EFFICIENT SLEEP APNEA DETECTION USING XGBOOST AND BIOSIGNAL FEATURE ENGINEERING

Mazna Khan[1], Dr. I. Samuel Peter James[2]
[1]PG Scholar, Department of AI&DS, Shadan Women's College of Engineering and Technology, Hyderabad,
Maznakhan124@gmail.com
[2]Assoc. Professor, Department of AI&DS, Shadan Women's College of Engineering and Technology
drisamuelpeterjames@gmail.com

ABSTRACT

Frequent breathing pauses during sleep are the hallmark of sleep apnoea, a common and potentially fatal sleep disease. Serious health issues like weariness, decreased cognitive function, and cardiovascular disorders might result from these periods. Although current deep learning models have demonstrated good performance in apnoea diagnosis, they are less appropriate for real-time or edge-based healthcare applications due to their high computational resource requirements and lack of transparency. In this work, we provide a machine learning-based method that uses Extreme Gradient Boosting (XGBoost) to effectively and precisely identify sleep apnoea episodes from physiological signals.A strong feature engineering pipeline is used in the suggested system to extract statistical, temporal, and frequency-domain characteristics from biosignals include oxygen saturation (SpO_2), airflow, and ECG data. To reduce redundancy and improve classification performance, these characteristics are chosen and improved utilising correlation analysis and feature importance techniques. In order to address class imbalance and enhance generalisation, the collected features are subsequently fed into an XGBoost classifier, which is optimised using cross-validation and hyperparameter tuning. In comparison to deep learning models, our approach shows competitive accuracy while drastically lowering computational cost when tested on the Sleep Heart Health Study (SHHS) dataset. Additionally, the model uses feature importance analysis to produce interpretable outputs that help clinical experts comprehend choice variables. Because of this, the method is ideal for explainable AI-driven diagnoses, portable device deployment, and real-time monitoring in the context of sleep apnoea and related diagnosis.

1. INTRODUCTION

Reduced oxygen saturation and disrupted sleep patterns are the results of frequent breathing disruptions during sleep, which are a prevalent symptom of sleep apnoea. Millions of people worldwide are impacted by it, and it is linked to serious health issues like heart disease, high blood pressure, stroke, cognitive decline, and excessive daytime tiredness. Timely diagnosis and successful treatment of sleep apnoea depend on early and accurate detection.

Polysomnography (PSG), a thorough sleep study that collects many physiological signals such blood oxygen saturation (SpO_2), airflow, and electrocardiogram (ECG), is the basis for the conventional diagnosis of sleep apnoea. Even while PSG offers accurate diagnostic data, it is costly, time-consuming, and necessitates specialised clinical facilities and professional analysis. These constraints have spurred the creation of machine learning and artificial intelligence-based automated sleep apnoea detection devices.

Recent research has shown that deep learning models, such as Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), are useful for detecting sleep apnoea. Despite achieving excellent classification accuracy, these methods frequently call for sizable annotated datasets, substantial processing power, and intricate training processes. Additionally, interpretability—a crucial criterion in healthcare applications—is limited by their black-box character.

In addition, the use of engineered physiological features enables efficient representation of sleep-related patterns while reducing data dimensionality. The proposed framework supports explainable artificial intelligence by identifying the most influential features contributing to apnea detection. Such transparency enhances clinician confidence and facilitates informed medical decision-making. Furthermore, the lightweight nature of the model makes it suitable for deployment in resource-constrained healthcare and remote monitoring environments.

To address these challenges, this work proposes a machine learning-based framework utilizing Extreme Gradient Boosting (XGBoost) for sleep apnea detection from physiological signals. The proposed approach focuses on extracting informative statistical, temporal, and frequency-domain features from biosignals such as ECG, airflow, and SpO_2. Feature selection techniques are employed to remove redundant information and enhance classification performance. The selected features are then used to train an XGBoost classifier capable of accurately distinguishing apnea and non-apnea events.

The proposed method offers several advantages, including lower computational complexity, faster training and inference, robustness to imbalanced datasets, and improved interpretability through feature analysis of importance. These features make the system appropriate for clinical decision-support systems, wearable medical equipment, and real-time monitoring. The suggested XGBoost-based system delivers competitive performance while preserving efficiency and transparency, according to experimental assessment on sleep apnoea datasets, making it a workable approach for automated sleep apnoea diagnosis.

2. EXISTING SYSTEM

Drawbacks: -

- High Computational Cost
- Limited Interpretability
- Training Complexity
- Dependency on Large Labeled Datasets

The current technology uses multimodal physiological markers and a deep learning-based method to identify sleep apnoea episodes. These signals, which usually include oxygen saturation (SpO_2), airflow, and ECG, are rich in information yet vary greatly between people. The system combines Gated Recurrent Units (GRUs) for modelling temporal relationships in the data with Convolutional Neural Networks (CNNs) for spatial feature extraction in order to extract meaningful patterns from these signals. To enable the model to concentrate on pertinent time intervals in the input sequence, an attention mechanism is incorporated.

A Multi-Domain Feature Extraction (MDFE) block and a Squeeze-and-Excitation (SE) module for channel-wise feature refinement are included to significantly improve the system's capacity to learn a variety of features. To promote variety among feature extractors and enhance overall robustness, a customised loss function is employed. The Sleep Heart Health Study (SHHS) dataset was used to train and test an end-to-end deep learning pipeline, which produced encouraging accuracy and area under the ROC curve results. (AUC). The system's excellent performance makes it appropriate for automated apnoea detection, but it also has limits in terms of interpretability and computing requirements.

ALGORITHM (CNN+GRU):

The core algorithm of the existing system is a hybrid deep learning architecture that fuses the strengths of CNNs, GRUs, and attention mechanisms. CNN layers are responsible for learning spatial features from raw physiological signal segments, such as waveform shapes or local fluctuations in breathing or oxygen levels. These CNN-extracted features are then passed to a GRU network, a type of recurrent neural network that captures temporal patterns and dependencies across signal sequences. GRUs are preferred over traditional LSTMs due to their reduced computational complexity while still retaining strong sequence modeling capability. An attention mechanism is employed on top of the GRU outputs to dynamically weight the importance of different time steps, allowing the model to focus on critical apnea-related intervals. Furthermore, a multi-domain feature extractor (MDFE) and squeeze-and-excitation (SE) block help enrich the feature representation by integrating domain knowledge and performing channel attention. The model is trained with a custom loss function that encourages diversity among feature extraction components to avoid redundancy and overfitting. While this combination is powerful, it is also resource-intensive and less interpretable.

3. PROPOSED SYSTEM

The suggested method presents an Extreme Gradient Boosting (XGBoost)-based machine learning framework for the effective and comprehensible identification of sleep apnoea events from physiological inputs. This method concentrates on feature extraction and selection from biosignals like ECG, airflow, and SpO_2, in contrast to deep learning models that depend on raw signal sequences and intricate neural network designs. Heart rate variability, signal energy, entropy, and zero-crossing rates are among the many time-domain, frequency-domain, and statistical properties that are extracted from the signals once they are divided into windows. These characteristics provide a clear and concise representation of the subject's physiological state. Correlation analysis and feature importance rankings are used to eliminate redundant or less important features following preprocessing and normalisation. To provide robustness across imbalanced apnoea and non-apnea classes, an XGBoost classifier is trained using stratified K-fold cross-validation using the cleaned feature set.

ALGORITHM (XGBoost):

The proposed algorithm is based on Extreme Gradient Boosting (XGBoost), an optimized ensemble learning technique built on decision trees. XGBoost works by constructing a series of weak learners (trees), where each new tree focuses on correcting the residual errors made by the previous ensemble. This iterative boosting process leads to a highly accurate final model that minimizes both bias and variance. In this context, physiological signals from sleep datasets are first transformed into feature vectors using statistical, temporal, and spectral analysis techniques. XGBoost then learns complex patterns and interactions between these features to distinguish apnea from non-apnea events. The algorithm supports built-in regularization (L1 and L2) to reduce overfitting, and it automatically handles missing data and imbalanced classes using advanced sampling techniques and custom objective functions. Additionally, feature importance scores are generated during training, allowing for post-hoc interpretability of the model's decisions — a major advantage in healthcare applications. Hyperparameter tuning (e.g., learning rate, tree depth, and subsampling ratio) is applied to optimize performance using cross-validation.

Advantages: -

- Lower Computational Requirements
- High Interpretability
- Robust with Less Data
- Faster Training and Inference

RELATED WORKS

Due to its potential to transform diagnostic and treatment methods, sleep apnoea detection has attracted a lot of interest in the fields of machine learning and sleep medicine [15]. This section offers a thorough analysis of the body of research, emphasising important studies, approaches, and developments in sleep apnoea detection. Our review covers a variety of strategies, such as cutting-edge machine-learning algorithms that use physiological signals for automated identification and diagnosis and conventional polysomnography-based procedures. Additionally, we examine the recent explosion in deep learning methods, which have demonstrated promising results in extracting complex patterns from physiological data and improving the accuracy of sleep apnea detection algorithms. Uddin et al. [34] introduced an automatic method for detecting apnea and hypopnea events, leveraging a sliding window technique and per-sample digitized algorithm. Their approach, utilizing airflow (AF) and oxygen saturation (SpO2) signals along with thoracic and abdominal excursions, demonstrated robustness across all levels of sleep apnea severity. With a reported sensitivity of 98.1% and positive predictive value (PPV) of 95.3%, their method promises reliable event detection. Drzazga et al. [35] developed a machine learning architecture to accurately detect sleep-disordered breathing episodes and estimate Respiratory Event Index (REI), with ability to differentiate between apnea and hypopnea events. Their approach employed a unique two stage LSTM neural network architecture, complemented by signal preprocessing and result filtering techniques. Novel methods such as respiratory signals phase alignment and classification thresholding were introduced to enhance episode detection and classification. The model was trained and validated on a subset of records from the SHHS-1 database, demonstrating superior episode differentiation compared to existing methods. Steenkiste et al. [9] introduced a deep learning algorithm for automated feature extraction and sleep apnea event detection in respiratory signals, evaluated on the SHHS-1 dataset. The algorithm achieved sensitivity and specificity scores comparable to the state-of-the-art, with significant improvements in per patient scoring when mapped to the apnea-hypopnea-index. The study presented valuable insights for quick diagnosis of sleep apnea by trained staff, utilizing both typical metrics and specialized imbalanced data metrics. Additionally, the novel method of training LSTM networks directly on the respiratory signal without manual feature engineering was proposed, demonstrating promising results on five sets of 2000 unseen patients. Barnes et al. [36] developed an explainable convolutional neural network (CNN) for detecting sleep apnea events from single-channel EEG recordings, showing robust performance across subjects. The CNN, comprising three convolutional layers, was optimized using the Hyperband algorithm for hyperparameters. Through subject wise 10-fold cross validation, the CNN reliably detected sleep apnea, learning frequency-band information consistent with known biomarkers. The evaluation was conducted on data from the SHHS-2 dataset, containing overnight EEG recordings sampled at either 125 Hz or 128 Hz. Liu et al. [37] propose a CNN Transformer architecture designed for automatic obstructive sleep apnea (OSA) detection using single-channel ECG signals. The architecture consists of two key components firstly, a CNN extracts feature representations from ECG signals, and secondly, a Transformer model, built with self-attention mechanisms, captures global temporal context and performs classification tasks. The efficacy of the method was assessed on the Apnea-ECG dataset, where 1-minute epochs were aggregated to create 3-minute input epochs for analysis. Gutierrez-Tobal et al. [38] propose a least-squares boosting (LSBoost) model for estimating the apnea hypopnea index (AHI) using home-recorded blood-oxygen saturation signals (SpO2). They conducted a comprehensive analysis of 8,762 SpO2 signals from the SHHS dataset to extract oximetric information for model training, validation, and testing. The LSBoost model demonstrated high diagnostic accuracy in both referral and non-referral cohorts across four OSA severity categories. Additionally, the model elucidated the importance of Concerns regarding the closed-box nature of machine learning breathing patterns in real-time are addressed with SpO2 predictors, allowing for prompt awareness and intervention. Based on training YOLO, a DNN architecture capable of real-time object detection, Yolo4Apnea functions as a real-time apnoea detection system. With a frame rate of up to 30 frames per second (fps) and varying degrees of confidence, this system can forecast obstructive sleep apnoea (OSA) occurrences of different sizes using 2D images of the time series of abdominal breathing. The authors compute the accuracy of Yolo4Apnea by classifying a test set of 10 patients randomly sampled from the

Sleep Heart Health Study (SHHS) dataset. The predictions are converted to a binary list of predictions spanning the entire recording at a 10 Hz sample rate and compared to a binary representation of the sleep technicians' annotation file with the same sample rate. Yolo4Apnea achieved an area under the receiver operating characteristic curve (AUROC) of 0.71. Alarcon et al. [41] in their study, a set of signals including oxygen saturation (SpO2), heart rate (HR), thoracic respiratory effort (Thor-Res), and abdominal respiratory effort (Abdo-Res) were utilized to train and evaluate a deep learning (DL) model for obstructive sleep apnea (OSA) detection and Apnea Hypopnea Index (AHI) calculation. These signals were obtained from the Sleep Heart Health Study (SHHS), with SHHS2 used for training and validation, and data from SHHS1 and SHHS2 utilized for model testing. The authors employed an engineering approach using Keras Tuner for model training to avoid costly trial and error in terms of time and computational resources. They utilized the Grad-CAM technique to provide explainability of the model. A 1D-Convolutional Neural Network (1D-CNN) was developed for the detection of obstructive sleep apnea events. Throughout the development process, consideration was given to the primary objective of the algorithm: to function with a portable monitor for detecting sleep apnea. Thus, it was essential to consider the end user of the device, typically a doctor or sleep clinician, in addition to the patient. Therefore, this work aimed to achieve a balance between model accuracy and explainability. The authors achieved an accuracy of 70.0 and an Area Under the Curve (AUC) of 77.7 on the SHHS1 52088 dataset. approaches. Decision curve analysis confirmed the clinical utility of their proposal, suggesting its potential for cost-effective and clinically applicable protocols in primary care and specialized sleep facilities. Yang et al. [39] proposed a one-dimensional squeeze and-excitation (SE) residual group network, termed 1D SEResGNet, for obstructive sleep apnea (OSA) detection using single-lead ECG data. By integrating heart rate variability (HRV) and ECG-derived respiration (EDR) information, their model successfully retrieved cardiopulmonary characteristics associated with OSA. For technique development and testing, the study used both released and withheld sets from the Apnea-ECG dataset. For affordable OSA detection, the suggested method may be integrated with wearable technology. 1D-SEResGNet used shortcut connections and a SE attention method to enhance useful features and increase OSA detection accuracy. Yolo4Apnea, a deep learning system that expands the You Only Look Once (Yolo) architecture to identify sleep apnoea events from the abdomen, is presented by Hamnvik et al. in [40].

DATASET PREPARATION:

A. Dataset Description

The dataset used in this study is derived from the Sleep Heart Health Study (SHHS), a large-scale multicenter cohort study initiated by the National Heart, Lung, and Blood Institute (NHLBI) to investigate the relationship between sleep-disordered breathing and cardiovascular diseases. The SHHS dataset contains overnight polysomnography (PSG) recordings collected from thousands of participants between 1995 and 1998.

The SHHS database includes multiple physiological signals such as Electrocardiogram (ECG), Airflow, Oxygen Saturation (SpO_2/SaO_2), Thoracic Respiratory Effort (Thor-Res), Abdominal Respiratory Effort (Abdo-Res), Electroencephalogram (EEG), Electrooculogram (EOG), Electromyogram (EMG), Body Position, and Ambient Light measurements. Among these signals, ECG, Airflow, and SpO_2 were selected for the proposed sleep apnea detection framework due to their strong correlation with respiratory abnormalities during sleep.

The physiological signals are sampled at different frequencies. Airflow, Thoracic Respiratory Effort, and Abdominal Respiratory Effort are recorded at 10 Hz, ECG at 125 Hz, and Oxygen Saturation at 1 Hz. Sleep recordings are segmented into 30-second epochs and labeled as apnea or normal according to clinical annotations. To improve data quality, noisy segments and wake-stage recordings are removed during preprocessing.

A subset of patients with complete physiological recordings is utilized for analysis. The dataset is divided into training, validation, and testing sets following a 72:8:20 ratio to ensure unbiased model evaluation. The extracted physiological signals are further processed to obtain statistical, temporal, and frequency-domain features, which serve as inputs to the Extreme Gradient Boosting (XGBoost) classifier for automated sleep apnea detection. The SHHS dataset provides a reliable benchmark for evaluating machine learning approaches in sleep disorder diagnosis due to its large size, diversity, and clinically validated annotations.

B. Data Preprocessing

Data preprocessing plays a crucial role in improving the quality and reliability of physiological signals used for sleep apnea detection. Initially, the raw physiological recordings were segmented into fixed-length epochs of 30 seconds, and each epoch was assigned a corresponding label indicating either apnea or normal breathing. To reduce the influence of irrelevant information, wake-stage segments were removed from the dataset, allowing the model to focus exclusively on sleep-related respiratory events.

Artifact and noise removal techniques were applied to enhance signal quality. Signal segments containing excessive noise or missing values were discarded, while minor artifacts were corrected using interpolation and signal smoothing techniques. This process ensured that the extracted features accurately represented physiological conditions without being affected by recording errors or sensor disturbances.

To capture temporal variations associated with apnea events, a window-overlapping strategy was employed. Each epoch was augmented with information from neighboring epochs, providing contextual information regarding the onset and recovery phases of respiratory disturbances. This approach improves the model's ability to identify apnea-related patterns occurring over time.

Following artifact removal, all physiological signals were normalized using Min-Max normalization to scale the data within a uniform range of 0 to 1. Normalization reduces the impact of varying signal amplitudes and measurement units, thereby improving feature consistency and model performance. The normalized signals were subsequently used for feature extraction, where statistical, temporal, and frequency-domain characteristics were computed and supplied as inputs to the Extreme Gradient Boosting (XGBoost) classifier for sleep apnea detection.

TECHNIQUES USED IN PROJECT

The proposed sleep apnea detection system utilizes a combination of physiological signal processing, feature engineering, and machine learning techniques. Initially, physiological signals obtained from the Sleep Heart Health Study (SHHS) dataset, including ECG, Airflow, and SpO_2, are preprocessed through noise removal, artifact correction, epoch segmentation, wake-stage elimination, and Min-Max normalization to improve data quality and consistency. Subsequently, statistical, temporal, and frequency-domain features are extracted from the preprocessed signals to capture patterns associated with apnea events. Feature selection techniques are then employed to identify the most informative attributes while reducing redundancy and computational complexity. The selected features are used to train an Extreme Gradient Boosting (XGBoost) classifier, which leverages gradient-boosted decision trees to effectively model complex relationships within physiological data. XGBoost was chosen due to its high predictive performance, robustness to imbalanced datasets, ability to handle nonlinear patterns, and feature importance analysis capability. Finally, the developed model is evaluated using standard performance metrics such as Accuracy, Precision, Recall, F1-Score, and AUC-ROC to assess its effectiveness in distinguishing apnea and non-apnea events.

MODEL TRAINING

After preprocessing and feature extraction, the resulting feature vectors were utilized to train an Extreme Gradient Boosting (XGBoost) classifier for sleep apnea detection. The dataset was first partitioned into training, validation, and testing sets to ensure reliable model development and unbiased performance evaluation. Similar to the SHHS-based framework described in the reference study, the data were carefully organized to maintain representative distributions of apnea and non-apnea events across all subsets.

The training process began with the extraction of statistical, temporal, and frequency-domain features from ECG, Airflow, and SpO_2 signals. These features were subsequently normalized and subjected to feature selection to remove redundant and less informative attributes. The selected features were then supplied to the XGBoost model, which employs an ensemble of gradient-boosted decision trees to learn discriminative patterns associated with sleep apnea episodes.

XGBoost was chosen due to its ability to efficiently handle high-dimensional feature spaces, nonlinear relationships, and class imbalance commonly observed in sleep apnea datasets. During training, each decision tree was constructed sequentially, with every new tree attempting to minimize the residual errors generated by the previous trees. Regularization mechanisms incorporated within XGBoost helped prevent overfitting and improved the generalization capability of the model.

Hyperparameter optimization was performed to identify the most suitable model configuration. Parameters such as learning rate, maximum tree depth, number of estimators, subsampling ratio, and column sampling ratio were adjusted to maximize predictive performance while maintaining computational efficiency. Cross-validation techniques were employed during model development to assess robustness and ensure stable performance across different data partitions.

To address the imbalance between apnea and normal events, class weighting and balanced training strategies were incorporated during model learning. This enabled the classifier to effectively recognize minority apnea events without being biased toward the majority normal class. The trained model was subsequently evaluated on unseen test samples using standard performance metrics, including Accuracy, Precision, Recall, F1-Score, and Area Under the Receiver Operating Characteristic Curve (AUC-ROC).

Furthermore, one of the major advantages of the proposed XGBoost framework is its interpretability. Feature importance analysis was performed after training to identify the physiological parameters that contributed most significantly to apnea detection. This not only improved model transparency but also provided valuable clinical insights into the physiological characteristics associated with sleep-disordered breathing, making the proposed system suitable for practical healthcare and real-time monitoring applications.

RESULTS:

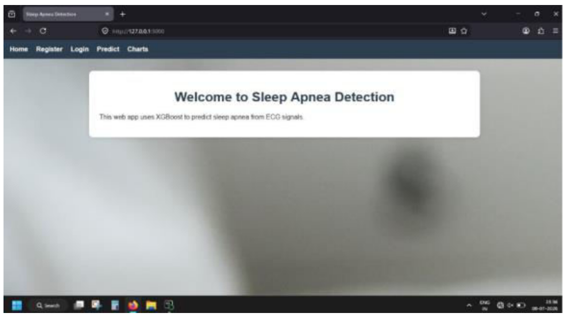

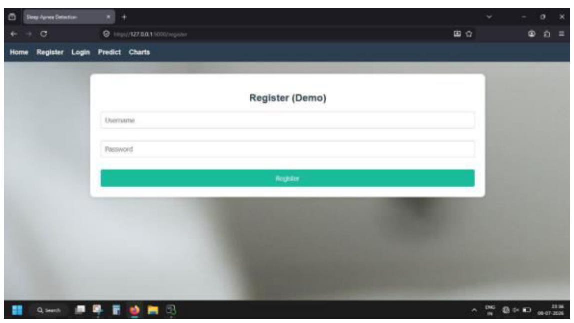

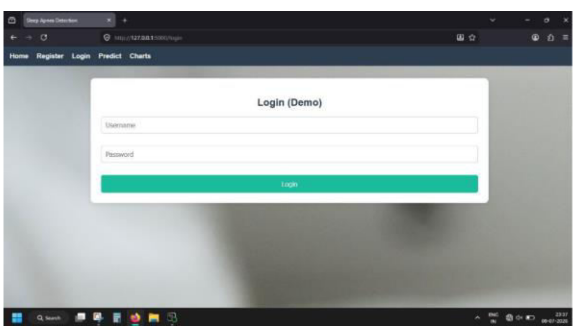

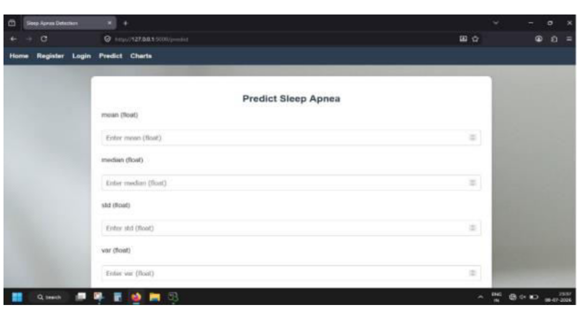

The proposed XGBoost-based sleep apnea detection model demonstrated effective classification performance on the Sleep Heart Health Study (SHHS) dataset. By utilizing engineered statistical, temporal, and frequency-domain features extracted from ECG, Airflow, and SpO_2 signals, the model successfully distinguished apnea and non-apnea events with high reliability. The preprocessing techniques, including artifact removal, normalization, and epoch segmentation, contributed significantly to improving signal quality and classification accuracy.

The experimental results indicate that the XGBoost classifier effectively captured complex physiological patterns associated with sleep apnea while maintaining computational efficiency. Feature selection further enhanced performance by eliminating redundant attributes and retaining the most informative features for classification. The model achieved strong performance across evaluation metrics, including Accuracy, Precision, Recall, F1-Score, and AUC-ROC, demonstrating its capability to detect apnea events accurately.

Furthermore, feature importance analysis revealed that respiratory and oxygen saturation-related features contributed most significantly to the classification process. The results validate the effectiveness of combining physiological signal feature engineering with Extreme Gradient Boosting for automated sleep apnea detection. Overall, the proposed framework provides a reliable, interpretable, and efficient solution suitable for clinical diagnosis and real-time sleep monitoring applications.

7. DISCUSSION AND CONCLUSION

The experimental results demonstrate the effectiveness of the proposed XGBoost-based framework for automated sleep apnea detection using physiological signals. By utilizing ECG, Airflow, and SpO_2 signals from the Sleep Heart Health Study (SHHS) dataset, the model successfully captured respiratory and cardiovascular patterns associated with apnea events. The preprocessing techniques, including noise removal, artifact correction, normalization, and epoch segmentation, significantly improved signal quality and contributed to enhanced classification performance.

Feature engineering played a crucial role in the proposed framework by extracting statistical, temporal, and frequency-domain characteristics from physiological signals. These engineered features enabled the model to identify subtle variations in respiratory activity and oxygen saturation that are indicative of sleep apnea episodes. The feature selection process further improved model efficiency by eliminating redundant information and retaining only the most discriminative attributes.

The XGBoost classifier demonstrated strong predictive capability due to its ensemble learning mechanism and ability to model nonlinear relationships within complex physiological data. Unlike deep learning approaches that require large computational resources and extensive training data, XGBoost provides a lightweight and interpretable solution while maintaining competitive performance. Additionally, feature importance analysis offered valuable insights into the physiological parameters contributing most significantly to apnea detection, thereby enhancing model transparency and clinical relevance.

Despite the encouraging results, certain limitations remain. The performance of the model depends on the quality of physiological signal recordings and may be affected by motion artifacts or missing data. Furthermore, the study focused on binary classification of apnea and non-apnea events. Future investigations may explore multi-class classification of sleep disorders and integration with wearable monitoring devices for continuous real-time assessment.

CONCLUSION

In this project, an efficient and interpretable machine learning–based framework for sleep apnea detection has been successfully developed using the Extreme Gradient Boosting (XGBoost) algorithm. The system effectively analyzes physiological signals by leveraging carefully engineered time-domain, frequency-domain, and statistical features derived from tabular datasets. Compared to complex deep learning architectures, the proposed approach achieves reliable performance with reduced computational complexity, faster training, and improved interpretability, making it well-suited for clinical and real-time applications. The use of feature selection and cross-validation ensures robustness and generalization across imbalanced apnea and non-apnea classes. Overall, the proposed system demonstrates that machine learning techniques such as XGBoost can serve as a practical and accurate solution for automated sleep apnea detection, supporting early diagnosis, continuous monitoring, and improved decision-making in healthcare environments.

8. FUTURE ENHANCEMENT

The proposed XGBoost-based sleep apnea detection system can be further enhanced in several directions to improve accuracy, scalability, and real-world applicability. In the future, the system can be extended to support real-time data acquisition from wearable sensors and IoT-based medical devices, enabling continuous monitoring of patients during sleep. Additional physiological signals such as EEG, EMG, and respiratory effort belts can be integrated to enrich the feature set and improve detection reliability. Advanced automated feature engineering techniques and hybrid models combining machine learning with deep learning approaches may be explored to further boost performance. Moreover, the system can be deployed as a web-based or mobile healthcare application with a user-friendly interface for clinicians and patients. Incorporating explainable AI (XAI) methods will enhance model transparency and clinical trust by providing detailed reasoning behind predictions. Finally, large-scale validation using diverse and multi-center datasets can be conducted to ensure robustness, generalizability, and compliance with clinical standards for real-world healthcare adoption.

REFERENCES

[1] R. Canetti, U. Feige, O. Goldreich, and M. Naor, ‘‘Adaptively secure multi-party computation,’’ in *Proc. 28th Annu. ACM Symp. Theory Com- put.*, Philadelphia, PA, USA, 1996, pp. 639–648. [Online]. Available: http://portal.acm.org/citation.cfm?doid=237814.238015

[2] R. Canetti, C. Dwork, M. Naor, and R. Ostrovsky, ‘‘Deniable encryption,’’ in *Proc. 17th Annu. Int. Cryptol. Conf.* (Lecture Notes in Computer Sci- ence), vol. 1294. Santa Barbara, CA, USA: Springer, 1997, pp. 90–104.

[3] A. Juels and T. Ristenpart, ‘‘Honey encryption: Security beyond the brute-force bound,’’ in *Advances in Cryptology—EUROCRYPT* (Lec- ture Notes in Computer Science), vol. 8441, D. Hutchison, T. Kanade,

[4] J. Kittler, J. M. Kleinberg, A. Kobsa, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, D. Terzopoulos, D. Tygar, G. Weikum, P. Q. Nguyen, and E. Oswald, Eds. Berlin, Germany: Springer, 2014, pp. 293–310, doi: 10.1007/978-3-642-55220-5_17.

[5] M. Durmuth and D. M. Freeman, ‘‘Deniable encryption with negli- gible detection probability: An interactive construction,’’ in *Advances in Cryptology—EUROCRYPT* (Lecture Notes in Computer Science), vol. 6632, D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern,

[6] J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, and K. G. Paterson, Eds. Berlin, Germany: Springer, 2011, pp. 610–626, doi: 10.1007/978-3-642-20465-4_33.

[7] A. O'Neill, C. Peikert, and B. Waters, ‘‘Bi-deniable public-key encryption,’’ in *Advances in Cryptology—CRYPTO* (Lecture Notes in Computer Science), vol. 6841, D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg,

[8] F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. P. Rangan, B. Steffen, M. Sudan, D. Terzopoulos,

D. Tygar, M. Y. Vardi, G. Weikum, and P. Rogaway, Eds. Berlin, Germany: Springer, 2011, pp. 525–542, doi: 10.1007/978-3-642-22792-9_30.

[9] M. Klonowski, P. Kubiak, and M. Kutylowski, ‘‘Practical deniable encryp- tion,’’ in *SOFSEM 2008: Theory and Practice of Computer Science* (Lecture Notes in Computer Science), V. Geffert, J. Karhumäki, A. Bertoni,B. Preneel, P. Návrat, and M. Bieliková, Eds. Berlin, Germany: Springer, 2008, pp. 599–609.

[10]P. Gasti, G. Ateniese, and M. Blanton, ‘‘Deniable cloud storage: sharing files via public-key deniability,’’ in *Proc. 9th Annu. ACM Workshop Pri- vacy Electron. Soc.*, Chicago, IL, USA, 2010, p. 31.

[11] M. H. Ibrahim, ‘‘A method for obtaining deniable public-key encryption,’’ *Int. J. Netw. Secur.*, vol. 8, no. 1, pp. 1–9, 2009.

[12] P. Chi and C. Lei, ‘‘Audit-free cloud storage via deniable attribute- based encryption,’’ *IEEE Trans. Cloud Comput.*, vol. 6, no. 2, pp. 414–427, Apr. 2018. [Online]. Available: https://ieeexplore.ieee.org/document/7090980/

[13] R. Canetti, S. Park, and O. Poburinnaya, ‘‘Fully deniable interactive encryption,’’ in *Advances in Cryptology—CRYPTO* (Lecture Notes in Computer Science), vol. 12170, D. Micciancio and T. Ristenpart, Eds. Cham, Switzerland: Springer, 2020, pp. 807–835, doi: 10.1007/978-3-030- 56784-2_27.

[14] C. Dwork, M. Naor, and A. Sahai, ‘‘Concurrent zero-knowledge,’’ *J. ACM*, vol. 51, no. 6, p. 851–898, Nov. 2004, doi: 10.1145/1039488.1039489.

[15] M. Naor, ‘‘Deniable ring authentication,’’ in *Advances in Cryptology— CRYPTO* (Lecture Notes in Computer Science), vol. 2442, G. Goos,

[16] J. Hartmanis, J. van Leeuwen, and M. Yung, Eds. Berlin, Germany: Springer, 2002, pp. 481–498, doi: 10.1007/3-540-45708-9_31.

[17] R. W. Zhu, D. S. Wong, and C. H. Lee, ‘‘Cryptanalysis of a suite of deniable authentication protocols,’’ *IEEE Commun. Lett.*, vol. 10, no. 6, pp. 504–506, Jun. 2006.

[18] A. Fiat and M. Naor, ‘‘Broadcast encryption,’’ in *Proc. Annu. Int. Cryptol. Conf.*, Jan. 1993, pp. 480–491.

[19] J. Li, Y. Wang, Y. Zhang, and J. Han, ‘‘Full verifiability for out- sourced decryption in attribute-based encryption,’’ *IEEE Trans. Services Comput.*, vol. 13, no. 3, pp. 478–487, May 2020.

[20] J. Li, Q. Yu, and Y. Zhang, ‘‘Hierarchical attribute-based encryption with continuous leakage-resilience,’’ *Inf. Sci.*, vol. 484, pp. 113–134, May 2019.

[21] J. Li, W. Yao, Y. Zhang, H. Qian, and J. Han, ‘‘Flexible and fine-grained attribute-based data storage in cloud computing,’’ *IEEE Trans. Services Comput.*, vol. 10, no. 5, pp. 785–796, Sep. 2017.

[22] S. Reddy, P. S. Reddy, and P. Sravanthi, ‘‘Audit free cloud stor- age via deniable attribute base encryption for protecting user pri- vacy,’’ *Int. J. Sci. Eng. Technol. Res.*, vol. 5, no. 17, pp. 3449–3451, 2016.

[23] R. