International Journal of Basic and Applied Sciences, 14 (2) (2025) 493-503

International Journal of Basic and Applied Sciences

Website: www.sciencepubco.com/index.php/IJBAS
https://doi.org/10.14419/7gtxb837
Research paper

Privacy-Preserving Deep Learning Models For Alcoholism Diagnosis Through EEG Data Analysis Using Differential Privacy Mechanisms

Varalatchoumy M [1], A. Srinivasa Reddy [2], S. Sagar Imambi [3], Gaurav Kumar Ameta [4], K. Thaiyalnayaki [5], S. K. Logesh [6], S. Dhivya [7], K Muthulakshmi [8] *

[1] *Department of Artificial Intelligence and Machine Learning, Cambridge Institute of Technology, Bengaluru, Karnataka 560036, India*
[2] *Department of Computer Science and Engineering (Data Science), CVR College of Engineering, Hyderabad, Telangana 501510, India*
[3] *Department of Computer Science & Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Andhra Pradesh 522302, India*
[4] *Department of Computer Science and Engineering, Parul Institute of Technology, Parul University, Vadodara, Gujarat 391760, India*
[5] *Department of Electronics and Communication Engineering, Chennai Institute of Technology, Chennai, Tamil Nadu 600069, India*
[6] *Department of Electrical and Electronics Engineering, Kongu Engineering College, Perundurai, Erode, Tamil Nadu 638060, India*
[7] *Department of Research and Innovation, Saveetha School of Engineering, SIMATS, Chennai, Tamil Nadu, India*
[8] *Department of Information Technology, Panimalar Engineering College, Chennai, Tamil Nadu 600123, India*
**Corresponding author E-mail: hod.aiml@cambridge.edu.in*

Received: May 10, 2025, Accepted: June 18, 2025, Published: June 30, 2025

Abstract

Alcoholism diagnosis through electroencephalogram (EEG) data analysis offers a promising alternative to traditional methods by identifying specific brain activity patterns associated with alcohol dependency. While deep learning techniques have demonstrated high accuracy in classifying EEG signals, privacy concerns related to sensitive medical data remain prevalent. Ensuring the privacy of patient data is critical for building trust and enabling the adoption of these tools in real-world clinical settings. This study develops deep learning models with enhanced privacy guarantees by incorporating differential privacy mechanisms, including (ϵ, δ)-Differential Privacy and Gaussian Differential Privacy (GDP). We compare their efficacy in preserving data privacy while maintaining model utility. Experiments show that convolu-tional and long-short-term-memory models optimized with Adam excel in utility and stability. GDP outperforms (ϵ, δ)-DP by requiring less noise, while DP-Adam surpasses DP-SGD in privacy and utility, particularly for fast convergence. Larger datasets further enhance this balance, emphasizing the importance of effective privacy mechanisms and sufficient data. By balancing privacy and utility, this work con-contributes a novel approach to privacy-preserving AI for sensitive health applications, emphasizing scalable models that maintain diagnostic accuracy.

Keywords: *Electroencephalogram (EEG); Gaussian Differential Privacy (GDP); Differential Privacy; Deep learning.*

1. Introduction

Alcoholism, or alcohol use disorder (AUD), is a complex condition whose diagnosis has traditionally relied on clinical assessments, questionnaires, physical exams, psychological evaluations, and laboratory tests, often making it subjective and time-consuming (Akbari et al., 2021). Accurate diagnosis is crucial for effective treatment, guiding therapies like Cognitive Behavioral Therapy (CBT), Motivational Enhancement Therapy (MET), and medications such as naltrexone or acamprosate (Shi et al., 2023). Precise diagnosis also considers psychological, social, and genetic factors, enabling personalized care and reducing relapse risk, while addressing co-occurring mental health issues like anxiety and depression (Sadiq et al., 2021a).

Considering the importance of AUD correct diagnosis, the need to develop an effective tool for this process has become the subject of many studies. More recently, the use of electroencephalogram (EEG) technology has introduced a novel approach to alcoholism diagnosis. EEG measures the electrical activity of the brain and has been found useful in identifying specific patterns that may be associated with alcohol dependency. Researchers have employed various EEG analysis methods to distinguish between the brain activity of alcohol-

ics and non-alcoholics. These methods analyze changes in brain wave patterns and can potentially offer a more objective and precise measure of brain function alterations due to alcohol use (Yeh et al., 2023; Sadiq et al., 2022a).

Machine learning methods have significantly improved the analysis of EEG data for diagnosing AUD by identifying complex patterns that may not be apparent to the human eye. However, traditional approaches often face limitations such as over-fitting, especially with small or non-diverse datasets, and challenges in extracting EEG features (Cohen et al., 2023; Pain et al., 2023). This has driven the integration of deep learning techniques, which can model higher-level abstractions through architectures like convolutional neural networks (CNNs) and recurrent neural networks (RNNs) (Salankar et al., 2022; Anuragi & Sisodia, 2020). While deep learning methods enhance AUD diagnosis from EEG data, they expose sensitive patient data to potential privacy risks due to the extensive data requirements and model complexity, raising concerns about compliance with regulations like HIPAA and GDPR (Sadiq et al., 2022b).

It is thus crucial to explicitly consider privacy aspects in diagnostic systems when working with medical data. Data-based diagnostic tools, therefore, need to ensure that the model is private in the sense that the patient data used to train the model remains private. This helps protect sensitive data and builds trust, making it more likely for the tool to be used in real-world clinical settings.

Therefore, there arises a crucial need to develop a diagnostic method that not only utilizes the predictive power of deep learning but also incorporates robust privacy-preserving mechanisms. Recent advancements in techniques such as differential privacy, which adds noise to obscure individual data points, offer promising avenues. These methods ensure that the model learns the necessary patterns to diagnose AUD effectively without compromising the privacy of the individual data contributions (Browarska et al., 2021). However, studies on the integration of differential privacy in EEG diagnostic methods based on deep learning are very limited. Thus, the pursuit of developing a technique that allows maintaining the privacy of the data and preserves the utility of the model holds significance.

To systematically address the research objectives, we propose a methodology composed of two phases: i) the development of a deep learning model for AUD diagnosis, and ii) the incorporation and evaluation of differential privacy techniques. The primary contributions of this work include the design of a deep learning model that classifies signals with high performance, with an architecture specifically adapted to capture the spatial and temporal characteristics of the signals. Additionally, this study provides a comprehensive comparison between two differential privacy approaches, based on how privacy is accounted for, and offers a detailed analysis of the privacy-utility trade-off in models with varying sample sizes and privacy guarantees. Experiments show that GDP achieves better utility than (ϵ, δ)- DP with less noise, while DP-Adam outperforms DP-SGD in privacy and utility. Larger datasets further enhance this balance, highlighting the value of robust privacy mechanisms and sample data. This work sets the stage for improving data utility while maintaining strong privacy protection in sensitive applications.

2. Literature review

Anuragi et al. (2021) introduce a novel framework for the automated identification of epileptic seizures utilizing time-segmented EEG signals. The writers merge Frequency Band Subspace Energy (FBSE) with Empirical Wavelet Transform (EWT) to extract features and utilize machine learning for classification purposes. The approach consists of breaking down EEG signals into various frequency bands, enabling the acquisition of intricate temporal and spectral data essential for identifying seizures. This study tackles an important clinical issue, which is the precise and prompt identification of epileptic seizures, thereby facilitating improved patient monitoring and treatment strategies.

