

Advancing Brainwave-Based Biometrics: A Large-Scale, Multi-Session Evaluation

Matin Fallahi*, Patricia Arias-Cabarcos^{§a}, Thorsten Strufe*

*KASTEL, Karlsruhe Institute of Technology (KIT), Germany

Email: {matin.fallahi, strufe}@kit.edu [§]Paderborn University, Germany

Email: pac@mail.upb.de

Abstract—The field of brainwave-based biometrics has gained attention for its potential to revolutionize user authentication through hands-free interaction, resistance to shoulder surfing, continuous authentication, and revocability. However, current research often relies on single-session or limited-session datasets with fewer than 55 subjects, raising concerns about the generalizability of the findings. To address this gap, we conducted a large-scale study using a public brainwave dataset comprising 345 subjects and over 6,007 sessions (an average of 17 per subject) recorded over five years using three headsets. Our results reveal that deep learning approaches significantly outperform hand-crafted feature extraction methods. We also observe Equal Error Rates (EER) increases over time (e.g., from 6.7% after 1 day to 14.3% after a year). Therefore, it is necessary to reinforce the enrollment set after successful login attempts. Moreover, we demonstrate that fewer brainwave measurement sensors can be used, with an acceptable increase in EER, which is necessary for transitioning from medical-grade to affordable consumer-grade devices. Finally, we compared our results to prior work and existing biometric standards. While our performance is on par with or exceeds previous approaches, it still falls short of industrial benchmarks. Based on the results, we hypothesize that further improvements are possible with larger training sets. To support future research, we have open-sourced our analysis code.

Index Terms—Brainwave Authentication, EEG Authentication, Biometrics

I. INTRODUCTION

Authentication is the cornerstone of securing digital interactions. While passwords have long been the dominant method, biometrics are increasingly preferred. According to the 2024 FIDO Survey¹ of 10,000 consumers, 28% preferred biometrics for signing into online accounts, apps, and smart devices, making it the most popular choice. Beyond established methods—such as fingerprint and face recognition—emerging biometric technologies are being explored to enhance usability and security [1], [2]. Brainwave authentication has gained attention for its hands-free operation, which enhances usability in devices like XR systems [3], [4]. It offers resistance to shoulder surfing by relying on distinct neural activity patterns rather than observable actions or physical characteristics [5]. Additionally, brainwave-based systems can enable continuous authentication, improving security by monitoring user identity

in active tracking [6]. Furthermore, research indicates that in cases of brainwave sample leakage, changing the stimulus can effectively enable template renewability and revocability, thereby maintaining security [3]. This contrasts with conventional biometric authentication, where compromised raw templates typically cannot be revoked or updated. However, brainwave authentication is still in its early stages, and both permanence and performance remain significant challenges.

Prior studies on brainwave authentication mainly rely on limited public datasets or small, privately collected datasets. Consequently, they either focus on single-session datasets [7], [8], [5], [9] or small-scale multi-session datasets [10], [11], [12], [13]. While single-session studies completely overlook session and temporal effects [14], [15], multi-session studies remain limited to a maximum of 55 subjects [13] and 270 sessions in total [16] (six sessions per subject) so far. These limitations increase the risk of overfitting recognition models and limit the generalizability of results, making it difficult to identify key factors that affect long-term reliability. Importantly, existing multi-session studies rely on a single EEG device, which prevents assessment of hardware variability, although users may access the same profile from different devices in real-world settings. Moreover, the lack of publicly accessible code and data hinders reproducibility and further development, in contrast to the face recognition field, which benefits from open data and collaborative research.

To address this gap and better understand how different parameters affect the performance of brainwave authentication over time, we analyzed the recently published PEERS dataset [17], originally designed for large-scale, longitudinal studies of memory-related electrophysiology. The dataset includes recordings from 345 subjects and 6,007 sessions collected over five years using three different headsets. This dataset contains more than six times the number of subjects and 22 times the number of sessions compared to previous brainwave authentication evaluations. Our contributions are:

- **Benchmark (Sec IV-A):** To identify the best pipeline for multi-session brainwave authentication, we conducted a comprehensive evaluation of hand-crafted and deep learning-based feature extraction methods. The results showed that deep learning methods reduced the error rate by up to two times.
- **Analysis (Sec IV-B):** To better understand the parameters influencing performance, we investigated the effects of test size, session intervals, and device type. Our analysis

^aCurrent Affiliation: European Commission, Joint Research Centre (JRC), Ispra, Italy.

¹<https://fidoalliance.org/wp-content/uploads/2024/10/Barometer-Report-2024-Oct-29.pdf>

revealed several key findings. First, a larger number of subjects in the evaluation set does not necessarily increase authentication error, while smaller datasets lead to higher uncertainty due to limited observations. Second, authentication errors increase over time by 100% after one year. This suggests that updating the enrollment dataset after successful logins may help reduce long-term error rates. Third, visualization of the embedding space shows that samples from the same session and recording device are relatively close. This highlights the influence of sessions and devices on the learning process, likely due to environmental noise, device setup, and physiological or behavioral variations over time. Finally, the analysis of the three devices used for data collection indicates that cross-device authentication is feasible when the feature extractor is trained on subjects who have collected data with both devices.

- **Enhancing Performance (Sec IV-C):** We investigated the trade-off between verification time and performance. For example, by increasing the verification time from 1 to 4 seconds, the error rate decreased by approximately 40%. Additionally, we examined the trade-off between enrollment sessions and error rates, observing that after the second enrollment session, the improvement in error rate diminished.
- **Channel Reduction (Sec IV-C):** To evaluate the system's performance under conditions resembling real-world scenarios, we reduced the number of channels in the PEERS dataset—collected using medical-grade devices—from 93 to 14, 7, and 4. These reductions correspond to sensor placements commonly found in consumer-grade EEG devices. The results showed a relatively small increase in error compared to the notable reduction in channels.
- **Industrial Standard Comparison (Sec IV-E):** For the first time, we compare brainwave authentication performance against international biometric standards. At required security levels, the false rejection rate (FRR) exceeds acceptable thresholds. However, our results show a logarithmic relationship between error rate and training subjects, suggesting that with more training data, these standards may be met at a practical FRR.
- **Open Sourced Code:** We published our source code, which, along with the public dataset, makes results easily reproducible and facilitates future error rate reductions by other researchers ².

II. BACKGROUND AND RELATED WORK

EEG measures electrical activity produced by neurons firing in the human brain, typically recorded at the scalp using specialized sensors [18], [19]. German physiologist and psychiatrist Hans Berger recorded the first human EEG in 1924 for medical purposes, which remains one of its primary applications [20], [21], [22]. Beyond medical use, EEG technology has advanced to support Brain-Computer Interfaces (BCIs), systems that enable direct communication between the brain and external devices [23]. Recently, EEG

data has also been found to include unique features linked to user identity, making it a promising method for user authentication [24]. Brainwave based authentication is hands-free, resistant to shoulder surfing, revocable, renewable, and supports continuous authentication.

Biometric authentication modalities should meet certain requirements to enable widespread usage. These requirements include universality (the modality must be present in all individuals), distinctiveness (it must uniquely identify an individual), permanence (it should remain consistent over time), and collectability (it must be measurable with available technology). Additional factors include performance (accuracy, speed, and resource efficiency), acceptability (user willingness to utilize the system), and circumvention resistance (difficulty in forging or bypassing) [25].

EEG is universal because each person's brain produces EEG signals [26]. These signals exhibit distinctiveness due to individual brainwave patterns, making them uniquely identifiable [27]. Advances in technology are rapidly improving collectability, with new EEG headsets entering the market annually, offering better signal quality and usability [28]. However, without real-time implementation, it is difficult to comprehensively assess acceptability and circumvention resistance, though some research attempts have been made in this direction [29], [30].

Currently, the primary focus of research is on performance and permanence, as these are challenging aspects of EEG-based authentication. The two are closely interrelated, with performance metrics reflecting the system's ability to maintain consistent accuracy over time. Factors such as hair growth, slight variations in sensor placement on the scalp, brain states (e.g., stress), and environmental conditions (e.g., noise) can significantly impact both performance and permanence in EEG-based authentication. However, most research studies evaluate the performance of brainwave authentication using single-session datasets (cf. [7], [8], [5], [9]), which neglect session invariance (permanence) completely. This limitation raises concerns that models may overfit to session-specific characteristics rather than capturing subject-unique features. Additionally, some studies (cf. [13], [14]) consider short time intervals, limited to periods of a few days only, which may only partially account for the aforementioned issues. While a limited number of papers explore time intervals spanning more than a week, these studies (cf. [10], [11], [16], [12]) rely on non-public datasets and omit sharing their source code, hindering reproducibility and further research. Moreover, such studies have been conducted with a relatively small sample size and limited sessions (up to 54 subjects [13] and 270 sessions [16]), which increases the risk of overfitting and restricts the generalizability of the results.

Therefore, a more comprehensive investigation of performance and permanence is needed using larger multi-session datasets to avoid overfitting. This would enable the examination of key parameters affecting performance over time and provide insights into improving brainwave authentication. Additionally, identifying the main challenges and ensuring reproducibility through open-source code sharing are essential for advancing the field.

²<https://github.com/kit-ps/NeuroShield/>

TABLE I: Demographic and session distribution of the dataset.

Metric	Value
Unique subjects	345
Unique sessions	6007
Average sessions per subject (\pm SD)	17.41 \pm 10.48
Female	161 (46.7%)
Male	136 (39.4%)
Unspecified sex	48 (13.9%)
Age, years (mean \pm SD)	27.0 \pm 15.9
Age, years (range)	17.0–85.0

III. METHOD

Biometric systems, including brainwave authentication, involve four key stages: data collection, preprocessing, feature extraction, and feature comparison. Like other biometric systems, they operate in two phases: enrollment and verification. In the enrollment phase, brainwave signals are collected from the user to generate a reference profile, known as a template. During verification, performed after enrollment, brainwave signals are captured again and compared with the reference profile to confirm the user's identity.

