism for biohashing of biometric and random identity inputs. IEEE transactions on pattern analysis and machine intelligence, 28(12), 1892-1901.
90. Wang, M., Yin, X., & Hu, J. (2024). Cancellable deep learning framework for EEG biometrics. IEEE Transactions on Information Forensics and Security, 19, 3745-3757.
91. Fidas, C. A., & Lyras, D. (2023). A review of EEG-based user authentication: trends and future research directions. IEEE Access, 11, 22917-22934.
92. Bernal-Romero, J. C., Ramirez-Cortes, J. M., Rangel-Magdaleno, J. D. J., Gomez-Gil, P., Peregrina-Barreto, H., & Cruz-Vega, I. (2023). A review on protection and cancelable techniques in biometric systems. Ieee Access, 11, 8531-8568.
93. Rashid, M., Sulaiman, N., PP Abdul Majeed, A., Musa, R. M., Ab. Nasir, A. F., Bari, B. S., & Khatun, S. (2020). Current status, challenges, and possible solutions of EEG-based brain-computer interface: a comprehensive review. Frontiers in neurorobotics, 14, 25.
94. Manisha, & Kumar, N. (2020). Cancelable biometrics: a comprehensive survey. Artificial Intelligence Review, 53(5), 3403-3446.
95. Politou, E., Alepis, E., & Patsakis, C. (2018). Forgetting personal data and revoking consent under the GDPR: Challenges and proposed solutions. Journal of cybersecurity, 4(1), tyy001.
96. Li, Y., Yang, H., Li, J., Chen, D., & Du, M. (2020). EEG-based intention recognition with deep recurrent-convolution neural network: Performance and channel selection by Grad-CAM. Neurocomputing, 415, 225-233.
97. Zhao, H., Chen, Y., Pei, W., Chen, H., & Wang, Y. (2021). Towards online applications of EEG biometrics using visual evoked potentials. Expert Systems with Applications, 177, 114961.
98. Boulgouris, N. V., Plataniotis, K. N., & Micheli-Tzanakou, E. (Eds.). (2009). Biometrics: theory, methods, and applications. John Wiley & Sons.
99. Salama, G. M., El-Gazar, S., Omar, B., & Hassan, A. A. (2024). Multimodal cancelable biometric authentication system based on EEG signal for IoT applications. Journal of Optics, 53(3), 1839-1853.

100. Wang, M., Wang, S., & Hu, J. (2022). PolyCosGraph: A privacy-preserving cancelable EEG biometric system. IEEE Transactions on Dependable and Secure Computing, 20(5), 4258-4272.