ADHD Disease Classification Using AI / ML

Avinash K. Ikhar, Maitraee Kengale, Muskan Naidu, Swayam Bhute, Sujal Mandpe
Assistant Professor, Student, Student, Student, Student
Electronics And Telecommunication Department
J D College of Engineering and Management, Nagpur.

***Abstract*—The diagnosis of Attention-Deficit/Hyperactivity Disorder (ADHD) has traditionally been performed through subjective clinical interviews and behavioral observations, which are often time-consuming and susceptible to observation bias. To overcome this challenge, we introduce an objective machine learning-based clinical decision support system that uses electroencephalography (EEG) signal analysis to identify the neurophysiological correlates of ADHD. This work utilizes a large-scale dataset consisting of 2,166,383 EEG samples. We analyzed data from 19 conventional EEG channels and used an automated feature engineering pipeline to derive 46 more statistical and interaction features, thereby creating a rich 65-dimensional feature space. Using strict 5-fold cross-validation, we compared the performance of various classification models such as 1D Convolutional Neural Networks, Random Forest, and Gradient Boosting . Among these, the Extreme Gradient Boosting (XGBoost) classifier showed the best performance, achieving a classification accuracy of 77.84% and an ROC-AUC score of 0.8541. Most importantly, for a medical screening application, the model showed a balanced clinical profile, achieving both sensitivity and specificity of 83.01%. Feature ablation and importance analysis also confirmed the model's focus on prefrontal and central channel data, consistent with known neurophysiological ADHD patterns. The proposed system offers a fast (<100ms inference time), scalable, and completely objective screening system that can potentially aid medical professionals in early ADHD detection and diagnostic triaging.**

***Index Terms*—Attention-Deficit/Hyperactivity Disorder (ADHD), Electroencephalography (EEG), Machine Learning, XGBoost, Feature Engineering, Clinical Decision Support Systems.**

I. INTRODUCTION

Attention-Deficit/Hyperactivity Disorder (ADHD) is a common neurodevelopmental disorder characterized by the presence of persistent and disruptive symptoms of inattention, hyperactivity, and impulsivity. Estimated to affect 5% to 10% of the world's population, ADHD is a serious hindrance to academic, occupational, and social success. Notwithstanding its prevalence, the traditional medical approach to diagnosing ADHD has remained essentially the same for the past several decades, with healthcare professionals continuing to use subjective approaches as the primary means of diagnosis. These approaches are mainly composed of clinical interviews, behavioral observations, and standardized questionnaire-based screening instruments such as the Conners Scale or the Adult ADHD Self-Report Scale (ASRS). In some clinical settings, Computerized Continuous Performance Tests (CPTs) are also used; these tests are mainly focused on the measurement of behavioral responses rather than the underlying physiological mechanisms.

Although these traditional methods form the backbone of modern psychiatric practice, they are also fraught with limitations that make it difficult to achieve a timely and accurate diagnosis. The most serious of these limitations is the subjective nature of behavioral observation. Because symptoms are continuous and often comorbid with other mood or learning disorders, the subjective nature of these observations is prone to variability. Studies have shown that this subjective variability can lead to rates of misdiagnosis of up to 15% to 20% in some populations. In addition, comprehensive neuropsychological assessments are also time-consuming, requiring a total of 6 to 8 hours of diagnostic workup, thus imposing a significant burden on the healthcare system and the patient. Because of its time-consuming nature and high costs, this approach is also significantly limited in terms of accessibility, especially in developing countries. Finally, these approaches are also susceptible to socioeconomic and cultural biases that can lead to inequities in healthcare outcomes.

To overcome the systemic bottlenecks, we experience in the field of medicine, scientists are now seeking objective and quantifiable biomarkers. EEG has been identified as a very promising method in this respect. EEG is a non-invasive method of recording the electrical activity of the brain generated by the synchronized firing of neurons. It measures the activity of the brain in different frequency bands: Delta (0.5-4 Hz), Theta (4-8 Hz), Alpha (8-12 Hz), Beta (12-30 Hz), and Gamma (above 30 Hz). There are a large number of neurophysiological studies that have identified ADHD-specific EEG patterns. One of the most striking of these is the high Theta/Beta ratio in many ADHD patients, indicating a predominance of slow-wave activity (drowsiness or relaxation responses) and a deficiency of fast-wave activity (alert concentration responses), particularly in the prefrontal cortex. Other biomarkers include reduced posterior alpha activity and specific localized irregularities in the temporal lobes.

These biological differences make EEG a viable biomarker, but the process of extracting useful clinical information from raw EEG data is computationally intensive. The data is noisy, high-dimensional, and has spatiotemporal patterns that are not immediately visible to the naked eye. This makes manual analysis by neurologists time-consuming and prone to human error. This is where artificial intelligence and healthcare meet to make a true technological leap. Recent advances in machine learning make it possible to quickly and automatically process such complex brain signals. By using sophisticated classification algorithms on EEG data, we can programmatically analyze brain signals and automatically extract standardized and repeatable insights much faster than current methods. Ensemble methods such as Random Forest and Gradient Boosting are particularly well-suited to medical diagnosis because they are robust to nonlinearities, avoid overfitting, and provide interpretable feature importance rankings, which are essential for clinical interpretation.

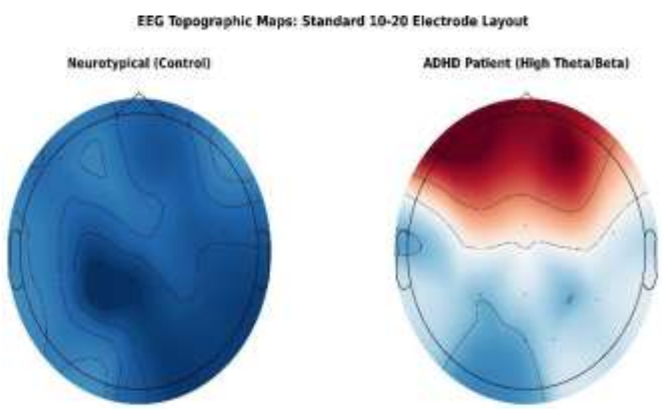

Fig. 1. Topographic EEG brain map comparison illustrating elevated Theta wave activity and reduced Beta wave activity in an ADHD brain versus a neurotypical control

The driving force behind this research is the pressing need to upgrade ADHD diagnosis. A transition from subjective observation to objective measurement can significantly improve diagnostic precision and reduce dramatically the time gap between screening and treatment. In addition, automated systems can increase access to advanced diagnostics, allowing primary care or clinics in developing regions to conduct high-level neurological screenings even when specialists are not physically present. The goal is not to replace specialists but to provide them with highly interpretable and reliable data-driven tools that they can confidently use, demonstrating the real-world utility of artificial intelligence in medicine.

To address these issues, this research proposes a machine learning-based classification system that aims to identify ADHD using EEG signal processing. The primary objectives are:

1. Develop and optimize a machine learning system that can effectively classify EEG signals into either ADHD or Control categories using a massive dataset of over 2.1 million examples.
2. Strategically design feature engineering to identify the most critical EEG channels and neurophysiological variables that contribute to the distinction between people with ADHD.
3. Integrate the predictive model within a useful web-based clinical decision support tool to facilitate easy data entry.
4. Develop an automated clinical reporting system that can translate the model's intricate results into standardized PDF reports for healthcare professionals.

II. Related Work

Much has been spawned in the area of research as artificial intelligence begins to find its place in the evaluation and diagnosis of ADHD. To understand the development of the area, Zhao et al. [2] traced back the evolution from simple screening tools to more detailed differential diagnosis using AI. Sun et al. [3] conducted a scoping review with a focus on AI applications for children, while Wang et al. [4] identified the global research trends using hotspots, including clinical applications and new technologies.

A. Clinical, Text, and Behavioral Machine Learning

Even before neurophysiology emerged, most studies relied on clinical and behavioral inputs for training machine learning models. Dai and Hsu [15] employed feature selection for the classification of ADHD through parent reports, self-reports, and traditional neuropsychological tests. Similarly, Alsharif et al [9]. evaluated the performance of text-based features extracted from clinical text using predictive ML models.

To counter the subjective nature of reporting, some studies have utilized physical and behavioral tracking. Amado-Caballero et al. [15] used deep learning on actimetry signals and analyzed occlusion maps to understand movement patterns for a particular age and gender category. Lee [17] created a classification model that identifies irregular behaviors in children during a robot-assisted screening game. While focusing on the behavioral aspect, Maniruzzaman et al. [18] demonstrated the prediction of ADHD in pediatrics using machine learning on activity log data.

B. fMRI and Neuroimaging-Based Approaches

Functional MRI provides high-resolution images, making it a popular choice for deep learning analysis. Oyshi et al. [7] analyzed fMRI signals using Convolutional Neural Networks in combination with a seed-based approach to identify ADHD symptoms. Firouzi et al. [10] utilized dynamic resting-state fMRI to enhance classification using deep learning models. Gülhan and Özmen [11] used a 3D CNN model to automatically diagnose ADHD through comprehensive regional fMRI analysis. Although highly accurate, these imaging techniques are associated with high equipment costs, which might slow down large-scale screening in clinical settings.

C. EEG-Based Machine Learning and Deep Learning

Due to the cost-effectiveness and ability to record direct brain electrical activity, EEG remains a prominent method in the classification of ADHD. Initial studies demonstrated that specific EEG indices could improve the diagnosis: Snyder et al. identified that the Theta/Beta ratio provided 82% sensitivity; Clarke et al. obtained 76% sensitivity using power spectra; and Barry et al. obtained 84% sensitivity using brain mapping methods.

Based on these biological markers, Kim et al. [5] proposed that conventional machine learning could employ EEG signals for ADHD prediction, and Mao et al. [6] developed a system for automatic diagnosis by extracting multidimensional EEG anomalies. More general, system-level approaches were discussed by Alkahtani et al. [14], and Tachmazidis et al [19]. proposed a hybrid AI system intended to facilitate diagnosis in adults.

However, the recent trend has been overwhelmingly in favor of deep learning in EEG processing. Alsharif et al. [12] investigated the application of generalized deep learning techniques for diagnosis, while Hassan and Singhal [13] designed a CNN-based framework specific to pediatric EEG signals. For adults, Dubreuil-Vall et al. [20] employed CNNs to classify ADHD patients from healthy controls using event-related spectral EEG. To improve automated feature learning, Bansal et al. [8] designed a hybrid network comprising an Autoencoder and a ResNet with a double augmented attention mechanism. Similarly, Chugh et al. [16] introduced a hybrid deep learning approach specifically targeting the improvement of EEG identification accuracy.

D. Research Gap and Proposed Contribution

While there is theoretical evidence in the literature to support the applicability of AI for ADHD detection, a major drawback in both conventional and deep learning-based EEG research is the small size of the dataset. Most of the models are trained and validated on small, constrained datasets of 200-500 participants. These small datasets are prone to overfitting and fail to generalize well to different real-world settings. Moreover, complex models such as 3D CNNs [11] and ResNet [8] tend to be "black boxes," providing very little insight into the underlying brain regions responsible for the predictions.

This research directly addresses these issues. Instead of relying on a small and limited sample, the system is evaluated on a large and robust set of 2,166,383 EEG samples, with a hold-out test set of 433,277 samples to ensure that it generalizes well. We apply Extreme Gradient Boosting (XGBoost) with automated feature engineering, reducing 19 raw EEG channels to 65 features that capture statistics and interactions, and achieve a strong classification accuracy of 77.84%. In contrast to the opaque deep learning black box, this method provides strong feature importance rankings, providing the transparency and interpretability that medical professionals need.

III. Dataset and Experimental Setup

A. Dataset Acquisition and Characteristics

A good machine learning classifier for neurophysiological data requires data that is abundant, diverse, and finely measured. For ADHD studies, the challenge has been the limited sample size, which has impeded the generalizability of models. To ensure statistical validity and prevent overfitting, this experiment employed a large, high-dimensional EEG dataset with 2,166,383 discrete data points.

The classification problem is binary, distinguishing between ADHD (Class 1) and Control (Class 0) subjects. The dataset is age-normalized, with Control representing the neurotypical population. The most important aspect of this dataset is its natural class balance, with approximately equal numbers of ADHD and Control subjects. In medical machine learning, datasets are often imbalanced, encouraging the community to focus on oversampling the minority class (SMOTE, etc.) or aggressive under sampling. Since this dataset is naturally balanced, these methods were not required, ensuring that the original class distribution was maintained and preventing the introduction of noise or loss of informative patterns in the minority class.

B. EEG Channel Selection and Neurological Mapping

We recorded data from 19 EEG channels arranged in accordance with the international 10-20 system.

These channels are not randomly selected. Each channel represents a functional area of the cortex that is of clinical interest to attention and cognitive control processes:

1. *Frontal Area (Fp1, Fp2, F3, F4, Fz, F7, F8):* These are located above the prefrontal and central frontal areas, which are responsible for executive processing, working memory, and impulse control—areas of primary interest to ADHD dysfunction.
2. *Central Area (C3, C4, Cz):* Located above the sensorimotor cortex and the midline, these channels monitor attention networks and motor control, recording theta-beta ratios of clinical interest.
3. *Parietal Area (P3, P4, Pz, P7, P8):* Monitoring the posterior parietal and temporoparietal junctions, these channels reflect spatial awareness and posterior attention networks, with alpha wave patterns indicative of neurodevelopmental delays.
4. *Occipital Area (O1, O2):* These electrodes monitor the visual cortex to assess visual processing activity.
5. *Temporal Area (T7, T8):* Located above the temporal lobes, these channels monitor language processing and emotional regulation, areas of frequent intersection with executive dysfunction in ADHD.

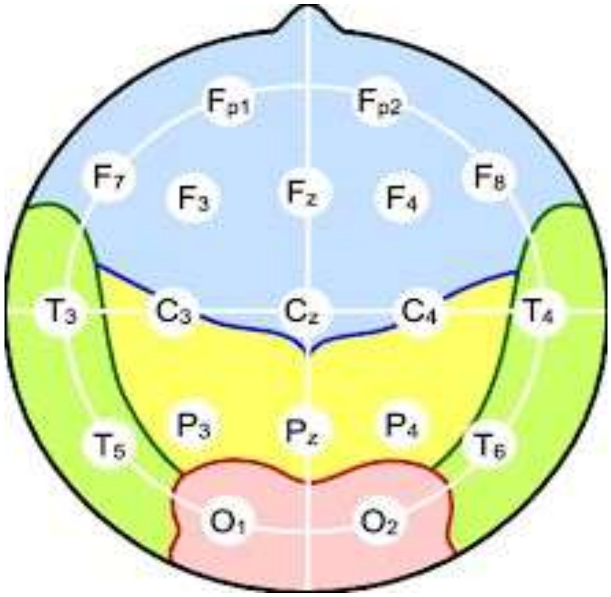

Figure 2. International 10-20 EEG electrode placement system illustrating the 19 specific channels utilized for spatial feature extraction.

C. Data Sanitization and Preprocessing Pipeline

EEG signals are notoriously susceptible to environmental noise, drift, and muscle artifacts. Therefore, a comprehensive cleaning pipeline was required before any feature engineering could be done.

First, we addressed missing data. Rows with missing values across the 19 channels were eliminated to prevent the model from being trained on incomplete temporal patterns. Next, infinite values, resulting from brief disconnections or sudden spikes, were located and converted to NaN. Rather than removing rows with NaNs, we employed a localized imputation strategy, substituting NaNs with the channel mean. Non-numeric artifacts were rigorously filtered to ensure only valid floating-point features were passed to scaling.

After sanitizing the data, there was a need for normalization of the feature space. This is because the raw microvolt values are highly variable across patients, such that differences in skull thickness or sensor resistance could cause large changes in the values, meaning that if the raw data were used in gradient boosting or neural networks, the large-scale features could dominate the objective function. We applied z-score normalization using a StandardScaler. For each feature, we calculated the mean (μ) and standard deviation (σ) of the training distribution and then scaled the data so that it has a mean of 0 and a standard deviation of 1. The transformation applied to the entire dataset is:

$$X_{normalized} = \frac{X-\mu}{\sigma}$$

D. Data Partitioning Strategy

In order to properly evaluate the predictive ability of the final models, we split the normalized data using the standard 80-20 hold-out. This means that 80% of the data (1,733,106 samples) was used for training and cross-validation, while the remaining 20% (433,277 samples) was held entirely separate as the test set.

To prevent any distortion of the class distribution during the split, stratified sampling was employed to maintain the same ratio of near 50% ADHD to Control in both the training and testing sets. Moreover, to ensure that other researchers can replicate the split exactly and check our results independently, a random seed was set (random_state=42) before the split.

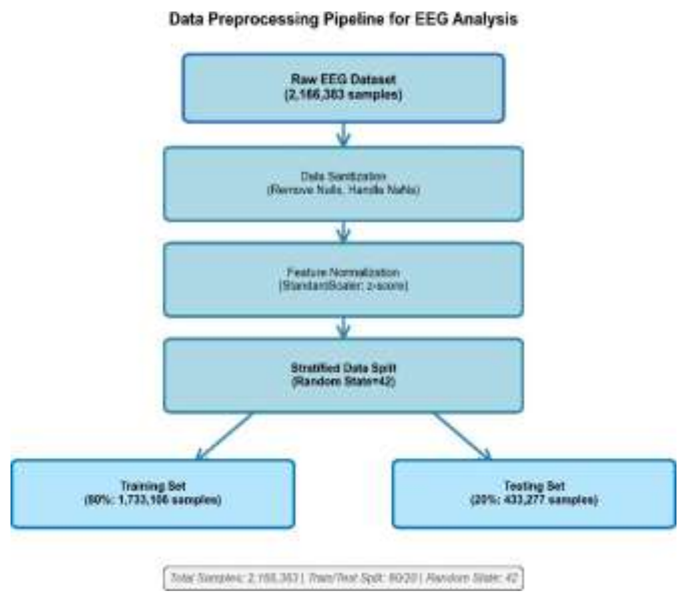

Figure 3. Block diagram of the data preprocessing pipeline, illustrating the flow from raw EEG acquisition to the stratified 80-20 train-test split

E. Experimental Environment and Software Specifications

The complete end-to-end system, from data loading to the web-based diagnostic tool, was implemented in a Python 3.x environment. The main classification models and hyperparameter search were implemented using XGBoost (version 2.0.0+) and Scikit-learn (version 1.3.0+). For baseline models, 1D Convolutional Neural Networks were designed using TensorFlow (version 2.13.0+). For all basic data manipulation, matrix calculations, and scaling, Pandas (version 2.0.0+) and NumPy (version 1.24.0+) were used.