A. Dataset Description

Our study utilizes the Penn Electrophysiology of Encoding and Retrieval Study (PEERS) dataset [17], provided by the University of Pennsylvania. Collected between 2010 and 2020, the dataset focuses on exploring EEG correlates of memory processes, particularly during tasks involving word memorization. In these tasks, participants are shown a sequence of words and asked to memorize them as stimulus. Stimulus presentation was synchronized with EEG data collection using event markers. We used the timestamps of word presentations on the screen to segment the EEG data into individual samples. EEG data, recorded as time series of scalp voltage signals, are collected during the sessions at sampling rate of 500 Hz. The dataset includes approximately 8.7 TB of EEG data from 345 subjects, each with more than one session, totaling over 6007 sessions. The participants range in age from 17 to 85 years, with an average age of 26.98 ± 15.86 years. The dataset includes 161 female and 136 male participants, while gender information for 48 subjects is unavailable (cmp. Table I). The PEERS dataset is particularly suited for our research as it includes multiple sessions per subject, enabling the investigation of inter-session variability, a critical factor for ensuring the robustness of real-world biometric systems. Data collection was conducted using three distinct medical-grade headsets: the 129-channel Geodesic Sensor Net (GSN 200 model)², the 129-channel HydroCel Geodesic Sensor Net,² and the 128-channel BioSemi³ headcap using the Biosemi ActiveTwo acquisition system. These diverse devices further enhance the dataset's relevance by reflecting real-world variability in EEG acquisition systems.

³<https://www.biosemi.com/headcap.htm>

²https://med.stanford.edu/content/dam/sm/lucasmri/documents/16_0824_EGI_geodesic_sensor_net.pdf

B. Preprocessing

As the channel layout of the different headsets varied, we unified them by mapping the channels to a well-known extended version⁴ of the ten-twenty electrode system [31], which includes 93 channels. Subsequently, the EEG data channels were reordered for each recording to align with the expected order during evaluation across all three devices. Following this, and in line with best practices for preprocessing in brainwave authentication [14], [5], [32], [12], a common average reference (CAR) was applied to reduce spatially correlated noise and improve signal quality. Next, we applied a bandpass filter between 1.0 Hz and 50.0 Hz using a zero-phase finite impulse response (FIR) filter with default parameters, as implemented in MNE-Python [33], to focus on frequency bands that contain EEG data relevant for authentication [34]. As the 50 Hz power-line frequency lies close to the upper end of the preserved frequency band, we additionally applied a notch filter at 50 Hz to suppress potential residual line noise⁵. After filtering, the data was normalized using a robust normalization approach, where the median was subtracted and the result divided by the interquartile range (IQR), mitigating the influence of outliers while preserving the temporal structure of the signal. The data then consisted of 1-second non-overlapping samples aligned to stimulus timestamps. Each sample was stored together with metadata including session date and hardware type.

C. Feature Extraction

After preprocessing, each sample contains 46,500 values (93 channels \times 500 Hz). Due to the high dimensionality and variability of EEG data, direct comparison is not feasible. Therefore, feature extraction is necessary to capture subject-specific information across sessions. Two main approaches exist: the first is handcrafted feature extraction, which uses predefined formulas based on expert knowledge and does not require training data. The second approach uses deep learning techniques, which require a separate training set to learn feature representations. In this method, the feature extractor is trained on labeled data to produce representations that generalize to new subjects.

handcrafted Feature Extraction: We use Power Spectral Density (PSD) to transform EEG signals from the time to frequency domain, capturing patterns linked to brain dynamics. PSD is commonly used in brainwave-based authentication systems [35], [5], [14], applying Fourier transformation to quantify signal power across frequency components [36]. We also extract AutoRegressive (AR) model coefficients, which capture temporal dependencies by modeling each point as a linear combination of previous values. AR features provide compact representations for authentication [5], [37]. Finally, we combine PSD and AR features to incorporate both frequency and temporal information.

⁴https://mne.tools/stable/auto_tutorials/intro/40_sensor_locations.html

⁵Residual power-line interference in EEG recordings typically occurs at 50 Hz (e.g., most of Europe) or 60 Hz (e.g., North America). We bandpass-filtered the data between 1 and 50 Hz and then applied notch filters at 50 Hz and 60 Hz to suppress residual line noise at both mains frequencies; because 60 Hz lies outside the preserved band, its notch filter effect on the data is negligible.

Deep learning feature extraction: We applied metric learning, which trains a network to map inputs into a feature space where samples with the same label are close and those with different labels are far [38]. This method uses an encoder to generate embeddings and a loss function that enforces distances based on labels. Metric learning losses are distance-based or softmax-based. Distance-based losses, operating on pairs, triplets, or mined batches, optimize embedding distances directly [39]. Softmax-based losses encourage class separation by clustering embeddings and generally need larger datasets to generalize well [39].

For softmax-based approaches, we selected two loss functions: (1) SoftTripleLoss, which addresses intra-class variations by introducing multiple learnable centers per class; and (2) ArcFaceLoss, which applies angular margin constraints to improve inter-class separability and robustness to noise. For distance-based approaches, we selected three loss functions: (1) Triplet Loss [40], which minimizes the distance between an anchor and a positive sample while maximizing the distance from a negative sample; (2) LiftedStructureLoss [41], which uses all positive and negative pairs within a batch and penalizes violations of distance constraints; and (3) SupConLoss (Supervised Contrastive Loss) [42], which extends contrastive learning by using class labels to group embeddings of the same class together and separate embeddings of different classes. This supervision helps SupConLoss align embeddings more effectively across the batch.

For Triplet Loss, we adopt semi-hard online triplet mining [40], which is recommended to emphasize informative (non-trivial) triplets and improve convergence. LiftedStructureLoss and SupConLoss are pair-based (and not triplet-based); accordingly, we employ Multi-Similarity mining [43] to select hard positive and negative pairs based on cosine similarity. ArcFaceLoss and SoftTripleLoss are proxy-based margin losses (softmax-based) that operate on class centers and therefore do not require mining, as they do not consume explicit pair or triplet indices. Samples are drawn across all sessions in the training set, and positive pairs include both within-session and cross-session pairs of the same subject. This allows the model to minimize both intra-session and inter-session embedding distances, supporting continuous authentication (stability within sessions) as well as multi-session authentication (robustness across sessions).

For each loss function, different encoder architectures were explored to optimize feature extraction. Specifically, we implemented six backbone EEG models — ResNet1D [44], ShallowNet [45], DeepConvNet [45], EEGNet [46], LSTM [47], and GRU [47] — as encoder networks. These models were selected to cover a diverse range of architectures, including convolutional and recurrent networks, and to balance model complexity and representational capacity for EEG signal processing. Details of the neural network architecture are in Appendix B.

D. Feature Comparison

After reducing brainwave samples to lower dimensional embeddings through feature extraction, the extracted features

should be compared, and a similarity score calculated to evaluate verification attempts. Two main approaches can be used. The first approach applies distance metrics; we select Euclidean, Cosine, and Manhattan distances because Euclidean captures absolute differences, Cosine measures angular similarity independent of magnitude, and Manhattan is more robust to outliers [4], [48]. The second approach trains binary and one-class classifiers to distinguish between positive and negative samples for each subject. However, binary classifiers require negative samples, which should not come from subjects included in the evaluation. We employ several binary classifiers, including logistic regression, linear discriminant analysis (LDA), support vector machine (SVM) with linear kernel, stochastic gradient descent (SGD) classifier with radial basis function (RBF) kernel approximation, and random forests. One-class SVM is also tested, although it is less suitable when negative classes are available; it is included due to its use in related work [12]. As shown in face recognition, distance-based approaches are expected to perform better if sufficient data is available to train the feature extractor [40], since metric learning aims to minimize distances between similar samples. The most suitable comparison method will be selected via systematic hyperparameter tuning, as described in the Section IV-A.

E. Threat model

Since EEG signals are not externally observable, shoulder surfing is not possible, and the adversary cannot impersonate a victim. Accordingly, authentication studies typically assume that adversaries attempt to gain access using their own brainwave recordings [8], [5]. Meanwhile, our goal is to investigate the reliability of brainwave-based authentication over time (multi-session setting), and to identify potential bottlenecks affecting this reliability. The low reliability of brainwave signals across sessions contributes to a high false acceptance rate, which poses a significant security threat and can be assessed using zero-effort attacks [49]. A zero-effort attack in biometric authentication occurs when an adversary uses their own biometric sample in an attempt to gain unauthorized access to another individual's account. In this context, the security objective of the biometric system is to minimize the probability that the authentication algorithm incorrectly verifies the adversary's identity as the legitimate user. Specifically, the likelihood of successful verification of a falsely claimed identity, based on the adversary's sensor readings, should be negligible. Therefore, zero-effort attacks are commonly used in studies on multi-session reliability involving brainwave data [12], [16], [10], [11], [14], [13].

F. Performance Metrics

We used the Equal Error Rate (EER) as a summary metric, indicating the point where the False Acceptance Rate (FAR) and the False Rejection Rate (FRR) are equal. The FAR represents the proportion of unauthorized attempts incorrectly accepted by the system, while the FRR refers to the proportion of legitimate attempts incorrectly rejected. EER has been widely used in biometric authentication papers [25], [50] to

report results and ensure comparability. However, in practical usage, it is preferred to keep FAR at very low rates to ensure system security, even at the cost of a higher FRR. Therefore, we also report the FRR at a FAR threshold of 1% for comparison with other brainwave authentication studies [5], [9], [51], as well as at 0.1%⁶ and 0.01%⁷ to align with biometric standards.

TABLE II: Results of hyperparameter trials on the validation set (15 subjects): number of pipeline occurrences within the top 1,000 trials, average EER computed for pipelines appearing in the top 1,000, and best EER across all 10,000 trials.

	Pipeline	Count	Avg	Best
Feature Extraction	ResNet1D + SupConLoss	98	10.20	8.05
	ResNet1D + ArcFaceLoss	65	10.65	8.09
	ResNet1D + LiftedStructureLoss	107	10.78	8.39
	ShallowNet + SoftTripleLoss	77	11.13	8.97
	ResNet1D + TripletMarginLoss	99	11.36	9.20
	ShallowNet + LiftedStructureLoss	34	10.85	9.25
	ShallowNet + SupConLoss	77	11.50	9.38
	ResNet1D + SoftTripleLoss	54	11.79	9.66
	ShallowNet + ArcFaceLoss	72	12.12	9.69
	DeepConvNet + SoftTripleLoss	90	12.19	10.23
	GRU + SoftTripleLoss	36	11.43	10.35
	LSTM + SoftTripleLoss	28	11.87	10.72
	DeepConvNet + ArcFaceLoss	26	12.10	11.11
	ShallowNet + TripletMarginLoss	47	12.30	11.50
	GRU + SupConLoss	28	12.62	12.04
	DeepConvNet + SupConLoss	10	12.69	12.34
	LSTM + LiftedStructureLoss	9	12.71	12.59
	DeepConvNet + LiftedStructureLoss	7	12.68	12.63
	GRU + LiftedStructureLoss	8	12.69	12.65
	EEGNet + LiftedStructureLoss	9	12.93	12.75
	GRU + ArcFaceLoss	19	12.89	12.79
	EEGNet + ArcFaceLoss	0	-	13.18
	LSTM + ArcFaceLoss	0	-	13.21
	EEGNet + SupConLoss	0	-	13.36
	LSTM + SupConLoss	0	-	13.61
	DeepConvNet + TripletMarginLoss	0	-	13.88
	EEGNet + SoftTripleLoss	0	-	14.49
	EEGNet + TripletMarginLoss	0	-	15.00
	LSTM + TripletMarginLoss	0	-	15.55
	GRU + TripletMarginLoss	0	-	18.74
PSD + AR	0	-	19.23	
AR	0	-	19.43	
PSD	0	-	29.34	
Feature Comparison	Euclidean Distance	86	11.18	8.05
	Cosine Distance	168	10.75	8.05
	Manhattan Distance	105	10.80	8.17
	Logistic Regression	171	11.70	8.36
	LDA	222	11.67	8.56
	Linear SVM	106	11.95	8.86
	SGD RBF Kernel	80	11.74	9.58
	Random Forest	62	12.45	11.49
	One-Class SVM	0	-	21.68

IV. RESULTS AND DISCUSSION

In order to evaluate brainwave authentication performance over time, we first select a pipeline comprising a feature extractor and a comparison metric through hyperparameter tuning. We then investigate important parameters such as test size, time interval, recording hardware, and subject embeddings visualized as distinct clusters through t-SNE plots. Next,

to move toward consumer-grade devices, we analyze performance with a reduced number of channels and contextualize our findings with state-of-the-art methods and authentication standards to identify potential directions for future research.

A. Benchmarking Feature Extraction and Compression

After introducing dataset and preprocessing methods we now want to turn to selecting suitable feature extraction and compression techniques, that are essential to complete the authentication pipeline.

1) *Experiment setup*: We hypothesize that the proposed method supports both inter-subject and inter-session generalization. Inter-subject generalization refers to the ability of a model trained on a subset of users to perform well on previously unseen users. Inter-session generalization refers to the ability of the model to embed data from different sessions of the same user close to each other in the feature space. Therefore, this is reflected by the data split and evaluation protocol. We divided the 345 available subjects into non-overlapping sets: 230 subjects for training the deep learning-based feature extractor, 15 for validation, and 100 randomly selected subjects for evaluation, on whom system performance is measured. In all experiments, unless specified otherwise, the first session of each subject in the evaluation set was used for enrollment, and the remaining sessions of the same subject for verification. This split ensures inter-subject evaluation, since no subject appears in more than one of the training, validation, and evaluation sets, and shows that the model does not require retraining when new users are added. Additionally, separating enrollment and verification by session enables the assessment of inter-session generalization.

2) *Hyperparameter tuning*: From each combination of model architectures and loss functions (6x5) mentioned in the methods (Section III), we trained a model using hyperparameter tuning on the optimizer, learning rate, batch size, and embedding dimension based on the training set, guided by performance on the validation set. The entire process took approximately 144 hours on four NVIDIA GeForce RTX 4090 GPUs.

After obtaining the feature extraction models, we conducted general hyperparameter tuning to select the best combination of feature extraction and feature comparison methods. The tuning process was guided by the validation set and employed a random search over 10,000 trials. To provide a clearer overview, Table II reports the results for the top 10% of trials on the validation set, thereby reducing the influence of high-error outliers on the average results. The hyperparameter tuning identified the “ResNet1D + SupConLoss + Euclidean distance” pipeline as the best configuration. This pipeline was then evaluated on a test set of 100 subjects, resulting in an EER of 10.28% ± 6.3%. *For the remainder of this paper, we use this pipeline as the default authentication pipeline.*

3) *Feature Extraction Methods*: Table II shows that our dataset is sufficient for deep learning approaches to outperform handcrafted features, which result in significantly higher error. Among the encoders, ResNet1D consistently achieved the top three ranks, followed by ShallowNet, which outperformed the remaining models. This suggests that convolutional

⁶<https://pages.nist.gov/800-63-3/sp800-63b.html>

⁷<https://pages.nist.gov/800-63-4/sp800-63b.html>

networks generally perform better than recurrent architectures. Among the convolutional models, simpler architectures like ShallowNet may benefit from reduced overfitting, while ResNet1D leverages residual connections to achieve better performance than models such as DeepConvNet. Although EEG-Net is relatively shallow, its specialized filters for frequency and spatial information did not translate into strong performance in our evaluation. Regarding loss functions, although SoftTripleLoss—with its multiple learnable centers per class—performed best with shallower and recurrent backbones (ShallowNet, DeepConvNet, LSTM, GRU), ResNet1D combined with SupConLoss achieved the highest overall performance across all architectures. SupConLoss combines the strengths of distance-based and softmax-based approaches by leveraging all positive and negative pairs in the batch, guided by class labels. This batch-level supervision enforces strong intra-class cohesion and inter-class separation, which ResNet1D—with its higher capacity—can fully exploit to learn highly discriminative EEG embeddings

4) *Comparison Methods:* The results (Table II) suggest that, in general, distance-based functions outperform shallow classifier approaches for feature comparison. This is not surprising, as loss functions in metric learning are designed to minimize distances between embeddings, typically using cosine or Euclidean distance. In contrast, shallow classifiers either require negative samples for training or, in the case of one-class classification, often show relatively poor performance due to the instability associated with one-class models. This is also why state-of-the-art face recognition methods use distance functions for feature comparison.

B. Detailed Performance Analysis

With the best-performing pipeline identified, the EER remains significantly lower than the 50% chance level but still reaches 10.28%, which is higher than the performance required for real-world deployment. This result may indicate that brainwaves are not sufficiently unique for authentication or that other underlying factors need to be understood to develop effective solutions. Therefore, we aim to investigate which parameters influence multi-session authentication more deeply by addressing the following research questions:

- 1) How does altering the number of participants in an EEG-based authentication evaluation set affect the overall comparability of the evaluation? To investigate this, we analyze how varying the number of evaluation subjects from 2 to 100 affects the EER.
- 2) How does EEG-based authentication performance change over time, and what does this suggest about brainwave pattern stability? To explore this, we examine how EER is influenced by varying verification and enrollment intervals.
- 3) How effectively are subjects clustered based on their embedding space derived from the feature extractor? This explores the visualization of embedding space.
- 4) How does different recording hardware affect the performance of brainwave-based authentication, and what are the best practices for ensuring reliability? This question

is investigated by measuring the EER on hardware-specific and generalized models as well as evaluating cross-hardware authentication scenario.

1) *Effect of the Number of Evaluation Subjects on EER:* When comparing the results of different experiments using metrics such as EER, number of subjects in the evaluation is typically reported as metadata. A larger test set generates more impostor scores in zero-effort scenarios, making it valuable to investigate the relationship between the number of subjects in the evaluation set and the resulting error. This analysis provides insights for better comparisons of results across different studies or our follow-up experiments with smaller test sets, such as when evaluating performance using headset-specific recordings. To achieve this, we randomly select N subjects, where N ranges from 2 to 99 (starting at $N = 2$ since at least one genuine user and one impostor are required). For each value of N , we compute the EER, repeat the procedure 50 times, and report the average EER per N .

Figure 2 presents the results of the experiment. The analysis revealed no significant relationship between the number of subjects and the EER. These findings align with observations in face recognition systems. For example, Friedman et al. [52] investigated how test size affects identification and authentication tasks. They reported a linear decline in identification performance with the logarithm of enrollment size but found that the EER for authentication tasks remained stable. This stability occurs because the EER is based on the Receiver Operating Characteristic (ROC) curve, which balances the FAR and FRR derived from impostor and genuine score distributions. While larger subject pools better define these distributions, they do not significantly influence the ROC. Therefore, we can conclude that having a higher number of subjects in the test set does not necessarily make authentication more challenging. However, both scenarios indicate that test sets with a small number of participants exhibit higher deviations in EER, which can lead to unreliable or misleading results. This underscores the importance of using larger test sets to obtain more dependable evaluations.

2) *Correlation Between EER and Time Interval:* To investigate permanence, we first visualize the distribution of evaluation session intervals in Figure 3. The test data includes a total of 1,699 unique sessions available for 100 evaluation subjects, where the first session (or the oldest session) per subject is considered as the enrollment session (100 sessions). The results show that the highest concentration of sessions occurred within the first week, with 129 sessions (8%) recorded. By the end of the first month (weeks 0 to 3), this number increased to 560 sessions (35%), indicating that one-third of all sessions were conducted in the first month. The second month (weeks 4 to 7) accounted for 398 sessions (24.8%), while the third month (weeks 8 to 11) added 201 sessions (12.5%). Beyond the third month, session frequency declined further, with 182 sessions (11.3%) recorded from Month 4 to the end of the first year. In the second year, session numbers were 89 sessions (5.5%), and remained similar at 71 sessions (4.4%) in the third year. The remaining years combined accounted for 4.8%.

To investigate the correlation between time intervals between sessions and authentication error, we extracted all

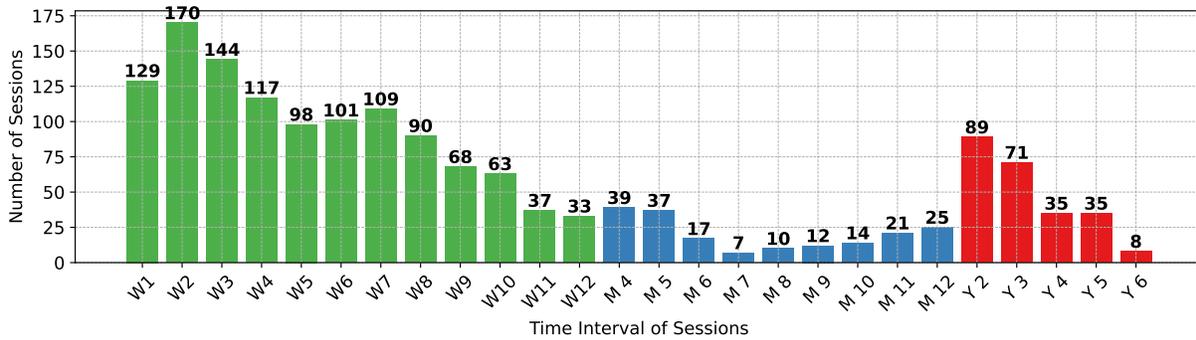


Fig. 1: Number of sessions in the evaluation dataset post-enrollment. W , M , and Y denote week, month, and year. Labels (e.g., $W3$) mark the number of sessions in that time unit, starting the day after the previous unit ends (e.g., $W3$: days 15–21).

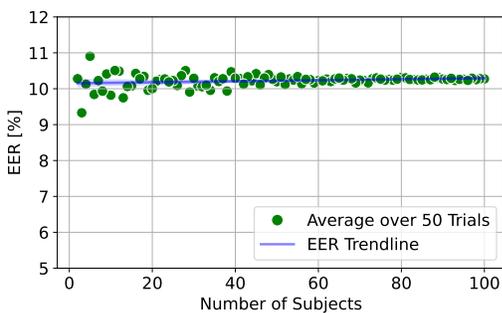


Fig. 2: EER vs. Number of Subjects: Each blue dot represents the average of 50 instances of randomly selecting N subjects and calculating the EER. The dot size represents the standard deviation of the EER values across these 50 calculations.

possible pairs of sessions within specific time intervals and calculated the EER for each. Figure 3 shows the results, demonstrating a positive correlation between time interval and error rate, as expected. Verifying identity becomes harder over time. The EER after one day is 6.7%, increasing to 7.6% after seven days. However, within the second month, the EER rises to 7.2%, which remains lower than the previously reported average of 10.28%, indicating stable performance over a reasonable time interval. Beyond the second month, the EER gradually increases, reaching 14.3% after one year. Performance remains relatively stable in the second and third years. In the fourth and fifth years, performance deteriorates significantly, with EER values of 24.6% and 47.5%, respectively. However, due to the limited data for these time points, the results are not fully reliable. Nonetheless, even the doubling of EER after just one year suggests substantial degradation, indicating the need to update the enrollment set at shorter time intervals.

3) *Visualize Embedding Space*: To better understand how well users are represented in distinct clusters, we project the embeddings into a 2D space using t-Distributed Stochastic Neighbor Embedding (t-SNE) and visualize the results. Figure 4 shows that the metric learning approach successfully separates subjects: the average distance between a subject's own samples is significantly lower (0.75) than the average dis-

tance to other subjects' samples (0.97), indicating the model's ability to capture individual differences. However, we observe two separate colonies within the embedding space, leading us to hypothesize that these colonies could be due to differences in the hardware used during data capture. This hypothesis is confirmed by Figure 5, which reveals that the data captured using BioSemi hardware is located separately from the rest, further illustrating the significant influence of hardware on the embedding distribution. The effect of hardware on the results will be discussed further in the next section.

To investigate the embedding space while accounting for color and spatial limitations, we randomly selected 20 subjects to prevent overcrowding the embedding space with excessive data.⁸ Figure 6 illustrates the embedding space for these selected subjects, demonstrating that the feature extraction method effectively identifies subject identities. However, a session effect is evident, where samples from the same session tend to form distinct clusters. To better visualize this phenomenon, the upper-right region of Figure 6 is magnified in Figure 7, with separate colors indicating different sessions, further confirming the session-based clustering. The session effect on clustering may arise from environmental noise, changes in hair length, slight variations in the positioning of brainwave channels on the head, hardware differences, and natural changes in brainwave patterns over time due to factors such as aging, personal experiences, or mental states (e.g., stress or focus levels) throughout sessions. As a consequence, most samples within a session tend to be either close to or far from the enrollment set at the same time, which can contribute to high variability in performance, as reflected in the relatively large standard deviation ($\pm 6.3\%$) observed in the reported EER (10.28%).

4) *Impact of Different Headsets on Performance*:⁹ Differences in EEG recording hardware and channel positions require further investigation to understand their impact on brainwave authentication. We set up to examine: (1) how well a model trained on one device performs on data from another, (2) whether models trained on multiple devices outperform

⁸BioSemi subjects were excluded to avoid large gaps in the plot and enhance detail visibility.

⁹To clarify, the results in IV.A (Table II) and Sections IV.B.1–IV.B.3 are computed on the combined data from all the three headsets.

device-specific models, and (3) how performs varies when enrollment and verification data come from different devices.

Cross Model: To address the first question, we trained models using data from each specific headset and grouped the test data based on hardware to evaluate performance across models. The results, summarized in Table III, show that headset-specific models perform well only when tested on data from the same device. For example, the model trained on HydroCe achieves an EER of 9.2% on HydroCe test data but 16.1% on Geodesic test data. To further investigate why some hardware-specific models perform better on their own data, we summarized the training and testing data in Table V, which reveals that the number of unique subjects in the training set has a stronger correlation with the EER. Therefore, weaker performance in some models could be related to the lower number of subjects available for learning features.

General Model: We trained a general model using data from all devices. The results in Table III show that the general model outperforms two hardware-specific models, even on their own data, but performs slightly worse than HydroCe. This comparison highlights two possible factors: the general model may benefit from learning cross-device patterns, or training on separate hardware-specific datasets may improve generalization. To examine this, we created a smaller training dataset by keeping only the most frequently used device per subject, ensuring each subject's data comes from a single device. The results in Table IV show a similar pattern, although overall error rates are higher. Table V also shows that HydroCe has significantly more training samples than other devices. These findings suggest that combining smaller datasets with consistent electrode positions can be beneficial for brainwave authentication, especially when large datasets are not available.

Cross Device: In real-world scenarios, individuals using brainwave-based authentication systems may rely on different recording devices. Therefore, authentication systems should maintain performance when enrollment and verification data come from different hardware. To evaluate this, we used the general model trained on the combined training set from all three headsets (230 subjects) and then filtered genuine scores in the evaluation set (100 subjects) by enrollment–verification device pairs to compute the EER for each cross-device setting. The results show an EER of 12.53% for Geodesic–HydroCe pairs and 12.56% for HydroCe–Geodesic pairs. In contrast, BioSemi–HydroCe and HydroCe–BioSemi pairs yield much higher EERs of 39.00% and 44.05%, respectively. For other combinations (e.g., Geodesic–BioSemi), no data were available in the evaluation set. The reason for this variation is that in the training set used for feature extraction, only 7 subjects had HydroCe–BioSemi session pairs and only 2 had BioSemi–HydroCe pairs out of 230 subjects, while 122 subjects had Geodesic–HydroCe session pairs. This limited overlap likely contributes to the poor generalization in the BioSemi–HydroCe condition and explains the formation of separate clusters in the t-SNE visualization shown in Figure 5, where BioSemi forms a distinct cluster.

TABLE III: EER (%) and standard deviation (\pm Std) for different datasets/models. The first column (D) lists the datasets, while the first row (M) lists the models. For example, the entry in the second row and third column (28.32 ± 11.7) represents the EER for the BioSemi model evaluated on the HydroCe dataset. "All" indicates results obtained using all data or models trained on data from all headsets.

D/M	Geodisi	HydroCe	BioSemi	All
Geodisi	9.52 ± 6.1	16.06 ± 8.4	29.33 ± 10.9	8.32 ± 5.2
HydroCe	18.56 ± 10.0	9.24 ± 7.4	28.32 ± 11.7	9.90 ± 8.4
BioSemi	32.63 ± 9.8	33.35 ± 10.5	23.52 ± 8.4	20.62 ± 7.6
All	22.60 ± 12.6	22.08 ± 14.9	28.70 ± 11.6	10.28 ± 6.3

TABLE IV: Table showing EER (%) with standard deviation (\pm Std). The first column (D) lists the datasets, and the first row (M) lists the models, where the training data for each subject is limited to their most frequently used hardware.

D/M	Geodisi	HydroCe	BioSemi	All
Geodisi	17.37 ± 7.8	16.63 ± 8.6	30.10 ± 11.5	10.55 ± 6.3
HydroCe	24.31 ± 9.6	10.91 ± 8.7	29.65 ± 12.2	10.09 ± 8.5
BioSemi	33.52 ± 10.0	36.31 ± 11.2	22.71 ± 8.0	21.55 ± 7.9
All	24.46 ± 11.5	23.27 ± 14.9	28.22 ± 13.4	21.09 ± 14.4

TABLE V: Summary of training datasets by hardware type, listing sessions and subjects with percentages in parentheses.

Hardware	Sessions	Subjects
Train Set: Full Dataset		
Hardware HydroCe	2038	169
Hardware Geodisi	1016	127
Hardware BioSemi	945	66
Train Set: Filtered Dataset		
Hardware HydroCe	1751	131
Hardware Geodisi	435	38
Hardware BioSemi	905	61

C. Enhancing Performance

Samples from the same session are relatively close in the embedding space, but in some cases, they are more dispersed. For example, the pink session in Figure 6. Therefore utilizing multiple verification samples can improve results, albeit at the cost of additional verification time. Table VI illustrates the tradeoff between the number of verification samples, EER, and verification time. The results show that using more samples reduces the error, but improvements beyond four samples show a slower pace in reducing error. Considering the time cost of approximately one second per sample, using four samples appears to be a reasonable balance for the averaging strategy.

Moreover, while comparing the enrollment session to all verification sessions offers useful insights into brainwave authentication, real-world scenarios typically involve a single initial enrollment session followed by one verification attempt after some time. Upon successful login, it may be possible to update the enrollment set with new data from the verification attempt, which can then be used in future verifications. Therefore, in Table VII, we consider the first session as enrollment and the second as verification and report the results. We observe an EER of 6.87% based on single verification sample

and 3.04% for four verification samples, showing significant improvement compared to the previous 10.28% and 6.22%. Additionally, we investigate the effect of multiple enrollment sessions. The results suggest that having two enrollment sessions improves EER to 4.87% and 1.54%, but further enrollment sessions show diminishing returns. Specifically, extending enrollment from two to three sessions reduces EER by only 0.08 pp with 1V (4.87→4.79) and slightly worsens with 4V (1.54→1.57). While four sessions can lower EER (1V: 3.78; 4V: 1.14), these gains are not sustained at five and six sessions (1V: 4.42/4.99; 4V: 1.43/1.39). Thus, beyond two sessions the marginal benefit is small or unstable relative to the added user effort. The average time interval between the two enrollment sessions was 28 days. Among the subjects, 88 had a session interval of 10 days or less, and 32 had intervals of only one or two days.

TABLE VI: Tradeoff between the number of samples (SN) and EER. Increasing the number of samples reduces the EER but increases the verification time, as shown in the table.

SN	1	2	4	8	16	32
EER	10.28	7.75	6.22	5.35	4.87	4.59
Time	1s	2s	4s	8s	16s	32s

TABLE VII: Tradeoff between the number of enrollment sessions (ES) and EER when using 1 verification sample (1V) and 4 verification samples (4V). The table also includes the number of available subjects (NAS) used in the evaluation.

ES	1	2	3	4	5	6
1V	6.87	4.87	4.79	3.78	4.42	4.99
4V	3.04	1.54	1.57	1.14	1.43	1.39
NAS	100	99	99	98	91	75

TABLE VIII: EER (%) when using the same channels as consumer-grade devices. E_x denotes the number of enrollment sessions, and xV denotes the number of verification samples. For the first two columns (Baseline and 4V), all possible verification sessions were considered. However, for the remaining columns, only the first verification session immediately following the last enrollment session was used.

ES	Baseline	4V	E1 1V	E1 4V	E2 1V	E2 4V
Emotiv	13.37	8.68	9.78	4.85	7.57	3.02
DSI-VR300	16.61	11.55	13.09	7.11	11.31	5.35
Muse 2	18.18	13.07	15.01	9.10	12.74	6.39

D. Channels Reduction

In this section, we analyze how reducing the number and placement of EEG channels affects biometric authentication performance. We consider two practically relevant scenarios: limiting channels to those available on specific consumer-grade headsets and restricting channels to subsets from distinct scalp regions under a fixed channel budget.

³The paper did not specify the exact number of days, so we used the same terminology. However, it typically refers to a very short time interval.

TABLE IX: Verification performance (EER [%]) for different scalp regions (frontal, central, temporal, parietal–occipital)

Region	Baseline	4V	E1 1V	E1 4V	E2 1V	E2 4V
Frontal	15.74	10.36	12.51	6.49	10.63	5.13
Central	16.36	11.31	12.41	6.59	10.67	4.94
Temporal	15.76	10.90	11.72	6.48	10.11	4.53
Par.–Occ.	15.11	9.95	11.67	5.82	9.52	4.33

1) *Consumer-grade headsets*: Despite advancements in brainwave sensors in recent years, consumer-grade devices still have a limited number of sensors. Therefore, we selected three devices to evaluate performance by focusing exclusively on the electrodes specific to each device. The selected devices include the Emotiv Epoc X, which features 14¹⁰ electrodes and costs 999 dollars, and the Muse 2, equipped with 4¹¹ electrodes and available at 269.99 euros. Both are budget-friendly options for consumer use. Additionally, we considered the DSI-VR300¹², which includes 7 electrodes and costs 10,000 dollars. This device, though expensive, has the potential to become more affordable in the future and could be integrated into XR devices. We simulated consumer-grade devices by selecting, from the PEERS dataset, the electrodes corresponding to each device’s channel layout. We subsequently trained a separate model for each configuration, following the same subject-disjoint split (230 subjects for training, 100 for evaluation). Table VIII represents the results with a lower number of channels. The baseline EER increases from 10.28% to 13.37%, 16.61%, and 18.18% for Emotiv Epoc X, DSI-VR300, and Muse 2, respectively. The reduction from 93 channels in our main dataset to 14, 7, and 4 channels suggests redundancy due to correlated EEG signals across electrodes. Moreover, using averaging techniques and multiple enrollment sessions can further improve the results.

2) *Regional channel subsets*: To analyze which scalp regions contribute most to authentication performance under a fixed channel budget, we grouped electrodes into frontal, central, temporal, and parietal–occipital regions according to their 10–20 labels. We first counted, in each region, how many such symmetric pairs were available and found that the smallest region contained six pairs; we therefore fixed the subset size to six pairs (12 channels) in all regions so that every regional model operated with the same number and type of electrodes. Separate models were trained on each regional subset using the same architecture and protocol as in the full-channel setting, and the resulting EERs are reported in Table IX. The parietal–occipital subset achieved the lowest error rates, followed by frontal and central subsets, while temporal channels yielded slightly higher EERs. This indicates that placing electrodes over parietal–occipital areas preserves authentication performance best, in line with previous EEG biometric studies[12], [53].

Overall, these channel-reduction experiments indicate that a subset of electrodes can retain much of the authentication performance while improving usability. Our consumer-grade

¹⁰AF3, F7, F3, FC5, T7, P7, O1, O2, P8, T8, FC6, F4, F8, AF4

¹¹TP9, AF7, AF8, TP10

¹²FCz, Pz, P3, P4, PO7, PO8, Oz

TABLE X: Summary of state-of-the-art brain biometric verification solutions evaluated on multisession datasets

Publication	Subjects	Sessions	Interval	Verification	Dataset	Open-source	Inter-subject	EER (%)
Das et al. [10], 2016	50	150 (3)	up to 49 days	30s	×	×	×	10 < EER < 22
Wu et al. [11], 2018	40	80 (2)	30 days	6s	×	×	×	5.75 < EER < 7.07
Maiorana et al. [16], 2018	45	270 (6)	up to 36 months	5s	×	×	×	6.6 ≤ EER ≤ 10.7
Maiorana [12], 2021	45	225 (5)	up to 15 months	5s	×	×	✓	4.8 ≤ EER ≤ 10.7
Debie et al. [13], 2021	54	108 (2)	different days ³	7.3s	✓	×	×	1.16 < EER < 1.93
Chaurasia et al. [14], 2024	54	108 (2)	different days ³	1s	✓	✓	✓	11.99
Proposed Method	345	6007 (~ 17)	up to 66 months	1s	✓	✓	✓	10.28

headset configurations show that a substantial fraction of the baseline performance can be preserved even with only 4–14 channels, suggesting that low-channel-count devices can be effective when properly configured. Regarding electrode placement, parietal–occipital sites demonstrate the best performance, which may be partly explained by the fact that they are less affected by eye-movement artifacts that mainly disturb electrodes on the forehead close to the eyes [54]. Therefore, these sites are particularly attractive for future consumer-grade devices. Moreover, as discussed in Sections IV-B3 and IV-B4, hardware characteristics affect performance, but this can be mitigated by providing sufficient data from target devices and incorporating such data into model training, which will be crucial for achieving reliable results. Finally, if developers adopt a common set of electrode locations, this would support data pooling across devices and studies, benefiting both authentication and other EEG-based applications.

E. Comparative Analysis

We first evaluate the proposed method against state-of-the-art brainwave authentication approaches. Next, we analyze its performance relative to international biometric standards to provide context and clarify the significance of the results.

1) Comparison with Brainwave Authentications Systems:

Table X summarizes the state-of-the-art in multi-session brainwave authentication studies.

Das et al. [10] is one of the earliest multi-session studies, where the authors used simple averaging of samples and compared them using cosine distance. Their reported EER ranges between approximately 10% and 22%. The best result for a one-week interval is around 10% EER, while for an average interval of 34 days, the EER is approximately 15%.

Wu et al.[11] and Debie et al.[13] do not report EER directly. Wu et al. report an FRR of 7.07% and FAR of 5.75% with a 6-second window, while Debie et al. report an FRR of 1.93% and FAR of 1.16% using 7.3 seconds of data with only a few days between sessions. Debie et al.[13] performance is calculated by averaging two scenarios: (1) first session as enrollment and second as verification, and (2) the reverse—an unrealistic setup that assumes access to future data. Both studies use two-session datasets and train a separate CNN per subject, making it unclear how session variability is learned. This design limits generalization and requires training a new model for each user, which is impractical and does not support inter-subject learning.

Maiorana et al. [16], [12] present a longitudinal evaluation with 45 subjects over six sessions collected across three years.

In their first study, they used handcrafted features with a Hidden Markov Model and reported EERs between 6.6% and 10.7% for 5-second verification samples. In the second study, the evaluation was limited to 15 months without a clear explanation. They adopted a metric learning approach, training a feature extractor using contrastive loss on 30 subjects, with a separate CNN per channel, and used a one-class SVM for classification. They reported EERs between 4.8% and 10.7% based on 15 evaluation subjects. However, this approach has key limitations: learning from single channels ignores inter-channel relationships, and one-class SVMs have shown limited reliability, as supported by our findings.

Chaurasia et al.[14] present a framework for EEG evaluation that investigates the multi-session scenario. They show that, on a small dataset, handcrafted features outperform deep learning approaches, reporting an EER of 11.99% using PSD and AR features combined with an SVM, compared to 20.68% EER using a CNN with triplet loss.

In terms of results, when using the first session for enrollment and the second for verification—with an average interval of 28 days—we achieve an EER of 3.04%, outperforming [10], [11], [14]. Maiorana et al.[16], [12] report 4.8% EER with a 5-second window and similar time interval (a month), while we obtain 3.04% with a 4-second window. Compared to Debie et al. [13], who report EER between 1.93% and 1.16% using a 7.3-second window and only a few days between sessions, our 2.30% EER with a 7-second window is measured across a average of 28-day interval. Overall, our results either outperform or are competitive with the state of the art.

However, several important issues are overlooked in related works: (1) Most of the studies use short intervals [10], [11], [13], [11] between sessions while our work considers longer intervals, making the evaluation more realistic for practical applications. (2) Except for Chaurasia et al.[14], none of the studies provide source code, making it impossible to reproduce results or build upon them. This lack of transparency makes it particularly difficult to verify findings, especially where reported performance differs from our results and expectations, such as the use of one-class SVM [12] and subject-specific CNNs [13], [11] (3) None of the reviewed studies examine the effect of recording hardware on authentication performance, which is a key factor for real-world deployment. (4) Some related works [10], [11], [13], [16] do not follow an inter-subject approach, where a general feature extractor is trained to avoid the need for handcrafted features or retraining a new network for each new subject—similar to best practices in face recognition systems [55], [40], [42]. (5) All studies use only one session for verification, even when multiple are available.

They typically repeat experiments with different session pairs (one enrollment - one verification), averaging the results. In contrast, using the first session for enrollment and the rest for verification reflects a more realistic deployment setting.

2) *Comparison with Biometric Industry Standards:* Brainwave authentication is in its early stages, making it crucial to assess progress and identify gaps for real-world implementation. For a biometric system to be deployed in practical applications, it must adhere to international industrial biometric standards. These standards emphasize a very low FAR to ensure security, while maintaining a reasonable FRR to keep the system usable. According to guidelines from NIST (2023)/ISO¹³ and the European Border Guard Agency Frontex¹⁴, biometric systems should operate at FAR $\leq 0.1\%$ (1 in 1,000). Stricter recommendations from FIDO¹⁵ and more recent NIST (August 2024)/ISO 2024 standards¹⁶ suggest a FAR $\leq 0.01\%$ (1 in 10,000). For usability, FRR standards propose a maximum FRR of $\leq 5\%$, allowing at least 19 successful logins out of 20 attempts for genuine users. Table XI presents the FRR at low FAR of 1%, 0.1%, and 0.01%. The results show higher FRR. For instance, using the first session as enrollment and all others as verification, with a duration of 4 seconds, the FRR is 50.39% at FAR 0.1% and 64.75% at FAR 0.01%. This means approximately 9 and 7 successful logins out of 20 attempts for legitimate users, respectively.

While our experiments show that brainwave authentication does not yet meet international standards, our previous experiment on the impact of headsets on performance (Sec.IV-B4) revealed that the headset with a higher number of subjects yielded a lower error rate, and we also know that biometrics that meet these standards typically employ datasets with thousands of subjects[56]. Therefore, we conducted an experiment to explore the relationship between the number of subjects used for training the feature extractor and the error rate. The results in Figure 8 indicate a linear relationship between the logarithm of the number of subjects and the EER. It is obvious that the linear decline eventually slows down. The current trend suggests that with a dataset of a few thousand subjects (from 2,000 to 16,000, depending on the device), it may be possible to approach biometric performance standards. This is reasonable when compared to the data volumes typically used in modalities like face recognition.

V. LIMITATIONS

Our study is subject to certain limitations at three levels: the authentication task, the dataset size, and the recoding hardware.

At the authentication task level, our stimuli were limited to a general text-based memorization task, which imposes non-negligible cognitive load, requires explicit focused attention during authentication, and may limit usability in practical settings; therefore, the usability of this specific task requires

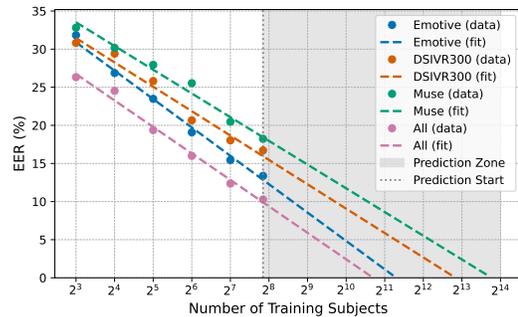


Fig. 8: EER vs. number of training subjects. Points show experimental results; dashed line shows linear prediction.

further investigation, despite existing general usability studies on brainwave authentication [29]. Ideally, multiple task types could be included to allow for performance comparisons and identification of the best-performing task. Additionally, we did not explicitly address replay attacks or EEG-based liveness detection. These aspects can be mitigated by adapting the approach of Lin et al. [3], which involves varying the stimulus and detecting corresponding changes in EEG responses.

At the dataset size level, we only had access to a single dataset including a notable number of sessions and subjects. However, best practices in state-of-the-art face recognition research recommend training on multiple large datasets and testing on an independent dataset to ensure a more robust evaluation [57]. To simulate this, we split our dataset into a subject-independent set for training feature extraction and evaluation. Moreover, we partitioned the dataset based on device, showing that cross-device evaluation significantly increases error (Tables III and IV.) Achieving robust performance requires large amounts of data to extract distinctive features and account for subject-specific, environmental, and hardware-related variations. For example, in the case of face recognition, Kim et al. [57] used 3 datasets comprising over 350,000 subjects and more than 14.9 million images to train a generalized model. Such large-scale datasets are not yet available in brainwave authentication.

At the hardware level, we simulated consumer-grade devices by selecting the same number and channel locations as those found in such devices. However, we did not have a measure to directly compare the quality of sensors. Although sensor quality will improve, the main challenge will remain their number and placement for effective use. It is worth noting that all other multi-session studies in the literature [16], [13], [11], [14] also employ medical-grade data.

VI. CONCLUSION AND FUTURE WORK

This study investigated factors influencing the performance of brainwave authentication over multiple sessions. Our findings highlight that the pairwise deep learning approach for feature extraction outperforms hand-crafted features. Interestingly, we observed that larger evaluation sets do not necessarily make authentication more difficult; instead, they reduce the uncertainty of the results. We acknowledge that the error rate

¹³<https://pages.nist.gov/800-63-3/sp800-63b.html>

¹⁴https://www.frontex.europa.eu/assets/Publications/Research/Best_Practice_Technical_Guidelines_ABC.pdf

¹⁵<https://fidoalliance.org/specs/biometric/requirements/Biometrics-Requirements-v4.0.1-fd-20240522.pdf>

¹⁶<https://pages.nist.gov/800-63-4/sp800-63b.html>

TABLE XI: Performance metrics for different scenarios with detailed enrollment and verification session information. (Verification samples refer to the number of samples from the same session used in a single verification attempt.)

Scenario			Metrics (%)			
Enrollment Sess.	Verification Sess.	Verification Samples	EER	FMR 1%	FMR 0.1%	FMR 0.01%
First	All Others	1	10.28 ± 6.30	43.33 ± 24.46	66.98 ± 25.30	80.65 ± 21.45
First	All Others	4	6.22 ± 5.64	27.86 ± 25.75	50.39 ± 30.40	64.75 ± 30.03
First	Second	1	6.87 ± 7.61	32.15 ± 32.29	54.04 ± 35.66	65.07 ± 34.52
First	Second	4	3.04 ± 6.23	20.71 ± 33.98	33.44 ± 40.58	41.26 ± 42.16
First Two	Third	1	4.87 ± 5.36	23.98 ± 28.29	43.59 ± 35.04	56.39 ± 35.15
First Two	Third	4	1.54 ± 2.74	12.24 ± 26.99	24.28 ± 35.33	32.11 ± 38.79
First Two	Third	16	0.68 ± 1.75	9.60 ± 26.67	16.08 ± 33.37	19.99 ± 36.31

will eventually increase as the number of individuals grows, since only a limited number of individual templates can be reliably distinguished within a finite embedding space. Expanding the embedding space could mitigate this limitation, though evaluating this would require datasets with substantially more users.

Authentication performance remained relatively stable during the first week, with error rates increasing only slightly in the second month. However, error rates could increase by up to 2.1 times after one year, suggesting that the enrollment dataset should be periodically updated to maintain performance. We also examined the impact of recording devices, emphasizing the importance of including data from different headsets in training the feature extractor. Additionally, we showed that increasing the number of enrollment sessions and verification time can help reduce the EER, though only up to a certain point. Our findings also demonstrate that brainwave authentication remains feasible even with a reduced number of EEG sensors. Finally, our results are competitive with the state-of-the-art in brainwave authentication, investigated with stronger assumptions in simpler scenarios.

For future work, there is an urgent need for a public multi-session dataset with a few thousand subjects to ensure brainwave authentication research meets international standards. This dataset should be collected using consumer-grade devices in multiple sessions, and incorporating multiple devices in the same session would be valuable for studying the cross-headset authentication challenge, and across different task paradigms (e.g., passive viewing, motor tasks) to study task similarity. On the technical side, further research should focus on developing deep learning architectures that better capture the unique features of brainwave data, as well as effective signal preprocessing approaches to remove noise. Specifically, the new trend of utilizing attention mechanisms and specialized transformer architectures [58], which have shown promising results in other biometric modalities (e.g., enhanced keystroke dynamics [59]), needs future investigation to see how they can be effectively integrated into EEG-based systems for improved feature extraction and long-term performance. Finally, adopting an open-source approach would enable researchers to build on each other's work, promoting reproducibility and collaboration. Further investigation of more advanced attacks on brainwave authentication is necessary.

ACKNOWLEDGMENT

This work was funded by the Topic Engineering Secure Systems of the Helmholtz Association (HGF) and supported by KASTEL Security Research Labs, Karlsruhe, and Germany's Excellence Strategy (EXC 2050/1 'CeTI'; ID 390696704).

REFERENCES

- [1] R. K. Nath, J. Tervonen, J. Närväinen, K. Pettersson, and J. Mäntyjärvi, "Acute stress data-based fast biometric system using contrastive learning and ultra-short eeg signal segments," in *Adjunct Proceedings of the 2023 ACM International Joint Conference on Pervasive and Ubiquitous Computing & the 2023 ACM International Symposium on Wearable Computing*, 2023, pp. 642–647.
- [2] S. A. Raurale, J. McAllister, and J. M. Del Rincón, "Emg biometric systems based on different wrist-hand movements," *IEEE Access*, vol. 9, pp. 12 256–12 266, 2021.
- [3] F. Lin, K. W. Cho, C. Song, W. Xu, and Z. Jin, "Brain password: A secure and truly cancelable brain biometrics for smart headwear," in *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*, 2018, pp. 296–309.
- [4] M. Fallahi, P. Arias-Cabarcos, and T. Strufe, "Beyond gaze points: augmenting eye movement with brainwave data for multimodal user authentication in extended reality," *Complex & Intelligent Systems*, vol. 12, no. 1, p. 39, 2026.
- [5] P. Arias-Cabarcos, M. Fallahi, T. Habrich, K. Schulze, C. Becker, and T. Strufe, "Performance and usability evaluation of brainwave authentication techniques with consumer devices," *ACM Transactions on Privacy and Security*, vol. 26, no. 3, pp. 1–36, 2023.
- [6] I. Nakanishi and T. Yoshikawa, "Brain waves as unconscious biometrics towards continuous authentication—the effects of introducing pca into feature extraction," in *2015 International Symposium on Intelligent Signal Processing and Communication Systems (ISPCS)*. IEEE, 2015, pp. 422–425.
- [7] T. Schons, G. J. Moreira, P. H. Silva, V. N. Coelho, and E. J. Luz, in *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications: 22nd Iberoamerican Congress, CIARP 2017, Valparaiso, Chile, November 7–10, 2017, Proceedings 22*. Springer, 2018, pp. 601–608.
- [8] A. J. Bidgoly, H. J. Bidgoly, and Z. Arezoumand, "Towards a universal and privacy preserving eeg-based authentication system," *Scientific Reports*, vol. 12, no. 1, p. 2531, 2022.
- [9] M. Fallahi, T. Strufe, and P. Arias-Cabarcos, "Brainnet: Improving brainwave-based biometric recognition with siamese networks," in *2023 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 2023, pp. 53–60.
- [10] R. Das, E. Maiorana, and P. Campisi, "Eeg biometrics using visual stimuli: A longitudinal study," *IEEE Signal Processing Letters*, vol. 23, no. 3, pp. 341–345, 2016.
- [11] Q. Wu, Y. Zeng, C. Zhang, L. Tong, and B. Yan, "An eeg-based person authentication system with open-set capability combining eye blinking signals," *Sensors*, vol. 18, no. 2, p. 335, 2018.
- [12] E. a. Maiorana, "Learning deep features for task-independent eeg-based biometric verification," *Pattern Recognition Letters*, vol. 143, pp. 122–129, 2021.
- [13] E. Debie, N. Moustafa, and A. Vasilakos, "Session invariant eeg signatures using elicitation protocol fusion and convolutional neural network," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 4, pp. 2488–2500, 2021.

- [14] A. K. Chaurasia, M. Fallahi, T. Strufe, P. Terhörst, and P. A. Cabarcos, "Neuroidbench: An open-source benchmark framework for the standardization of methodology in brainwave-based authentication research," *Journal of Information Security and Applications*, vol. 85, p. 103832, 2024.
- [15] G. Huang, Z. Hu, W. Chen, S. Zhang, Z. Liang, L. Li, L. Zhang, and Z. Zhang, "M3cv: A multi-subject, multi-session, and multi-task database for eeg-based biometrics challenge," *NeuroImage*, vol. 264, p. 119666, 2022.
- [16] E. Maiorana and P. Campisi, "Longitudinal evaluation of eeg-based biometric recognition," *IEEE transactions on Information Forensics and Security*, vol. 13, no. 5, pp. 1123–1138, 2017.
- [17] M. J. Kahana, L. J. Lohnas, M. K. Healey, A. Aka, A. W. Broitman, P. Crutchley, E. Crutchley, K. H. Alm, B. S. Katerman, N. E. Miller *et al.*, "The penn electrophysiology of encoding and retrieval study," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 2024.
- [18] E. Niedermeyer and F. L. da Silva, *Electroencephalography: basic principles, clinical applications, and related fields*. Lippincott Williams & Wilkins, 2005.
- [19] L. Hu and Z. Zhang, *EEG signal processing and feature extraction*. Springer, 2019.
- [20] V. Gupta, N. K. Saxena, A. Kanungo, A. Gupta, P. Kumar, and Salim, "A review of different eeg classification/detection techniques for improved medical applications," *Int. Journal of System Assurance Engineering and Management*, vol. 13, no. 3, pp. 1037–1051, 2022.
- [21] H. Adeli and S. Ghosh-Dastidar, *Automated EEG-based diagnosis of neurological disorders: Inventing the future of neurology*. CRC press, 2010.
- [22] L. F. Haas, "Hans berger (1873–1941), richard caton (1842–1926), and electroencephalography," *Journal of Neurology, Neurosurgery & Psychiatry*, vol. 74, no. 1, pp. 9–9, 2003.
- [23] A. Kawala-Sterniuk, N. Browarska, A. Al-Bakri, M. Pelc, J. Zygarlicki, M. Sidikova, R. Martinek, and E. J. Gorzelanczyk, "Summary of over fifty years with brain-computer interfaces—a review," *Brain Sciences*, vol. 11, no. 1, p. 43, 2021.
- [24] Q. Gui, M. V. Ruiz-Blondet, S. Laszlo, and Z. Jin, "A survey on brain biometrics," *ACM Computing Surveys (CSUR)*, vol. 51, no. 6, pp. 1–38, 2019.
- [25] Z. Rui and Z. Yan, "A survey on biometric authentication: Toward secure and privacy-preserving identification," *IEEE access*, vol. 7, pp. 5994–6009, 2018.
- [26] S. Marcel and J. d. R. Millán, "Person authentication using brainwaves (eeg) and maximum a posteriori model adaptation," *IEEE Trans. on pattern analysis and machine intelligence*, vol. 29, no. 4, pp. 743–752, 2007.
- [27] B. Kaliraman, P. Singh, and M. Duhan, "Use of eeg as a unique human biometric trait for authentication of an individual," in *International Conference on Advanced Communication and Computational Technology*. Springer, 2019, pp. 277–286.
- [28] S. Yang and F. Deravi, "On the usability of electroencephalographic signals for biometric recognition: A survey," *IEEE Transactions on Human-Machine Systems*, vol. 47, no. 6, pp. 958–969, 2017.
- [29] M. Fallahi, P. Arias-Cabarcos, and T. Strufe, "On the usability of next-generation authentication: A study on eye movement and brainwave-based mechanisms," in *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 2025.
- [30] M. Wang, S. Wang, and J. Hu, "Cancellable template design for privacy-preserving eeg biometric authentication systems," *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 3350–3364, 2022.
- [31] D. Silverman, "The rationale and history of the 10-20 system of the international federation," *American Journal of EEG Technology*, vol. 3, no. 1, pp. 17–22, 1963.
- [32] W. Peng, "Eeg preprocessing and denoising," *EEG Signal Processing and Feature Extraction*, pp. 71–87, 2019.
- [33] A. Gramfort, M. Luessi, E. Larson, D. A. Engemann, D. Strohmeier, C. Brodbeck, L. Parkkonen, and M. S. Hämäläinen, "Mne software for processing meg and eeg data," *neuroimage*, vol. 86, pp. 446–460, 2014.
- [34] R. Suppiah and A. Prasad Vinod, "Biometric identification using single channel eeg during relaxed resting state. iet biometr 7: 342–348," 2018.
- [35] A. E. K. Khalil, J. A. Perez-Diaz, J. A. Cantoral-Ceballos, and J. M. Antelis, "Unlocking security for comprehensive electroencephalogram-based user authentication systems," *Sensors*, vol. 24, no. 24, p. 7919, 2024.
- [36] P. D. Welch, "The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms," *IEEE Transactions on Audio and Electroacoustics*, vol. 15, no. 2, pp. 70–73, 1967.
- [37] H. Nai-Jen and R. Palaniappan, "Classification of mental tasks using fixed and adaptive autoregressive models of eeg signals," in *The 26th annual international conference of the IEEE engineering in medicine and biology society*, vol. 1. IEEE, 2004, pp. 507–510.
- [38] B. Kulis *et al.*, "Metric learning: A survey," *Foundations and Trends® in Machine Learning*, vol. 5, no. 4, pp. 287–364, 2013.
- [39] S. Horiguchi, D. Ikami, and K. Aizawa, "Significance of softmax-based features in comparison to distance metric learning-based features," *IEEE Trans. on pattern analysis and machine intelligence*, vol. 42, no. 5, pp. 1279–1285, 2019.
- [40] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proc.s of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 815–823.
- [41] H. Oh Song, Y. Xiang, S. Jegelka, and S. Savarese, "Deep metric learning via lifted structured feature embedding," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4004–4012.
- [42] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, "Supervised contrastive learning," *Advances in neural information processing systems*, vol. 33, pp. 18 661–18 673, 2020.
- [43] X. Wang, X. Han, W. Huang, D. Dong, and M. R. Scott, "Multi-similarity loss with general pair weighting for deep metric learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5022–5030.
- [44] Y. Zheng, Z. Liu, R. Mo, Z. Chen, W.-s. Zheng, and R. Wang, "Task-oriented self-supervised learning for anomaly detection in electroencephalography," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2022, pp. 193–203.
- [45] R. T. Schirrmester, J. T. Springenberg, L. D. J. Fiederer, M. Glasstetter, K. Eggenesperger, M. Tangermann, F. Hutter, W. Burgard, and T. Ball, "Deep learning with convolutional neural networks for eeg decoding and visualization," *Human brain mapping*, vol. 38, no. 11, pp. 5391–5420, 2017.
- [46] V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, "Eegnet: a compact convolutional neural network for eeg-based brain-computer interfaces," *Journal of neural engineering*, vol. 15, no. 5, p. 056013, 2018.
- [47] X. Zhang and L. Yao, *Deep learning for EEG-based brain-computer interfaces: Representations, algorithms and applications*. World Scientific, 2021.
- [48] Y. Jiang, P. Shen, L. Zeng, X. Zhu, D. Jiang, and C. Chen, "Cancelable biometric schemes for euclidean metric and cosine metric," *Cybersecurity*, vol. 6, no. 1, p. 4, 2023.
- [49] A. J. Mansfield and J. L. Wayman, "Best practices in testing and reporting performance of biometric devices," 2002.
- [50] S. Ayeswarya and K. J. Singh, "A comprehensive review on secure biometric-based continuous authentication and user profiling," *IEEE Access*, 2024.
- [51] F. Navaratnam-Blair, K. Wagstaff, G. Miller, M. Cumberbatch, and C. Rethore, "Beyond reality: Is the long-awaited vr revolution finally on the horizon?" National Research Group, April 2022, design by Grace Stees, Illustrations by Hannah Robinson.
- [52] L. Friedman, H. Stern, V. Prokopenko, S. Djanian, H. Griffith, and O. Komogortsev, "Biometric performance as a function of gallery size," *Applied Sciences*, vol. 12, no. 21, p. 11144, 2022.
- [53] D. La Rocca, P. Campisi, B. Vegso, P. Cserti, G. Kozmann, F. Babiloni, and F. D. V. Fallani, "Human brain distinctiveness based on eeg spectral coherence connectivity," *IEEE transactions on Biomedical Engineering*, vol. 61, no. 9, pp. 2406–2412, 2014.
- [54] V. Ronca, R. Capotorto, G. Di Flumeri, A. Giorgi, A. Vozzi, D. Germano, V. D. Virgilio, G. Borghini, G. Cartocci, D. Rossi *et al.*, "Optimizing eeg signal integrity: A comprehensive guide to ocular artifact correction," *Bioengineering*, vol. 11, no. 10, p. 1018, 2024.
- [55] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proc. of the IEEE/CVF conf. on computer vision and pattern recognition*, 2019, pp. 4690–4699.
- [56] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao, "Ms-celeb-1m: A dataset and benchmark for large-scale face recognition," in *ECCV*, 2016.
- [57] M. Kim, A. K. Jain, and X. Liu, "Adaface: Quality adaptive margin for face recognition," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 18 750–18 759.

- [58] K.-N. Nguyen, S. Rasnayaka, S. Wickramanayake, D. Meedeniya, S. Saha, and T. Sim, "Spatio-temporal dual-attention transformer for time-series behavioral biometrics," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 6, no. 4, pp. 591–601, 2024.
- [59] D. Senarath, S. Tharinda, M. Vishvajith, S. Rasnayaka, S. Wickramanayake, and D. Meedeniya, "Behaveformer: A framework with spatio-temporal dual attention transformers for imu-enhanced keystroke dynamics," in *2023 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2023, pp. 1–9.
- [60] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [61] K. K. Ang, Z. Y. Chin, H. Zhang, and C. Guan, "Filter bank common spatial pattern (fbcsp) in brain-computer interface," in *2008 IEEE international joint conference on neural networks (IEEE world congress on computational intelligence)*. IEEE, 2008, pp. 2390–2397.
- [62] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.

APPENDIX A SESSION COUNT DISTRIBUTION

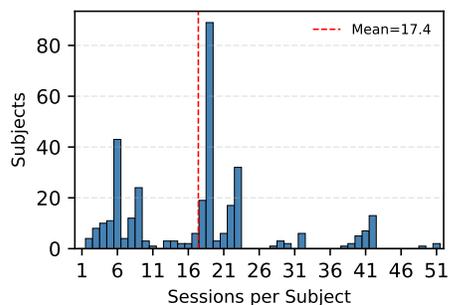


Fig. 9: Histogram of session counts per subject.

Figure 9 shows the session count distribution in our dataset.

APPENDIX B ARCHITECTURAL DETAILS OF DEEP LEARNING FEATURE ENCODERS

This appendix provides detailed specifications for the Convolutional Neural Network (CNN) architectures employed as feature encoders in our benchmarking study. In our experiments, CNN-based models trained with Supervised Contrastive Loss (SupConLoss) [42] generally achieved the best performance among the evaluated feature extraction methods, with the strongest results obtained using ResNet1D.

A. Best-Performing Architecture: ResNet1D

The core concept of the Residual Network was introduced by He et al. [60] to address the degradation problem in very deep neural networks. By incorporating skip connections (residual blocks), ResNet allows information to bypass one or more layers, enabling the stable training of networks with significantly greater depth. In our work, we adopt the ResNet1D backbone proposed by Zheng et al. [44] for EEG anomaly detection, which is based on a ResNet-34-style architecture with temporally oriented convolutional kernels for multi-channel EEG segments.

Architecture Detail: ResNet1D is a deep CNN built from residual blocks that operate along the temporal dimension of the EEG. An initial temporal convolutional block is followed by several stages of stacked residual blocks. Each residual block contains one or more 1D convolutional layers with batch normalization (BN) and rectified linear unit (ReLU) activations, together with a skip connection from input to output. After the last residual stage, global average pooling over time is applied and the resulting feature vector is passed through a fully connected layer, which provides the final feature embedding for comparison.

B. ShallowConvNet

Introduced by Schirrneister et al. [45], the ShallowConvNet was designed to closely mimic the successful Filter Bank Common Spatial Pattern (FBCSP) pipeline [61] in an end-to-end deep learning framework. The goal was to demonstrate that a simple deep-learning structure could perform as well.

Architecture Detail: This architecture consists of two convolutional layers, specifically engineered to extract frequency and spatial information, followed by operations that implement log-variance features:

- 1) Temporal filtering (Block 1): A convolution with kernel size 1×25 is applied along the temporal dimension. This layer learns temporal filters independently across each EEG channel.
- 2) Spatial filtering (Block 2): A convolution with kernel size $Channels \times 1$ is used to linearly combine the temporal features across electrodes, followed by batch normalization (BN).
- 3) Feature extraction: The output of the spatial filter then passes through a squaring operation ($x \mapsto x^2$), followed by mean pooling (size 1×75 , stride 15), and a logarithmic activation ($\log(x)$). These operations directly implement the log-variance feature computation analogous to FBCSP.

C. DeepConvNet

Also introduced by Schirrneister et al. [45], DeepConvNet was designed as a more generic CNN architecture, closer to models used in computer vision such as VGG [62]. The motivation was to explore whether less specialized, deeper networks could capture a wider variety of abstract features from the raw EEG signal, rather than being biased towards the amplitude-based features commonly used in specific task decoding.

Architecture Detail: The network is constructed with four sequential blocks, using smaller kernels and max pooling for feature compression:

- 1) Block 1: Temporal convolution (size 1×10) followed by a spatial convolution (size $Channels \times 1$), an exponential linear unit (ELU) non-linearity, and max pooling (size 1×3 , stride 3).
- 2) Blocks 2, 3, 4: These three blocks are stacked consecutively. Each consists of a convolution (size 1×10), an ELU activation, and max pooling (size 1×3 , stride 3).

These blocks use small convolutional kernels to extract progressively abstract and compressed features.

- 3) Output: The final output is flattened and passed through a dense (fully connected) layer, which provides the feature embedding for comparison.

D. EEGNet

Lawhern et al. [46] introduced EEGNet to create a compact and generalized CNN that could perform well across multiple brain-computer interface (BCI) paradigms (e.g., P300, SMR, ERN) and datasets. The primary design constraint was to achieve this with a minimal number of trainable parameters, making the model suitable for deployment in resource-constrained or real-time BCI systems.

Architecture Detail: EEGNet achieves its efficiency through the factorization of standard convolutions into depthwise and separable convolutions:

- 1) Block 1 (temporal-spatial filter): This block consists of a temporal convolution followed by a depthwise convolution. Temporal filters learn frequency features. The depthwise filters then apply a unique spatial filter to the output of each temporal filter, extracting frequency-specific spatial patterns.
- 2) Block 2 (feature mixing): This block employs a separable convolution, which is the combination of a depthwise convolution (providing a temporal summary for each feature map) and a pointwise convolution (a 1×1 convolution). The pointwise convolution acts as an efficient final layer, learning how to optimally combine or mix the feature maps across channels.



Thorsten Strufe is Professor of IT Security at KIT, adjunct professor for Privacy and Network Security at TU Dresden, Deputy Speaker of the Cluster of Excellence CeTI (Centre for Tactile Internet with Human-in-the-Loop), and PI in the German Competence Center for IT Security KASTEL, the Research Training Group RoSI (TU Dresden), and the 5G-Lab Germany. Prior to his position at KIT, he was full professor of privacy and security at TU Dresden from 2014 to 2019, received an offer for the Chair of Security and Privacy at University of Freiburg,

was assistant professor for Peer-to-Peer Networks at TU Darmstadt from 2009 to 2014, and deputy professor for Secure Distributed Systems at the University of Mannheim throughout 2011. During this time, he was PI in several research groups and collaborative projects, such as the SFB 912 HAEC (Highly Adaptive Energy-Efficient Communication, TU Dresden), the SFB 1053 MAKI (Multi-Mechanism Adaptation for the Future Internet, TU Darmstadt), the research group 733 QuaP2P (Improving the Quality of Peer-to-Peer Systems, TU Darmstadt) and various collaborative projects within the funding programmes of EU and BMBF.



Matin Fallahi is currently pursuing the Ph.D. degree in computer science with the KASTEL – Institute of Information Security and Dependability, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany. The research focuses on biometric authentication, in particular brainwave-based user authentication using electroencephalography (EEG). This includes the security and robustness of EEG-based biometrics, usability and continuous authentication in extended reality (XR) scenarios, and privacy-preserving processing of biometric data. Fallahi has coauthored

several publications on brainwave-based biometrics and related topics in usable security and privacy.



Patricia Arias-Cabarcos is a Professor of IT Security at the University of Paderborn, where she leads the Human-centered IT Security (HITS) group. Her research focuses on usable authentication, behavioral data privacy, data transparency, and related topics in usable security and privacy. Before joining Paderborn, she was a Senior Researcher at the Karlsruhe Institute of Technology (KIT) from 2019 to 2021, a Humboldt Fellow at the University of Mannheim between 2017 and 2019, and an Assistant Professor at Universidad Carlos III de Madrid (2013–2018),

where she also obtained her Ph.D. in Telematic Engineering. She has additionally held positions as a visiting postdoctoral researcher at TU Darmstadt and as a research intern at NEC Laboratories Europe in Heidelberg.