Farsi et al. (2021a) investigate deep learning methods for classifying EEG signals to distinguish between brain activity of alcoholics and non-alcoholics. The authors suggest a method based on deep learning to classify EEG data collected from people with alcohol dependence. Utilizing sophisticated neural networks, the research shows that deep learning models can effectively detect alterations in brain patterns caused by alcohol, which can be applied for diagnostic purposes. The study emphasizes the promise of machine learning and neural networks in enhancing the diagnosis and tracking of alcohol use disorders.

Li and Wu's (2022) advances the classification of EEG signals to differentiate between normal and alcoholic brain conditions utilizing deep learning. This research supports earlier findings in alcohol detection through EEG data, concentrating on a deep learning method to improve classification precision. Through the training of models on an extensive dataset containing both normal and alcoholic EEG signals, the authors emphasize the reliability of deep learning algorithms in diagnosing via EEG. Their results indicate that this method can be employed in clinical environments to track the impact of alcohol on brain activity over time, potentially aiding in early intervention strategies for alcohol-related issues.

Ahmadi et al. (2017) aim to identify alcoholism through EEG signals by extracting features of functional brain networks. The research extends beyond simple EEG analysis by exploring how alcohol consumption alters functional connectivity among various brain regions. The authors suggest a technique for feature extraction that identifies these alterations in connectivity, which are subsequently categorized using machine learning algorithms. Their research establishes the basis for more advanced methods of diagnosing alcoholism, adding to the expanding area of neuroimaging in addiction studies. This research highlights the significance of comprehending the fundamental brain network alterations in alcohol dependence and how these insights can guide treatment strategies.

Farsi et al. (2021b) coincides with reference #3 and offer a comparable emphasis: classification of alcoholic EEG signals using deep learning techniques. The research is part of a continuous investigation into using sophisticated machine learning techniques to detect brain activity patterns caused by alcohol. Their research helps enhance detection systems for alcohol dependence, which may be incorporated into wearable BCI systems for ongoing monitoring. These systems may offer immediate insights into how alcohol impacts brain function, enabling prompt interventions.

In this study, Balle et al. (2020) examine privacy amplification within the realm of differential privacy, particularly through the application of subsampling methods. The authors employ rigorous analyses through couplings and divergences to improve the security of privacy-preserving techniques, especially in data sharing contexts where sensitive data such as EEG signals could be present. This study is vital for creating privacy-preserving techniques for health data, guaranteeing that personal details in applications like BCI and EEG monitoring stay secure while still being beneficial for analysis. The research offers a robust mathematical foundation for enhancing privacy in different machine learning contexts.

Sherstinsky (2020) offers a comprehensive examination of Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks, which are essential models in deep learning for analyzing sequential data. The document describes the theoretical foundations of RNNs and LSTMs, detailing their architecture and operations, alongside their uses in areas such as speech recognition, natural language processing, and forecasting time-series data. This study is especially important for analyzing EEG signals, as sequential data patterns are essential for comprehending brain activity through time. LSTMs are particularly effective at capturing the temporal dependencies inherent in EEG signals, rendering them suitable for real-time brain monitoring and disorder identification applications.

Brunnhuber et al. (2020) reviewed the development of home-based video-EEG telemetry, emphasizing its transition from clinical to in-home use. The study outlines technological advances such as wireless EEG systems and mobile data transmission, enabling more comfortable and cost-effective long-term monitoring. The authors highlight benefits like increased ecological validity and identify challenges in standardization, data security, and interpretation.

Ko et al. (2020) proposed a method that combines eye-blink recognition with single-channel forehead EEG to enhance fatigue prediction during sustained attention tasks. Their results show that integrating blink metrics improves the accuracy of cognitive fatigue detection, offering a low-cost and practical solution for applications in safety-critical environments like driving or workplace monitoring.

Dwork et al. (2006) introduced the concept of differential privacy, providing a rigorous mathematical framework for protecting individual data. By calibrating noise based on a function's sensitivity, their approach ensures robust privacy guarantees in data analysis. This foundational work underpins many modern privacy-preserving algorithms used in sensitive fields such as healthcare and government statistics.

Porjesz et al. (2005) explored the use of EEG-based neurophysiological markers to study alcoholism, focusing on traits like reduced P300 amplitude and altered brain oscillations. These markers were shown to be heritable and useful in identifying individuals at genetic risk. The study highlighted their potential as endophenotypes for understanding the neurocognitive effects of alcohol dependence and guiding early diagnosis and intervention.

3. Data and processing

3.1. Data

The dataset analyzed in this study originates from the Neurodynamic Laboratory at the State University of New York, hosted by the UCI repository. It encompasses data collected from 122 subjects, divided into two groups: 77 diagnosed with alcoholism and 45 control subjects. Each subject underwent up to 120 trials involving EEG recordings, with the trials of alcoholic and control subjects labeled correspondingly. EEG data were acquired using 64 electrodes positioned on the subjects' scalps. The signals were sampled at a frequency of 256Hz throughout 1 second, corresponding to a 3.9-msec epoch for each of the 256 samples collected per trial. Figure 1 shows a sample of the raw data for alcoholic and control participants for 10 electrodes.

The experimental protocol involved presenting subjects with visual stimuli. Participants were exposed either to a single stimulus (S1) or two stimuli (S1 and S2), sourced from the Snodgrass and Vanderwart picture set. For the two-stimuli condition, presentations were either congruent (S1 identical to S2) or incongruent (S1 differing from S2).

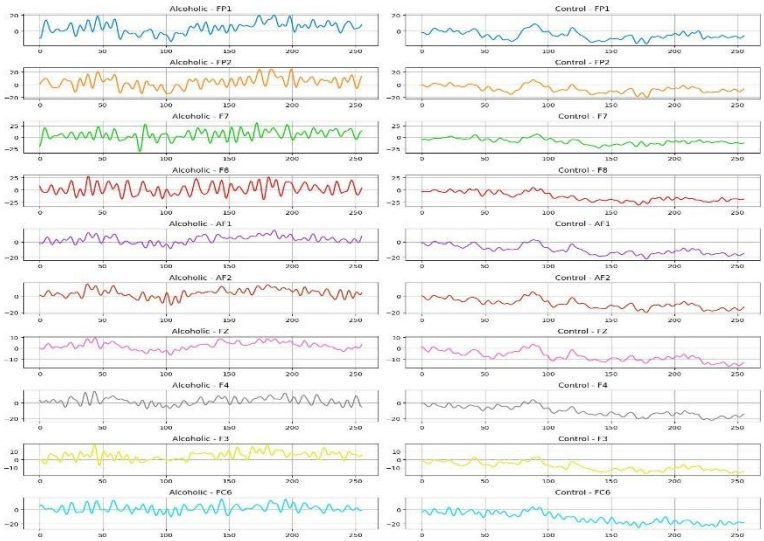

Fig. 1: Sample Of Raw Data for an Alcoholic Participant (Left) and A Control Participant (Right) – First 10 Channels.

The UCI EEG Alcoholism dataset, while widely used, presents several limitations that affect its generalizability and robustness. A major concern is the lack of demographic information, such as age, gender, or ethnicity, which restricts the evaluation of model fairness across diverse populations. Additionally, the dataset includes only 122 subjects, limiting the sample diversity and potentially leading to overfitting. The diagnostic labels provided are also somewhat ambiguous, lacking detailed clinical criteria or information on comorbid conditions, which may introduce label noise. Furthermore, each EEG trial is restricted to a short, 1-second recording, which constrains the temporal scope of brain activity analysis. The experimental protocols and equipment, although standardized, may not align with current clinical practices, further reducing the applicability of findings to real-world settings. These factors collectively limit the dataset's utility for developing broadly generalizable or clinically deployable models.

3.2. Data processing

The data preprocessing pipeline began with the decompression of the raw EEG files. The EEG data was stored in compressed .tar.gz and .gz formats, and each file was extracted using standard decompression techniques. Any files that contained error messages such as "S2 match err" or had extreme values outside the expected range (values exceeding ±200μV) were ignored to ensure data quality. After extracting and verifying the integrity of each trial file, only the valid trials were processed further (10,617 trials).

The next step involved applying a band-pass filter to remove unwanted frequency components. Specifically, a Butterworth filter (Yeh et al., 2023) with a low cut-off frequency of 0.5 Hz and a high cutoff frequency of 50 Hz was used to retain relevant brainwave frequencies. This filter was applied individually to each of the 64 channels of the EEG signals.

After filtering, a Fast Fourier Transform (FFT) was applied to the time-domain data to convert it into the frequency domain, extracting useful frequency-based features. Once the transformation was completed, the data was normalized across all channels using a standard scaling approach to ensure that each channel had a mean of zero and a unit variance, reducing any bias due to varying signal amplitudes.
To further enhance the dataset and improve model generalization, data augmentation techniques were applied. This included adding random Gaussian noise and introducing slight amplitude changes to simulate variability in the signal, increasing the robustness of the model during training. The final preprocessed data was then used as input for the subsequent modeling phase. Figure 2 shows a sample of the processed data for alcoholic and control participants for 10 electrodes.

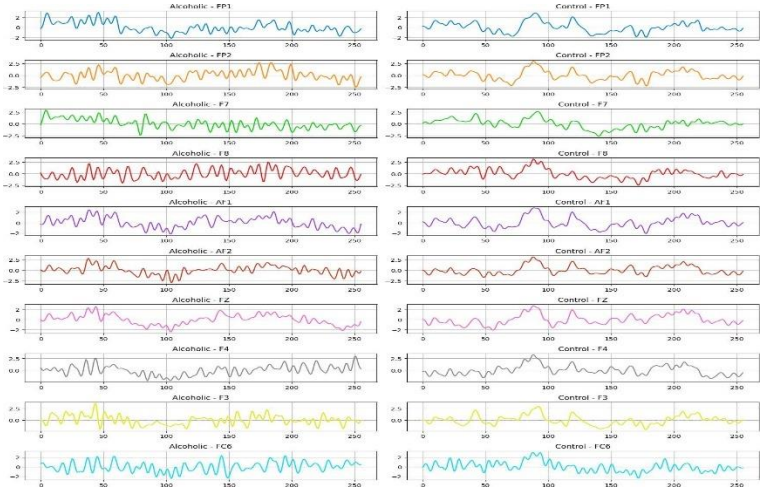

Fig. 2: Sample of Processed Data for an Alcoholic Participant (Left) and A Control Participant (Right) – First 10 Channels.

4. Deep learning model

Convolutional Neural Networks (CNNs) are deep learning models effective at capturing spatial hierarchies and patterns in data (Sadiq et al., 2022a). Initially designed for image recognition, CNNs have been applied to fields like signal processing, speech recognition, and EEG time series data (Sadiq et al., 2022a). By using convolutional filters, CNNs detect local features such as edges or textures, progressively extracting complex patterns while preserving spatial relationships, making them ideal for structured data. Each layer typically includes a convolution followed by a pooling to downsample data and focus on key features, enhancing robustness against noise and distortions. CNNs also pair well with architectures like RNNs or LSTMs to capture both spatial and temporal dependencies, enabling comprehensive sequence modeling (Cohen et al., 2023; Pain et al., 2023).
Recurrent Neural Networks (RNNs) are designed to process sequences of data, making them particularly effective in domains where context and temporal dependencies are key, such as language modeling, time series prediction, and speech recognition (Salankar et al., 2022; Anuragi & Sisodia, 2020). Unlike feedforward networks, RNNs have loops that allow information to persist, maintaining a hidden state that captures sequence data. At each time step, the RNN uses the current input and the previous hidden state to generate an output, which also influences future inputs (Shi et al., 2023). However, standard RNNs, like CNNs, often face challenges with long-range dependencies due to vanishing and exploding gradients during training (Salankar et al., 2022).
To overcome the limitations of traditional RNNs, Long Short-Term Memory networks (LSTMs) were introduced. LSTMs are a special kind of RNN capable of learning long-term dependencies (Salankar et al., 2022). They achieve this through a complex architecture that includes several gates: the input gate, which determines how much of the new input should be added to the cell state, the forget gate, which decides what information should be removed from the cell state, and the output gate, which con-trols the output of the cell state to the next hidden state and the final output. In summary, these gates manage the flow of information into and out of the cell, deciding what to keep in memory and what to discard, making LSTMs effective at capturing long-range dependencies in sequence data (Sadiq et al., 2022b).

4.1. Deep learning architectures

Considering the characteristics of RNNs, CNNs, and LSTMs, different models and architectures were tested in this study. Initially, an architecture based on the methodology proposed by (Salankar et al., 2022) was implemented, which aimed to classify EEG signals with the same dataset. In addition, many other architectures were tested, and four relevant models were selected for comparison, each designed to leverage different aspects of these neural networks.
The first model, Model 1, was an LSTM-based network inspired by the methodology proposed by (Salankar et al., 2022). This model consisted of two LSTM layers, each followed by dropout layers to prevent overfitting, which is crucial given the high variability in EEG data. The first LSTM layer, with 100 neurons and a ReLU activation function, focused on learning temporal patterns from the raw EEG signals. The second LSTM layer employed a sigmoid activation function to help the network retain long-term dependencies. Dropout layers with a rate of 0.4 were used to improve generalization. The network concluded with a dense layer for feature integration and a sigmoid output for classifying the EEG signals as alcoholic or control. Figure 3 shows the architecture of the neural network, taking the processed EEG signals as input data.

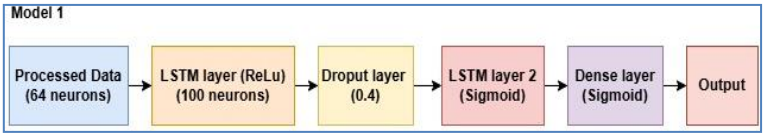

Fig. 3: Model 1: LSTM architecture.

The second model, Model 2, was a CNN designed to capture spatial features from the EEG signals. This model consisted of three Conv1D layers with 32, 64, and 128 filters, respectively. Each Conv1D layer was followed by batch normalization and max pooling to reduce the dimensionality while preserving important features. Dropout layers were included to prevent overfitting. After flattening the data, a dense layer with 128 neurons refined the learned features, and a final sigmoid layer classified the EEG signals into control or alcoholic. Figure 4 shows the architecture of the neural network.

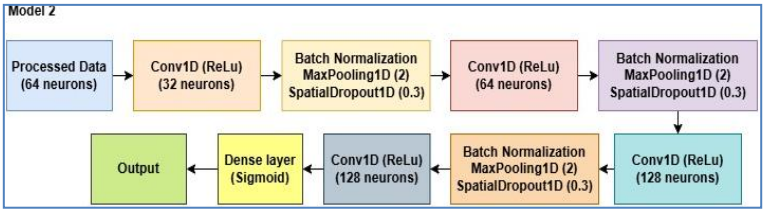

Fig. 4: Model 2: CNN Architecture.

The third model, Model 3, combined CNN and LSTM architectures (CNN-LSTM) to capture both spatial and temporal features of EEG signals for classification. The model consisted of two Conv1D layers with Leaky ReLU activation, batch normalization, max pooling, and spatial dropout to prevent overfitting. A bidirectional LSTM layer with 32 units was used to capture temporal dependencies, followed by a dense layer with 32 neurons for feature processing. Afterward, a second dense layer with a sigmoid activation function was added for final binary classification. Figure 5 shows the architecture of the neural network.

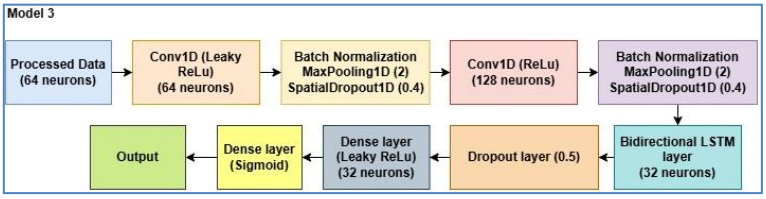

Fig. 5: Model 3: CNN-LSTM Architecture.

The last model, Model 4, combined a RNN with a bidirectional LSTM to capture both short and long-term dependencies in EEG signals. The model started with a SimpleRNN layer (64 units) followed by Leaky ReLU, batch normalization, and dropout (0.5) to prevent overfitting. A bidirectional LSTM layer (32 units) was added for capturing long-term dependencies, followed by another dropout layer. The model ended with two dense layers: one with 32 neurons and a final sigmoid layer for binary classification. Figure 6 shows the architecture of the neural network.

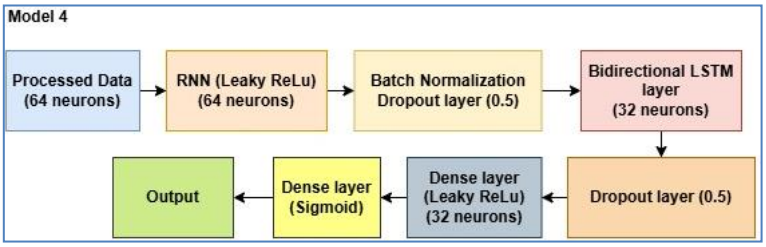

Fig. 6: Model 4: RNN-LSTM Architecture.

4.2. Experimental setup

Data Subset. Initially, this process was developed using a subset of 1,220 trials out of the total 10,617 trials, selecting 10 trials per patient. The deep learning models were built, and hyperparameters were tested using this subset. Once validated, the methodology was then applied to the complete dataset.

Data Partioning. The experimental dataset was partitioned into training, validation, and testing sets to evaluate the performance of the proposed models. GroupKFold with 5 splits was employed to ensure that no patient's data appeared in more than one fold, thus preserving the integrity of the evaluation process. Within each fold, the GroupShuffleSplit was utilized to further segregate the training set by allocating 10% of it to validation. This additional split within the training set ensured that training data remained completely distinct from validation data, maintaining the robustness of the validation methodology. The resultant data distribution was 70% for training, 10% for validation, and 20% for testing.

Optimization Algorithms. The robustness of the deep learning models was further assessed by implementing two different optimizers: Adam and SGD. Each of the four proposed architectures was tested with both optimizers to investigate their impact on model performance. For Models 2-4, SGD used a fixed learning rate of 0.001, while Adam was initialized at 0.0003 due to its adaptive learning capabilities. In contrast, Model 1 required a learning rate of 0.0001 to ensure a stable training behavior for both optimizers.

Implementation. The algorithms were implemented using TensorFlow 2.15. They were trained using batch sizes of 32, 64, 128, over 100, 200, and 300 epochs, and different activation and loss functions were tested to optimize model performance. To ensure reproducibility, a fixed random seed was used across all experiments for initializing the model parameters and shuffling the dataset. Class weights adjustment was implemented to address class imbalance, and early stopping was employed as a regularization technique to terminate training when the validation loss ceased to improve, thereby preventing overfitting.

Evaluation Metrics. To thoroughly assess the efficacy of the models, several metrics were employed: accuracy, precision, recall, and the F1 score. Accuracy provided a straightforward measure of overall performance, while precision and recall offered insights into the models' ability to identify true positive cases against the false positives and negatives, crucial for medical diagnostic applications. The F1

score, a harmonic mean of precision and recall, was particularly valued for its balance between precision and recall, providing a more comprehensive picture of model performance across various thresholds.

4.3. Experimental results

First experiments and hyperparameter tuning were done using the sample of 1,220 trials to limit computation time. Table 1 showcases the performance metrics for the models proposed. The performance metrics were calculated as the mean of the results obtained for each fold, ensuring a robust evaluation across different data subsets.

Table 1: Performance with A Sample Size of 1,220 Trials

Model	Optimizer	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Epochs
Model 1	Adam	64.84	71.39	66.28	64.01	300
Model 1	SGD	46.31	71.25	32.60	30.72	300
Model 2	Adam	82.21	89.89	80.89	85.10	300
Model 2	SGD	85.16	90.17	85.83	87.91	300
Model 3	Adam	89.18	93.22	89.34	91.22	300
Model 3	SGD	87.13	91.19	88.19	89.61	300
Model 4	Adam	83.44	90.07	82.87	86.26	300
Model 4	SGD	82.95	88.28	84.14	86.12	300

Once the optimal parameters were selected, the established models were trained with the complete dataset. Table 2 shows the performance metrics obtained for the same algorithms tested with 10,617 trials.

From Tables 1 and 2, we observe a significant discrepancy between these results and those reported by (Salankar et al., 2022), which reported a 93% accuracy. This level of accuracy could not be replicated with Model 1 due to several limitations in the base article. The referred study described the model architecture and training setup, stating that each participant's data consisted of 64 channels sampled for one second and outlining the use of an LSTM network with dropout and a dense layer for classification.

Table 2: Performance with the Complete Dataset of 10,617 Trials

Model	Optimizer	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Epochs
Model 1	Adam	47.36	56.98	41.07	34.13	100
Model 1	SGD	47.35	55.07	40.62	33.32	300
Model 2	Adam	96.57	98.28	96.29	97.28	100
Model 2	SGD	91.62	97.65	88.98	93.11	300
Model 3	Adam	97.63	98.60	97.66	98.13	100
Model 3	SGD	82.34	94.37	76.88	84.71	300
Model 4	Adam	96.18	97.76	96.2	96.97	100
Model 4	SGD	91.95	95.65	91.46	93.50	300

Also, Adam optimizer was implemented with learning rates ranging from 0.0001 to 0.1 and trained the model over 100 epochs (Salankar et al., 2022). However, it lacked critical details, such as whether data splitting for training, validation, and testing was done randomly or without repeating patients, as well as whether a separate validation set was used. Furthermore, there was no repository or code provided to test or validate their results. Despite exhaustive trials, the accuracy results consistently fell below 70%. Consequently, this model did not provide a reliable baseline, contrary to our initial expectations. Due to these results, we performed a wide search for new models that could potentially improve performance, which led to Models 2-4.

In general, the results obtained show a significant improvement in performance metrics when training the models with the complete dataset of 10,617 trials (Table 2) compared to the sample size of 1,220 trials (Table 1). This is a consequence of a broader distribution of data for training provided by a larger dataset. Increasing the number of trials helps the model learn more complex patterns, reducing overfitting and improving generalization.

Additionally, it is possible to observe that Model 3 with Adam optimizer presents superior performance metrics, regardless of the number of samples used. This algorithm implements a CNN with LSTM architecture, which enables a robust analysis of both spatial and temporal characteristics of neural activity, making it well-suited for processing EEG data. Figures 7 and 8 illustrate the performance of Model 3 optimized with Adam, for one of the five test folds, using 1,220 trials and 10,617 trials, respectively (the complete accuracy and loss trends for each dataset are displayed in Figures A1 and A2). The accuracy graphs show a consistent increase in both training and validation accuracy, which suggests strong generalization capabilities of the models. On the other hand, the loss graphs show a significant drop in the early epochs, with both training and validation losses stabilizing over time. This steady performance throughout the training period highlights the effectiveness of the chosen architecture and hyperparameters. Specifically, Figure 8 shows higher accuracy values reached within fewer epochs and lower loss values compared to Figure 7, suggesting that using more trials during the training process results in a higher model utility.

Overall, the algorithms optimized with Adam demonstrated superior performance compared to those using SGD, as shown in Tables 1 and 2. Additionally, the SGD optimizer requires a larger number of epochs to achieve similar results as the Adam optimizer when using the complete dataset for training (Table 2).

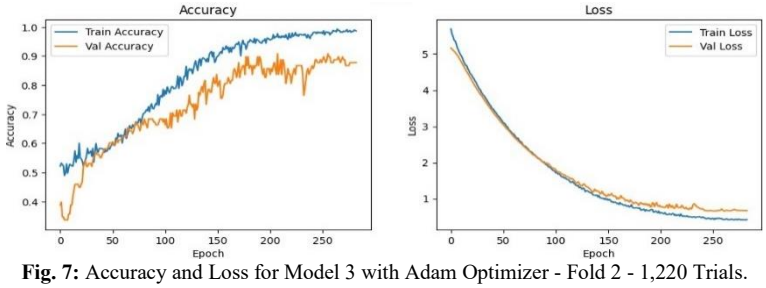

Fig. 7: Accuracy and Loss for Model 3 with Adam Optimizer - Fold 2 - 1,220 Trials.

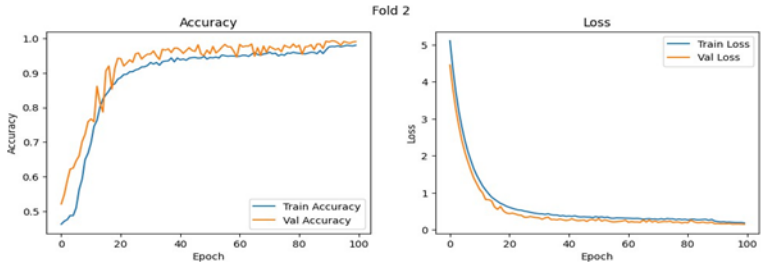

Fig. 8: Accuracy and Loss for Model 3 with Adam Optimizer - Fold 2 - 10,617 Trials.

According to Sadiq et al. (2021a), Adam adapts the learning rates based on the average of recent gradients in the computation, ensuring a more effective optimization process. Moreover, Adam incorporates momentum by keeping track of past gradients, which stabilizes the updates. This results in faster convergence and improved handling of problems with sparse gradients or noisy data, as outlined by (Yeh et al., 2023). Therefore, Adam's methodical approach to adjusting learning rates and leveraging momentum enhances model performance across various tasks and datasets.

When fewer trials are used, differences between algorithms become more apparent, allowing for a clearer distinction in performance. Therefore, hyperparameter tuning focused on models trained with the subset of 1,220 trials. Table 2 highlights that two dense layers achieve the best performance, with the highest accuracy (89.18%) and F1 score (91.22%). Binary Cross-Entropy consistently outperformed Categorical Cross-Entropy. Dropout values of 0.4–0.5 balanced overfitting and generalization effectively. A batch size of 64 delivered the best results, while increasing to 128 slightly reduced performance. The learning rate of 0.0003 ensured stable convergence across experiments. LeakyReLU combined with Sigmoid consistently performed better than ReLU, emphasizing the importance of carefully tuned hyperparameters for optimal performance.

5. Differential privacy mechanism

We now move into the second phase of the study, where we integrate differential privacy techniques to ensure data privacy while looking to avoid significantly compromising model performance.

Differential privacy (DP) provides a solid standard for guaranteeing privacy in algorithms that operate on aggregate databases. Two approaches will be investigated: (ϵ, δ)- Differential Privacy ((ϵ, δ)-DP) and Gaussian Differential Privacy (μ-GDP). These methods introduce Gaussian noise to protect individual data points in the training dataset, but differ in how they count privacy. The objective of this study is to compare these techniques to determine which one offers a more effective balance between privacy protection and the preservation of model utility.

5.1. (ϵ, δ)-Differential privacy

Differential privacy is defined in terms of adjacent databases, which in typical experiments differ by a single entry. The concept of (ϵ, 0)-differential privacy, or ϵ-DP, offers a strong privacy guarantee (Browarska et al., 2021).

5.2. μ-Gaussian differential privacy

The differential privacy ensures that an adversary has a limited ability to identify the presence or absence of any individual in a dataset, depending on the values of ϵ and δ. Let S and S′ be neighboring datasets; then, the adversary seeks to distinguish between two probability distributions, M (S) and M (S′). Considering this, differential privacy can be defined through a hypothesis testing problem (Sherstinsky, 2020):

H0: the true dataset is S versus H1: the true dataset is S′.

Authors (Sherstinsky, 2020) proposed using a trade-off function f between type I error (the probability of incorrectly rejecting H0 when it is true) and type II error (the probability of incorrectly accepting H0 when H1 is true). The function f quantifies the minimum type II error achievable for a given type I error level α, integrating concepts from differential privacy. In f -DP, privacy is thus evaluated in terms of the function f. The larger the trade-off function is, the more difficult the hypothesis testing problem is; hence, there is a stronger privacy guarantee (Anuragi & Sisodia, 2020).

5.3. Optimization algorithms for differential privacy

Differentially Private Stochastic Gradient Descent (DP-SGD) and Differentially Private Adam (DP-Adam) are adaptations of the classic SGD and Adam optimization algorithms, respectively, designed to provide differential privacy guarantees during the training of machine learning models (Browarska et al., 2021). Both methods achieve privacy by introducing noise to the gradients during the backpropagation phase and applying gradient clipping to limit the influence of any single data point, thus reducing the risk of sensitive information leakage (Browarska et al., 2021)

for the tests. Next, we analyze the results of incorporating (ϵ, δ)-DP and GDP techniques into deep learning models. Finally, we conduct statistical tests to evaluate the impact of key parameters on model performance and privacy guarantees.

5.4. Hyperparameters for differential privacy

For the implementation of differential privacy mechanisms, several parameters play crucial roles in balancing privacy and model performance. First, the learning rate determines the speed at which the model updates its weights in response to the error it observes, impacting both convergence and stability (Li & Wu, 2022). Second, the clipping norm sets a threshold for the L2 norm of the gradient for each training example, limiting the sensitivity of the algorithm to individual data points and thereby controlling the influence of any single observation on the model's updates (Farsi et al., 2021b). The noise multiplier specifies the amount of Gaussian noise added to the clipped gradients, directly influencing the degree of privacy and the performance of the trained model (Browarska et al., 2021). A higher noise

level increases privacy but may decrease model utility. Lastly, the number of microbatches is crucial in differentially private optimization because it affects how noise is added to gradients and how the privacy budget is managed (Browarska et al., 2021). By splitting a batch into microbatches, the noise is applied to the average gradient for each microbatch, influencing the balance between privacy and model utility. These parameters must be carefully tuned to achieve an optimal trade-off between privacy protection and effective learning in DP-SGD and DP-Adam implementations.

5.5. Experiments

5.5.1. Experimental setup

The deep learning architectures were evaluated using the differential privacy approaches. For the testing process, the models were evaluated using both DP-SGD and DP-Adam optimizers. Privacy accounting was conducted following the frameworks of (ϵ, δ)-DP and μ-GDP. For (ϵ, δ)-DP, the associated ϵ value was calculated using a δ value of 10−5, while for μ-GDP, ϵ was computed based on the value of μ obtained.

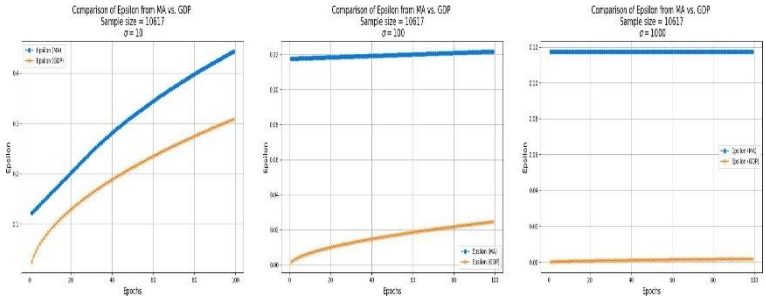

Fig. 9: Comparison of ϵ Values from Moments Accountant (MA) vs. Gaussian Differential Privacy (GDP) for Different Noise Levels (Σ) with A Sample Size of 10,617.

Moments Accountant (MA), a method for tracking cumulative privacy loss during training, which allows for tighter and more accurate privacy analysis.
Figure 9 illustrates the evolution of ϵ over 100 epochs for three noise settings: σ = 10, σ = 100, and σ = 1000, respectively.
Experiments were conducted using three data samples: a sample of 1,220 data points, a sample of 5,508 data points, and a full dataset consisting of 10,617 data points. The same performance metrics were used for evaluation. Additionally, early stopping was applied to prevent overfitting and to ensure optimal training performance.
Parameter tuning was performed to optimize privacy and saturation point as quickly as MA (Sadiq et al., 2022a). Remarkably, despite these high noise levels, model performance metrics for Model 3 remain unaffected, compared to the results presented in Table 3 without differential privacy, suggesting that the model's ability to generalize is not compromised by noise perturbations.

Table 3: Performance Evaluation and Privacy Preservation with Different Noise Levels for Model 3 with DP-Adam

Noise	Accuracy	Precision	Recall	Noise	Accuracy	Precision
10	97.14	98.37	97.12	97.74	0.445	0.310
50	97.30	98.44	97.29	97.86	0.134	0.053
100	96.99	98.60	96.64	97.61	0.122	0.025
500	97.16	98.59	96.94	97.76	0.118	0.004
1000	97.30	98.37	97.37	97.87	0.118	0.002
3000	97.40	98.59	97.31	97.95	0.118	0.0004
5000	97.15	98.47	97.04	97.75	0.118	0.00022
20000	97.15	98.57	96.94	97.75	0.118	0.0000199

Considering that high noise levels do not significantly impact model performance, σ was set to 20,000, providing strong privacy guarantees by achieving small ϵ values. Table 4 presents the results for Models 2-4 with each optimizer (DP-SGD and DP-Adam) and sample size tested at 1,220, 5,508, and 10,616. The best performance was obtained with an L2 norm clip of 1 and 64 microbatches. As seen in the table, increasing sample sizes leads to improved model performance, due to better generalization and stability, and it also enhances privacy, as a smaller value of ϵ is obtained for GDP. In most cases, DP-Adam demonstrates better results than DP-SGD, and Model 3 shows the best results with the Adam optimizer and the complete dataset, indicating its strong capacity to capture complex patterns even under differential privacy constraints. Model 2 also demonstrates solid performance, demonstrating its utility with privacy mechanisms. Furthermore, the results show that DP-SGD outperforms standard SGD without differential privacy, highlighting not only its effectiveness in maintaining privacy but also its potential to enhance model performance through the noise-induced regularization effect.

5.5.2. Statistical analysis

Regression models were constructed across 90 experiments to quantify the influence of factors such as sample size, optimizer algorithm, and model architecture on performance metrics and privacy guarantees. By employing multiple linear regression, we sought to evaluate how important each predictor is and to understand their effects and possible interactions.

Table 4: Overall Performance Results with Differential Privacy Mechanisms

Sample Size	Optimizer	Model	Noise	L2 Norm Clip	Microbatches	Accuracy (%)	Precision (%)	Recall	F1-Score (%)	Time (min)	Epsilon MA	Epsilon GDP
1220	SGD	2	20000	1	64	83.64	91.03	82.56	86.54	1.9	0.117	0.00014
1220	SGD	3	20000	1	64	77.54	86.71	76.59	81.19	4.1	0.117	0.00014
1220	SGD	4	20000	1	64	63.85	78.79	60.87	66.00	3.1	0.117	0.00014
1220	Adam	2	20000	1	64	79.36	87.84	78.54	82.77	1.4	0.117	0.00014

1220	Adam	3	20000	1	64	81.46	88.58	81.56	84.89	4.2	0.117	0.00014
1220	Adam	4	20000	1	64	66.33	79.38	64.33	70.31	3.7	0.117	0.00014
5508	SGD	2	20000	1	64	94.34	96.12	94.92	95.53	7.2	0.117	0.00004
5508	SGD	3	20000	1	64	90.04	96.39	87.65	91.80	19.7	0.117	0.00004
5508	SGD	4	20000	1	64	78.55	89.76	74.83	81.53	33.2	0.117	0.00004
5508	Adam	2	20000	1	64	93.78	96.25	93.89	95.05	7.2	0.117	0.00004
5508	Adam	3	20000	1	64	93.54	96.86	92.86	94.82	19.1	0.117	0.00004
5508	Adam	4	20000	1	64	88.07	93.18	87.69	90.32	31.1	0.117	0.00004
10617	SGD	2	20000	1	64	96.93	98.14	97.02	97.58	17.2	0.117	0.0000199
10617	SGD	3	20000	1	64	95.88	98.51	94.97	96.70	39.1	0.117	0.0000199
10617	SGD	4	20000	1	64	82.85	93.25	78.59	85.01	58.9	0.117	0.0000199
10617	Adam	2	20000	1	64	96.82	98.01	96.98	97.49	18.5	0.117	0.0000199
10617	Adam	3	20000	1	64	97.21	98.48	97.13	97.78	41.8	0.117	0.0000199
10617	Adam	4	20000	1	64	94.17	96.80	93.95	95.35	65.5	0.117	0.0000199

Table 5 presents the results of a multiple linear regression analysis used to identify the influence of various predictors on model accuracy. Significance codes are provided as a quick reference for the statistical relevance of each predictor, with asterisks denoting levels of significance. Here we observe that the sample size has a positive and significant effect on the accuracy, as indicated by its estimate of 0.0019 and a highly significant p-value ($p < 2e\text{-}16$). This suggests that as the sample size increases, model accuracy improves. On the other hand, the predictor DP-SGD shows a negative impact on accuracy with an estimate of -3.012 and a p-value of 0.00162, indicating a statistically significant decrease in accuracy when compared to the reference level, DP-Adam. The predictors Model 3 and Model 4 show varying effects: Model 3 is not statistically significant ($p = 0.17905$), while Model 4 has a strong negative effect on accuracy, with an estimate of -11.84 and a highly significant p-value ($p < 2e\text{-}16$). Overall, the results highlight the importance of the sample size and the negative influence of DP-SGD and Model 4 on accuracy.

Table 5: Regression Coefficients for Predictors of Model Accuracy

Predictor	Estimate	Std. Error	t value	Pr(> |t|)
(Intercept)	81.08	1.158	70.029	< 2e-16 ***
Sample size	0.0019	0.00012	16.150	< 2e-16 ***
DP-SGD	-3.012	0.925	-3.256	0.00162 **
Model 3	-1.535	1.133	-1.355	0.17905
Model 4	-11.84	1.133	-10.453	< 2e-16 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1.

Table 6 presents the regression coefficients for predictors of the model's F1-score. As well as the results shown in Table 5, sample size has a positive and significant effect (estimate: 0.0017, $p < 2e\text{-}16$), while Optimizer DP-SGD negatively impacts F1-Score (estimate: 2.990, $p = 0.00283$). Model 3 is not significant ($p = 0.27775$), but Model 4 shows a strong negative influence (estimate: -11.07, $p = 1.33e\text{-}14$).

Table 7 summarizes the regression coefficients for predictors of ϵ from GDP. Sample size shows a negative and highly significant effect (estimate: 1.250e08, $p < 2e\text{-}16$). The other predictors (Optimizer DP-SGD, Model 3, and Model 4) show no significant effect, as indicated by their p-values of 1. The privacy guarantee remains constant regardless of the architecture and optimizer used.

Table 6: Regression Coefficients for Predictors of Model F1-Score

Predictor	Estimate	Std. Error	t value	Pr(> |t|)
(Intercept)	84.11	1.217	69.131	< 2e-16 ***
Sample size	0.0017	0.00013	13.499	< 2e-16 ***
Optimizer SGD	-2.990	0.9722	-3.076	0.00283 **
Model 3	-1.301	1.191	-1.092	0.27775
Model 4	-11.07	1.191	-9.300	1.33e-14 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1.

Table 7: Regression Coefficients for Predictors of ϵ from GDP

Predictor	Estimate	Std. Error	t value	Pr(> |t|)
(Intercept)	1.389e-04	5.777e-06	24.05	< 2e-16 ***
Sample size	-1.250e-08	6.009e-10	-20.81	< 2e-16 ***
Optimizer SGD	-2.402e-20	4.616e-06	0.00	1
Model 3	3.627e-20	5.653e-06	0.00	1
Model 4	3.782e-20	5.653e-06	0.00	1

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1.

6. Ethical considerations and practical challenges

6.1. Ethical concerns

The use of EEG data for neurological classification tasks, such as distinguishing between alcoholic and non-alcoholic subjects, raises several ethical considerations that must be addressed to ensure responsible research practices.

- Informed Consent:

The dataset used in this study originates from the Neurodynamic Laboratory at the State University of New York, distributed via the UCI Machine Learning Repository. It is assumed that proper informed consent was obtained from all participants at the time of data collection. Nevertheless, secondary data users must be diligent in adhering to ethical guidelines surrounding the use and sharing of human subject data. This includes anonymizing the dataset, refraining from re-identification attempts, and using the data strictly for research and educational purposes.

- Data Bias and Representativeness:

The dataset includes 122 participants, with a greater number of alcoholic subjects (77) compared to control subjects (45). This class imbalance, while mitigated through techniques such as class weighting during model training, may introduce bias into the classification results. Furthermore, the demographic diversity of the participants (e.g., age, gender, ethnicity) is not specified, which may limit the generalizability of the findings. Models trained on such datasets might perform poorly on more diverse real-world populations, thus raising concerns about fairness and equity in medical applications.

- Clinical Misuse and Interpretability:

Another ethical concern lies in the potential misuse or overreliance on automated EEG classification tools in clinical settings. Deep learning models, particularly those involving CNN-LSTM architectures, often operate as "black boxes," offering limited interpretability. Deploying such models without appropriate clinical oversight could lead to misdiagnosis or inappropriate treatment decisions. Hence, these tools should be used to support—not replace—clinical judgment, and any diagnostic application should involve clinicians in the loop.

6.2. Practical challenges

The integration of deep learning in analyzing sensitive health data, such as EEG signals for AUD detection, raises several ethical concerns, chief among them being privacy, informed consent, and algorithmic bias. As highlighted by Floridi and Taddeo (2018), data ethics encompasses the moral obligations involved in collecting, analyzing, and deploying data-driven technologies, particularly in contexts involving vulnerable populations and high-stakes decisions.

- Computational Resources and Training Time:

Training deep learning models, especially on large EEG datasets with high temporal resolution, demands substantial computational power. Models like the CNN-LSTM architecture (Model 3), while delivering superior performance, require significant memory and processing time, particularly when multiple optimizers and hyperparameter combinations are tested. In our experiments, batch sizes, epochs, and optimizer configurations were optimized iteratively, which is computationally intensive. This imposes a barrier to adoption in resource-limited research environments or clinical institutions without dedicated hardware such as GPUs or TPUs.

- Data Preprocessing Complexity:

EEG data preprocessing is a critical but time-consuming task. Ensuring signal integrity through filtering, transformation (e.g., FFT), and artifact removal requires domain expertise. Any errors or oversights in this pipeline may significantly degrade model performance. Additionally, standardization and reproducibility of preprocessing procedures across studies remain ongoing challenges in the field of neuroinformatics.

- Clinical Integration and Validation:

Despite promising experimental results, translating EEG-based classification models into clinical practice remains a nontrivial task. Clinical validation requires extensive testing across varied populations and settings. Moreover, regulatory approval for machine learning systems in healthcare typically demands clear documentation, interpretability, and evidence of safety and efficacy. Finally, clinicians may face challenges in trusting and interpreting deep learning predictions unless supported by explainable AI (XAI) techniques and user-friendly interfaces.

This research is an example of the interdisciplinary integration of computer science and health science, consistent with the Journal's mission to advance data-driven innovation in biomedical applications. Using deep learning techniques and privacy-preserving methods such as differential privacy, this work addresses key challenges in clinical neuroscience, especially the detection of alcohol use disorder (AUD) from EEG signals. The proposed models not only maximize diagnostic accuracy but also maintain data privacy standards suitable for use in ethical healthcare contexts. Furthermore, expanding these methods will support the development of safe, adaptive diagnostic tools that can be integrated into global healthcare systems, particularly in settings where patient trust, data protection, and resource constraints are paramount. This interdisciplinary approach highlights the potential of AI-based solutions to transform public health and clinical decision-making.

Future research can build on these findings by exploring advanced privacy mechanisms such as Zero-Concentrated Differential Privacy (zCDP), which provides stronger evidence and less noise than traditional (ε,δ)-DP, thus improving model efficiency. Federated learning – another approach that allows isolated training with local EEG data without sharing patient data – is ideal for privacy-related clinical settings. To increase generalizability and reduce bias, validation should be performed across the entire population using multiple EEG data from different regions. Furthermore, the integration of privacy-preserving models into real-time EEG monitoring systems allows continuous detection of AUD in clinical and remote settings. This will require models that are lightweight and efficient, with a good balance between accuracy, privacy, and real-world feasibility.

7. Conclusions

Achieving a balance between privacy and utility is a central challenge in deep learning, especially when working with sensitive data, as ensuring the privacy of patient information is critical for building trust and enabling the adoption of these tools in clinical settings. This work aims to incorporate differential privacy mechanisms in deep learning models for the detection of AUD in EEG signals. Experiments demonstrate that GDP stands out by requiring less noise than (ϵ, δ)-DP, with the Moments Accountant, to achieve an equivalent privacy guarantee, thus preserving higher utility. Moreover, DP-Adam excels in maintaining strong performance while ensuring privacy, compared to DP-SGD, making it particularly effective for deep learning models requiring fast convergence and adaptive learning rates. Additionally, increasing the sample size further enhances both model performance and privacy guarantees. Larger datasets enable noise distribution with minimal impact on utility, reinforcing the importance of sufficient data in privacy-preserving models. Overall, choosing the right privacy mechanism and hyperparameters is critical to optimizing the privacy-utility tradeoff.

This study opens the door to exploring additional differential privacy techniques, which might offer tight guarantees and flexibility, and Zero-Concentrated Differential Privacy (zCDP), known for reduced noise requirements and enhanced utility preservation. Future work could also focus on differential privacy mechanisms that consider the inclusion or exclusion of larger data groups. As this research involves up to 120 trials per patient, evaluating the impact of including or excluding 120 data points on privacy guarantees could offer deeper insights into privacy scalability and effectiveness for structured datasets.

References

[1] Akbari, H., Sadiq, M. T., Payan, M., Esmaili, S. S., Baghri, H., & Bagheri, H. (2021). Depression detection based on geometrical features extracted from SODP shape of EEG signals and binary PSO. Traitement du Signal, 38, 43–46. https://doi.org/10.18280/ts.380102.

[2] Shi, W., Feng, H., Zhang, X., & Yeh, C.-H. (2023). Amplitude modulation multiscale entropy characterizes complexity and brain states. Chaos, Solitons & Fractals, 173, 113646. https://doi.org/10.1016/j.chaos.2023.113646.

[3] Sadiq, M. T., Akbari, H., Rehman, A. U., Nishtar, Z., Masood, B., Ghazvini, M., Hamedi, N., & Kaabar, M. K. A. (2021). Exploiting feature selection and neural network techniques for identification of focal and nonfocal EEG signals in TQWT domain. Journal of Healthcare Engineering, 2021, 1–24. https://doi.org/10.1155/2021/6283900.

[4] Yeh, C. H., Zhang, C., Shi, W., Lo, M. T., Tinkhauser, G., & Oswal, A. (2023). Cross-frequency coupling and intelligent neuromodulation. Cyborg and Bionic Systems, 4, 34. https://doi.org/10.34133/cbsystems.0034.

[5] Sadiq, M. T., Akbari, H., Siuly, S., Li, Y., & Wen, P. (2022). Alcoholic EEG signals recognition based on phase space dynamic and geometrical features. Chaos, Solitons & Fractals, 158, 112036. https://doi.org/10.1016/j.chaos.2022.112036.

[6] Cohen, S., Katz, O., Presil, D., Arbili, O., & Rokach, L. (2023). Ensemble learning for alcoholism classification using EEG signals. IEEE Sensors Journal, 23, 17714–17724. https://doi.org/10.1109/JSEN.2023.3279904.

[7] Pain, S., Roy, S., Sarma, M., & Samanta, D. (2023). Detection of alcoholism by combining EEG local activations with brain connectivity features and graph neural network. Biomedical Signal Processing and Control, 85, 104851. https://doi.org/10.1016/j.bspc.2023.104851.

[8] Salankar, N., Qaisar, S. M., Pławiak, P., Tadeusiewicz, R., & Hammad, M. (2022). EEG based alcoholism detection by oscillatory modes decomposition, second order difference plots and machine learning. Biocybernetics and Biomedical Engineering, 42, 173–186. https://doi.org/10.1016/j.bbe.2021.12.009.

[9] Anuragi, A., & Sisodia, D. (2020). Empirical wavelet transform based automated alcoholism detecting using EEG signal features. Biomedical Signal Processing and Control, 57, 101777. https://doi.org/10.1016/j.bspc.2019.101777.

[10] Sadiq, M. T., Aziz, M. Z., Almogren, A., Yousaf, A., Siuly, S., & Rehman, A. U. (2022). Exploiting pre-trained CNN models for the development of an EEG-based robust BCI framework. Computers in Biology and Medicine, 143, 105242. https://doi.org/10.1016/j.compbiomed.2022.105242.

[11] Browarska, N., Kawala-Sterniuk, A., Zygarlicki, J., Podpora, M., Pelc, M., Martinek, R., & Gorzelanczyk, E. J. (2021). Comparison of smoothing filters' influence on quality of data recorded with the Emotiv Epoc Flex brain–computer interface headset during audio stimulation. Brain Sciences, 11, 98. https://doi.org/10.3390/brainsci11010098.

[12] Anuragi, A., Sisodia, D., & Pachori, R. B. (2021). Automated FBSE-EWT based learning framework for detection of epileptic seizures using time-segmented EEG signals. Computers in Biology and Medicine, 136, 104708. https://doi.org/10.1016/j.compbiomed.2021.104708.

[13] Farsi, L., Siuly, S., Kabir, E., & Wang, H. (2021). Classification of alcoholic EEG signals using a deep learning method. IEEE Sensors Journal, 21, 3552–3560. https://doi.org/10.1109/JSEN.2020.3026830.

[14] Li, H., & Wu, L. (2022). EEG classification of normal and alcoholic by deep learning. Brain Sciences, 12(6), 819. https://doi.org/10.3390/brainsci12060778.

[15] Ahmadi, N., Pei, Y., & Pechenizkiy, M. (2017). Detection of alcoholism based on EEG signals and functional brain network features extraction. IEEE International Conference Proceedings, 2017-June, 179–184. https://doi.org/10.1109/CBMS.2017.46.

[16] Farsi, L., Siuly, S., Kabir, E., & Wang, H. (2021). Classification of alcoholic EEG signals using a deep learning method. IEEE Sensors Journal, 21(3), 3552–3560. https://doi.org/10.1109/JSEN.2020.3026830.

[17] Balle, B., Barthe, G., Gaboardi, M., & Hsu, J. (2020). Privacy amplification by subsampling: Tight analyses via couplings and divergences. Journal of Privacy and Confidentiality, 10(S1), 1–39. https://doi.org/10.29012/jpc.726.

[18] Sherstinsky, J. (2020). Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network. Physica D: Nonlinear Phenomena, 404, 132306. https://doi.org/10.1016/j.physd.2019.132306.

[19] Brunnhuber, F., Slater, J., Goyal, S., Amin, D., Thorvardsson, G., Freestone, D. R., & Richardson, M. P. (2020). Past, present and future of home video-electroencephalographic telemetry: A review of the development of in-home video-electroencephalographic recordings. Epilepsia, 61(S1), S3–S10. https://doi.org/10.1111/epi.16578.

[20] Ko, L. W., Komarov, O., Lai, W. K., Liang, W. G., & Jung, T. P. (2020). Eyeblink recognition improves fatigue prediction from single-channel forehead EEG in a realistic sustained attention task. Journal of Neural Engineering, 17(3), 036015. https://doi.org/10.1088/1741-2552/ab909f.

[21] Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In Theory of Cryptography Conference (TCC), Lecture Notes in Computer Science (Vol. 3876, pp. 265–284). Springer. https://doi.org/10.1007/11681878_14.

[22] Porjesz, B., Rangaswamy, M., Kamarajan, C., Jones, K. A., Padmanabhapillai, A., & Begleiter, H. (2005). The utility of neurophysiological markers in the study of alcoholism. Clinical Neurophysiology, 116(5), 993–1018. https://doi.org/10.1016/j.clinph.2004.12.016.

[23] Floridi, L., & Taddeo, M. (2018). What is data ethics? Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 376(2133), 20180081. https://doi.org/10.1098/rsta.2018.0081.