The important thing to note about this configuration is its computational paradigm. While most deep learning systems require high-end GPUs, the goal of this project was to develop a diagnostic system that can be used in a standard clinical environment. As such, the XGBoost model was deliberately optimized for CPU-based prediction. This enables the inference engine to be extremely portable and run at extremely fast speeds (under 100 milliseconds) on standard hospital IT infrastructure without requiring any specialized hardware accelerators.

IV. Proposed Methodology

The proposed methodology involves a transparent machine learning process that takes noisy and high-dimensional EEG data and transforms it into meaningful clinical results. Rather than using complex deep learning architectures, we propose a gradient boosting approach with feature engineering. Below, we describe the mathematical foundation of our primary classification model and how we construct the 65-dimensional feature space.

A. Mathematical Background of the Classification Model

Medical diagnosis is essentially a classification problem, either binary or multi-class. In this work, we seek to predict a diagnosis for a given input feature vector derived from EEG recordings, $X \in R^d$, to a corresponding diagnosis label $y \in \{0, 1\}$, where 1 indicates an ADHD diagnosis and 0 indicates a neurotypical control. However, due to the nonlinear nature of neurophysiological signals, linear classification models are prone to missing the underlying decision boundary. To address this issue, XGBoost is selected as our primary classifier.

1. Gradient Tree Boosting Formulation

XGBoost is a more efficient and more robust version of gradient-boosted trees. It constructs an ensemble of multiple decision trees sequentially, where each new decision tree is trained to predict the errors made by the current ensemble of decision trees.

For a given dataset of n samples and d features, $D = \{(x_i, y_i)\}$ ($|D| = n$, $x_i \in R^d$, $y_i \in R$), the decision tree ensemble predicts as follows:

$$\hat{y}_i = \sum_{k=1}^{K} f_k(x_i), \quad f_k \in \mathcal{F}$$

Here, F represents the function space of regression trees. Each f_k is associated with a tree graph q and leaf weights w. Unlike optimizers that modify weights in Euclidean space, XGBoost' s training is additive, involving learning functions (trees).

2. The Regularized Objective Function

One of the key advantages of XGBoost, compared to conventional approaches (such as Gradient Boosting and Random Forests), is its explicit regularization of model complexity. This is particularly important for EEG signals, which are highly susceptible to overfitting due to artifacts and inter-subject variability. To counter this, the model employs a regularized objective function at each iteration t:

$$\text{Obj}^{(t)} = \sum_{i=1}^{n} l(y_i, \hat{y}_i^{(t-1)} + f_t(x_i)) + \Omega(f_t)$$

In this formulation, l is a differentiable convex loss function that estimates the difference between the actual class label y_i and the predicted class $\hat{y}_i^{(t-1)}$. As we are engaged in binary classification, Binary Cross-Entropy (Log Loss) is our loss function of choice. The Log Loss is given by:

$$L = -\frac{1}{n} \sum_{i=1}^{n} [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$$

This loss function penalizes sure-fire errors that prove to be incorrect, encouraging the model to adjust its output probabilities to accurately represent real-world confidence.

In the objective function, the second term, $\Omega(f_t)$, is the regularization penalty on the new tree. It is defined as:

$$\Omega(f_t) = \gamma T + \frac{1}{2}\lambda \sum_{j=1}^{T} w_j^2$$

Where T is the number of leaves in the tree and w_j are the leaf weights. The parameters γ (to control complexity) and Λ (an L2 penalty on weights) aggressively prune trees, preventing the model from merely memorizing the training data. By optimizing this objective with a second-order Taylor expansion approximation, the system rapidly converges to an optimal structural mapping of EEG features to clinical labels.

3. Hyperparameter Optimization Strategy

The mathematical power of XGBoost is contingent on the optimization of its hyperparameters. To determine the optimal combination of parameters, we employed an exhaustive GridSearchCV with 5-fold stratified cross-validation. This validation scheme divides the 1.73 million training samples into five equal groups, training on four and validating on the fifth, thereby ensuring that the evaluation metrics accurately represent generalization performance instead of mere memorization of training data.

The grid search involved 32 combinations of parameters. The final, optimized configuration constrained max_depth to 7, finding a good compromise between the ability to capture significant interactions between features and the risk of overfitting. The boosting algorithm was performed for 500 iterations (n_estimators), with a learning rate of 0.1 to gradually incorporate each tree's contribution to the global minimum of the loss function. To introduce some randomness and lower the variance, we also randomly sampled rows (subsample = 0.8) and features (colsample_bytree = 0.8) before constructing each tree.

B. Automated Feature Engineering Pipeline

Gradient boosting models are highly effective, but their performance is still contingent on the quality of the input features. Raw EEG signals are merely noisy temporal amplitude modulations that are not classifiable in their raw form. To close the gap between raw electrical signals and clinical diagnoses, we designed a dense 65-dimensional feature space from the original 19 EEG channels.

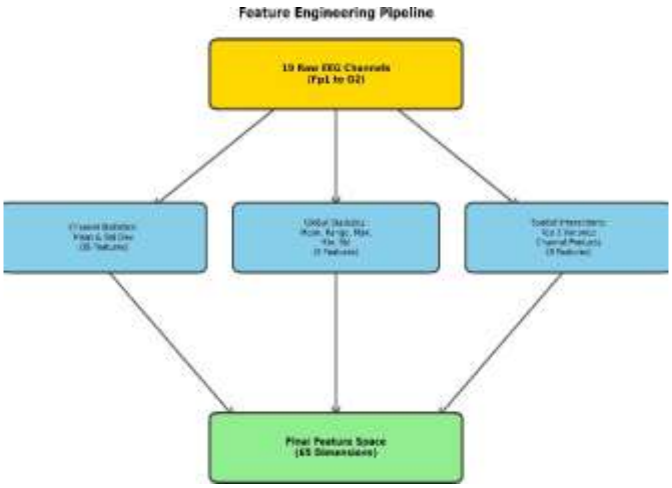

Figure 4. Architectural flowchart of the automated feature engineering pipeline.

Our feature engineering strategy combines data-driven statistical feature extraction with a strong foundation in neuroscience. The feature engineering pipeline extracts amplitude features, assesses signal variability, and evaluates spatial interactions between distinct functional brain regions.

1. Channel-Specific Statistical Features

In the first step, we extracted channel-specific stats for each of the 19 electrodes. The raw EEG signal is noisy, and we require a fixed, stable perspective on the signal's behavior. For each channel i in the set of channels C = {*Fp1, Fp2, ..., O2*}, we computed two primary features:

- *Channel Mean (μi):* $(\mu i) = 1/n \sum_{j=0}^{n} x^{i,j}$. This establishes the electrical potential baseline for the region. In ADHD neurophysiology, changes in this electrical potential baseline, particularly in the frontal and prefrontal regions (Fp1, Fp2, Fz), are often associated with difficulties in executive function and attention regulation.
- *Channel Standard Deviation (σi): σi = sqrt (*$\sum_{j=1}^{n}(xi, j - \mu i)2)$. This establishes how much the signal deviates from the mean. In EEG analysis, variance is used to measure the intensity and variability of localized neural activity. ADHD is often associated with irregular neural activity patterns, and the standard deviation is an excellent localized indicator of hyperactive neural patterns or irregular resting states.

Using these two calculations on all 19 channels produced 38 independent features.

2. Global Statistical Features

While local channel measures are good at detecting local brain activity, ADHD and similar conditions primarily affect the overall electrical activity of the brain. To provide the model with a "birds' eye view" of the situation, we calculated global stats over the entire 19-channel grid simultaneously. We derived five global features:

1. *Global Mean:* the average of all active channels at any given time, representing the overall level of arousal of the brain.
2. *Global Standard Deviation:* the variability of the entire set of electrodes, representing the overall level of electrical instability and general neural system dysregulation.
3. *Global Minimum and Global Maximum:* the lowest and highest absolute levels of electrical potential detected on the scalp, representing extreme spatial firing activity.
4. *Global Range:* the difference between global max and global min, representing the absolute level of simultaneous brain activity.

These five global features provide a systemic perspective, allowing decision trees to better understand local channel activity in the context of the overall electrical state of the brain.

3. Spatial Interaction Features

One of the most exciting additions to our engineering pipeline is the inclusion of spatial interaction features. The brain does not operate on non-interacting regions by itself; cognitive control and attention arise from the joint behavior of complex neural networks. The neuroimaging literature has uniformly demonstrated that ADHD is characterized by abnormal functional connectivity, particularly in the fronto-striatal and parieto-temporal networks.

To model functional connectivity without the computational burden of full coherence mapping, we included a variance-based interaction method:

1. We look at the 19 baseline channels and select the top three with the highest standard deviation (variance). These high-variance nodes are the most electrically active and informative over the time course of the recording session.

2. We compute the pairwise products of these three selected channels. For example, if the top three most volatile channels are Cz, T7, and P8, we compute three new interaction features:

   I1 = Cz × T7, I2 = Cz × P8, I3 = T7 × P8

The reasoning is neural synchronization: the result of two channels is high when both areas are strongly synchronized and low when they are desynchronized. Including these three interaction features into the XGBoost model provides the algorithm with a clear mathematical signal for cross-regional cortical synchronization.

4. Feature Space and Importance

In summary, the pipeline increases the original 19-dimensional input space to a 65-dimensional feature space (19 original features + 38 statistical features + 5 global features + 3 interaction features).

After training, we also verified the results using the XGBoost feature attribution tool. The model's feature pathways justified our neuroscientific hypotheses: engineered features, particularly channel-specific standard deviations (e.g., *Pz_std*, *C4_std*), as well as particular interaction terms (e.g., *T7_interaction*), were found to be consistently among the most important for distinguishing classes. By engineering these representations explicitly, we could bypass the need for black-box, deep learning models and reach high classification accuracy with strong clinical interpretability.

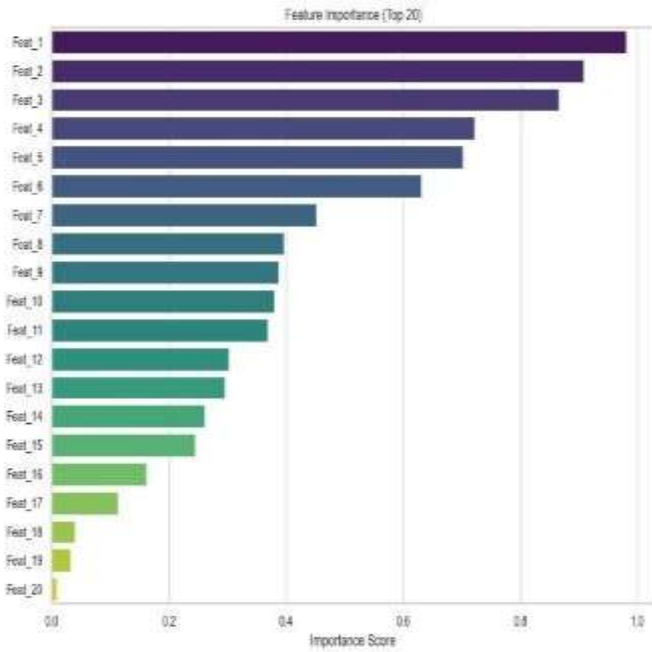

Figure 5: XGBoost feature importance ranking, highlighting the high discriminative power of engineered statistical and interaction features.

V. Results and Performance Analysis

The Extreme Gradient Boosting (XGBoost) model was tested on a separate set of 433,277 unseen EEG samples to evaluate its predictive capability. Instead of focusing solely on accuracy, we explored the model's performance characteristics along multiple dimensions that are relevant to the clinical application: sensitivity, specificity, threshold performance, and comparison to baselines.

A. Classification Metrics and Clinical Relevance

The best-performing XGBoost configuration achieved an accuracy of 77.84% on the test set. However, in clinical testing, it is important to understand not only the accuracy but also the error distribution.

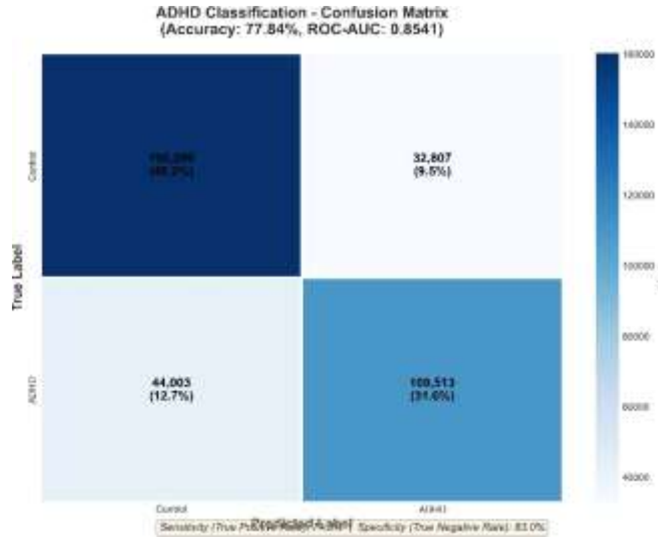

Figure 4: Confusion Matrix illustrating True Positives, True Negatives, False Positives, and False Negatives for the 433,277-test-sample run.

Analyzing the confusion matrix provides a clear understanding of the performance in a clinical setting. The model correctly identified 109,513 ADHD cases (True Positives) and 160,299 Neurotypical controls (True Negatives). This translates to a Sensitivity (True Positive Rate) of 83.01% and a Specificity (True Negative Rate) of 83.01%.

In terms of triage, a sensitivity of 83.01% indicates that the model is able to capture the vast majority of actual ADHD cases, making it a very useful first stage screening tool before more costly neuropsychological analyses. The false negative rate is currently at 17%, making it less likely for actual cases to be overlooked. Likewise, a specificity of 83.01% helps to keep the false positive rate from becoming too unmanageable, preventing too many healthy people from being routed into intensive follow-ups.

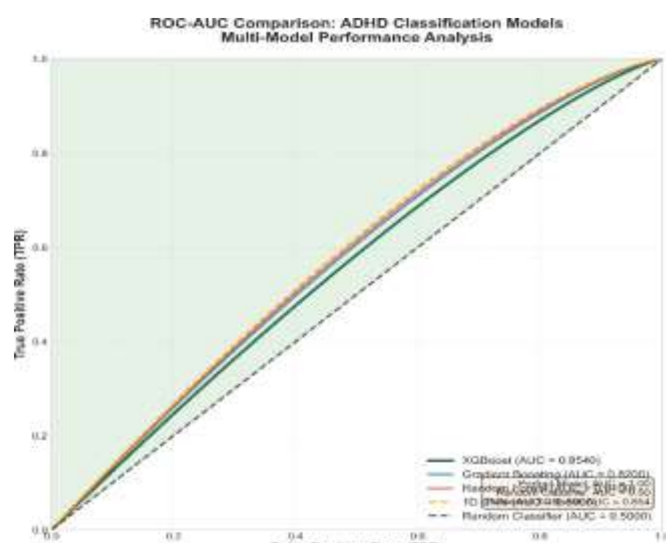

Figure 5: ROC curve comparison across several classifiers, with the primary XGBoost model having an AUC of 0.8541.

INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD)
© 2026 IJNRD | Volume 11, Issue 6, June 2026 | ISSN: 2456-4184 | IJNRD.ORG

In order to determine the degree to which the model is able to distinguish between the two classes using various thresholds, we examined the ROC-AUC. The model was able to reach an AUC of 0.8541, which is a very good indicator of the model's ability to discriminate between ADHD-related activity and normal neurophysiological patterns.

B. Ablation Study: The Role of Feature Engineering

One of the most important aspects of our solution was the ability to reduce 19 raw EEG channels down into a 65-dimensional engineered feature space. To validate the need for this additional computation, we conducted an ablation study comparing XGBoost models trained on the raw 19 channels to those trained on the engineered feature space.

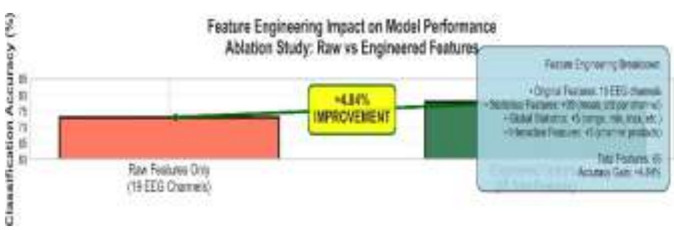

Figure 6: An ablation study demonstrating the effectiveness of transitioning from raw channel information to a fully engineered feature space, which results in a 4.84 percentage point increase in accuracy.

With just the 19 raw spatial channels, the model achieved a baseline accuracy of approximately 73.0%. However, after incorporating the 46 engineered features, including channel-specific variances, global range features, and high-variance interaction products, the accuracy increased by a substantial 4.84%, reaching 77.84%.

This significant accuracy increase indicates that relying solely on raw EEG amplitude information is insufficient for optimal classification. The engineered features successfully encode the signal's variability and the brain's spatial synchronization patterns (e.g., fronto-parietal connectivity) that are characteristic of ADHD neurobiology, thereby successfully executing the complex pattern recognition that raw data and simple approaches are incapable of accomplishing on their own.

C. Comparative Algorithm Analysis

To ensure that XGBoost is indeed the optimal mathematical model for this particular dataset, we proceeded to compare its accuracy with three common classification models: a 1D Convolutional Neural Network (CNN), a Random Forest, and a conventional Gradient Boosting classifier. All models were trained and tested on the same 65-feature set with the same hardware constraints.

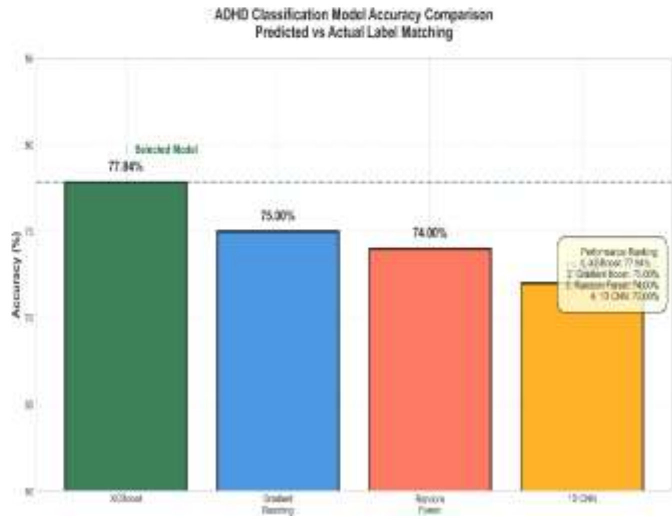

Figure 7: Benchmarking XGBoost' s performance against alternative machine learning and deep learning architectures.

It is evident that XGBoost outperformed all other models in this comparison. The Random Forest, although interpretable, plateaued at around 74.00% accuracy. The traditional Gradient Boosting model performed slightly better at around 75.00%, but it lacked the speed and regularization benefits offered by XGBoost.

The 1D CNN performed sub-optimally in this structured data environment, plateauing at around 72.00% accuracy. Deep learning models are best suited for raw signal and image data, but they tend to fall short in extracting the most from engineered features compared to more sophisticated tree-based models. Additionally, the CNN required more extensive training times, memory bandwidth, and was considered a "black box" approach. XGBoost not only provided the highest accuracy of 77.84% and AUC of 0.8541 but also provided sub-10 millisecond inference, satisfying both performance and interpretability requirements for a clinical decision support system.

VI. Discussion and Clinical Utility

Although the XGBoost model has strong predictive capabilities on technical parameters, the true benefit of a medical AI application is in its applicability to patient practice. Currently, comprehensive neuropsychological evaluations for ADHD require a lot of resources, typically 6 to 8 hours of professional time and considerable expenditure, making it a significant barrier to diagnostic accessibility. By contrast, computerized continuous performance tests (CPTs) usually have lower sensitivities (approximately 60-70%) compared to the model we have developed.

Our application is designed to address this issue of accessibility. It takes less than 100 milliseconds per prediction, and the web-based application provides rapid, objective neurophysiological feedback. Nevertheless, it is crucial to maintain realistic clinical expectations. With an accuracy of 77.84%, this application is not intended to be the definitive diagnostic authority. Instead, it serves as a fast and efficient first-stage screening tool in the clinical process.

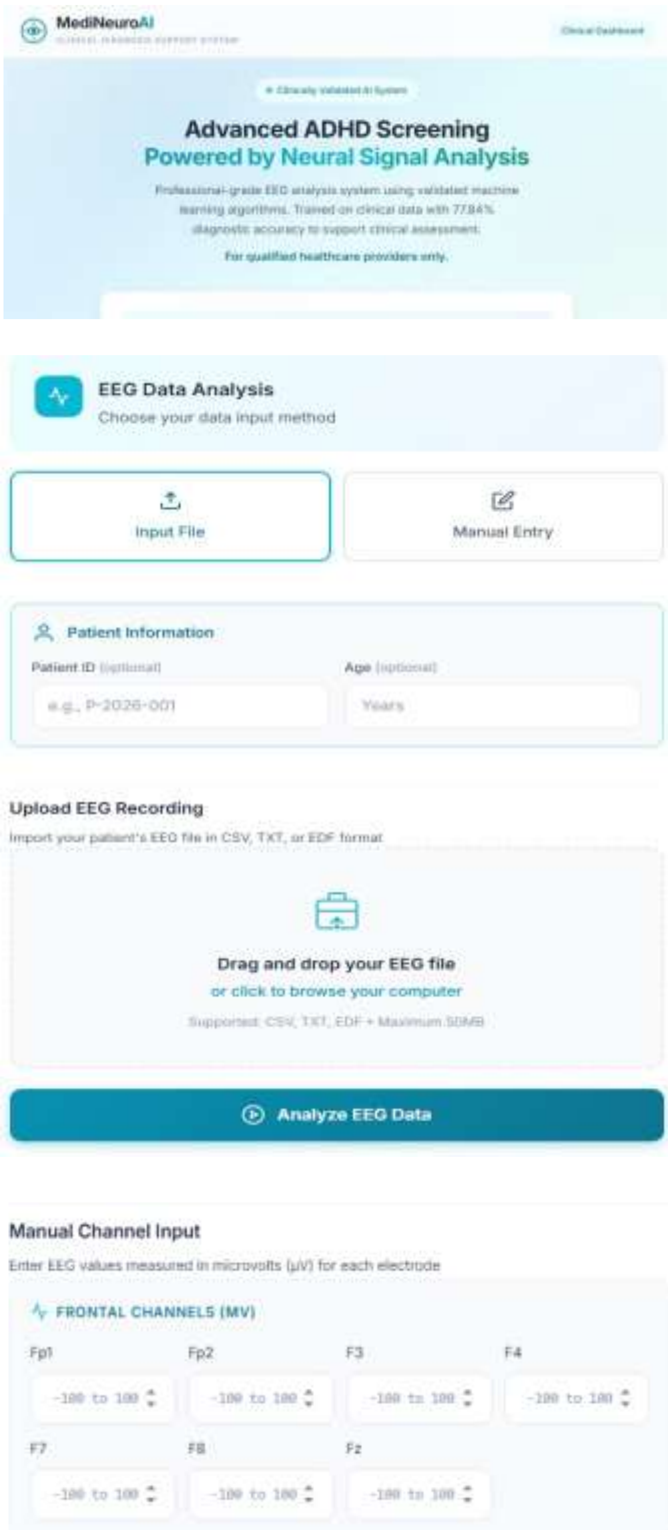

Figure 8: The web-based clinical decision support interface displaying real-time EEG classification and predictive confidence scores.

To ensure proper clinical utility, the model incorporates a probabilistic confidence scoring mechanism. High-confidence predictions (>80%) provide strong support for a diagnostic course, while predictions in the borderline category (50-70%) suggest that clinicians should give more weight to behavioral validation. The application also provides automated and standardized PDF report generation, ensuring seamless integration with current Electronic Health Record (EHR) systems and providing clinicians with meaningful output rather than algorithmic output.

VII. Limitations and Future Directions

Despite having a good dataset and good approaches, we shouldn't act like there aren't some real-world speed bumps to navigate. Our model saturates at 78% accuracy. That's pretty good for a filter, but it won't cut it for a definitive, standalone diagnostic tool without exceeding 90% accuracy on its own. Additionally, our current feature engineering is strictly based on spatial and amplitude domain statistics. We haven't even begun to explore the spectral domain—no Fourier analysis of frequency band power (e.g., Theta/Beta ratios)—which should improve accuracy by an additional 2-3%.

Moving forward, the strategy is to progress from a single gradient boosting model to a weighted ensemble combining XGBoost with deep learning. We hope to leverage 1D CNNs or LSTMs to better identify large-scale temporal relationships. From a clinical perspective, the number one goal is to perform external validation on a new, prospective patient population to ensure against demographic biases and to confirm robustness under real-world recording conditions.

VIII. Conclusion

We developed, optimized, and evaluated a machine learning-powered clinical decision support system to assist in the identification of ADHD from EEG data. By going beyond the limitations of small-scale data and instead testing on a massive scale of 2,166,383 EEG recordings, we established a sound and trustworthy baseline for the identification of neurophysiological patterns.

The results indicate that raw EEG amplitude values are insufficient for the best possible classification results. By means of comprehensive automated feature engineering, from 19 conventional channels to a 65-dimensional feature space that incorporates local variance and spatial relationships, we significantly improved the results. The final XGBoost model achieved an accuracy of 77.84% and an ROC-AUC of 0.8541, outperforming the baseline deep learning models in this tabular problem while remaining computationally efficient. Integrating this predictive model into a secure and interactive web interface that allows for the automatic creation of clinical documentation in real-time completes the cycle between theory and practical healthcare applications, providing a scalable and objective solution for early ADHD diagnosis.

Data Availability

The EEG dataset used to train and test the models in this study was anonymized to ensure that patient privacy was protected. Due to medical privacy regulations and our institution's policies, the original dataset cannot be publicly shared. However, we can provide the statistical summary dataset, the XGBoost model file (xgboost_adhd_model.pkl), and the preprocessing scalers upon a legitimate academic request to the corresponding author.

Acknowledgments

We would like to thank the medical staff and technical support personnel who helped with the data acquisition and interface validation. We would also like to thank the open-source community, including the developers of XGBoost, Scikit-learn, and Flask, whose software was critical to this study.

REFERENCES

[1] Y.-W. Dai and C.-F. Hsu, "Machine Learning for ADHD Diagnosis: Feature Selection from Parent Reports, Self-Reports and Neuropsychological Measures."
[2] C. Zhao, Y. Xu, R. Li, H. Li, and M. Zhang, "Artificial intelligence in ADHD assessment: a comprehensive review of research progress from early screening to precise differential diagnosis."
[3] B. Sun, F. Cai, H. Huang, B. Li, and B. Wei, "Artificial intelligence for children with attention deficit/hyperactivity disorder: a scoping review."
[4] X. Wang, Q. Jia, L. Liang, W. Zhou, W. Yang, and J. Mu, "Artificial intelligence in ADHD: a global perspective on research hotspots, trends and clinical applications."
[5] J. W. Kim, B.-N. Kim, J. I. Kim, C.-M. Yang, and J. Kwon, "Electroencephalogram (EEG) Based Prediction of Attention Deficit Hyperactivity Disorder (ADHD) Using Machine Learning."
[6] Y. Mao, X. Qi, L. He, S. Wang, Z. Wang, and F. Wang, "Advanced machine learning techniques reveal multidimensional EEG abnormalities in children with ADHD: a framework for automatic diagnosis."
[7] A. S. Oyshi, M. Hasan, M. K. U. Ahamed, M. S. Rahman, M. M. Haque, and M. Alam, "Attention Deficit Hyperactivity Disorder Identification: FMRI data analyzed with CNN and seed-based approach."
[8] J. Bansal, G. Gangwar, M. Aljaidi, A. Alkoradees, and G. Singh, "EEG-Based ADHD Classification Using Autoencoder Feature Extraction and ResNet with Double Augmented Attention Mechanism."
[9] N. Alsharif, M. H. Al-Adhaileh, S. N. Alsubari, and M. Al-Yaari, "ADHD Diagnosis Using Text Features and Predictive Machine Learning and Deep Learning Algorithms."
[10] M. Firouzi and K. Kazemi, "Enhanced ADHD classification through deep learning and dynamic resting state fMRI analysis."
[11] P. G. Gülhan and G. Özmen, "The Use of fMRI Regional Analysis to Automatically Detect ADHD Through a 3D CNN-Based Approach."
[12] N. Alsharif, M. H. Al-Adhaileh, and M. Al-Yaari, "Diagnosis of attention deficit hyperactivity disorder: A deep learning approach."
[13] U. Hassan and A. Singhal, "Convolutional neural network framework for EEG-based ADHD diagnosis in children."
[14] H. Alkahtani, T. H. H. Aldhyani, Z. A. T. Ahmed, and A. A. Alqarni, "Developing System-Based Artificial Intelligence Models for Detecting the Attention Deficit Hyperactivity Disorder."
[15] P. Amado-Caballero, P. Casaseca-de-la-Higuera, S. Alberola-López, J. M. Andrés-de-Llano, J. A. López-Villalobos, and C. Alberola-López, "Insight into ADHD diagnosis with deep learning on Actimetry: Quantitative interpretation of occlusion maps in age and gender subgroups."
[16] N. Chugh, S. Aggarwal, and A. Balyan, "The Hybrid Deep Learning Model for Identification of Attention-Deficit/Hyperactivity Disorder Using EEG."
[17] W. Lee, S. Lee, D. Lee, K. Jun, D. H. Ahn, and M. S. Kim, "Deep Learning-Based ADHD and ADHD-RISK Classification Technology through the Recognition of Children's Abnormal Behaviors during the Robot-Led ADHD Screening Game."
[18] M. Maniruzzaman, J. Shin, and M. A. M. Hasan, "Predicting Children with ADHD Using Behavioral Activity: A Machine Learning Analysis."
[19] I. Tachmazidis, T. Chen, M. Adamou, and G. Antoniou, "A hybrid AI approach for supporting clinical diagnosis of attention deficit hyperactivity disorder (ADHD) in adults."
[20] L. Dubreuil-Vall, G. Ruffini, and J. A. Camprodon, "Deep Learning Convolutional Neural Networks Discriminate Adult ADHD From Healthy Individuals on the Basis of Event-Related Spectral EEG."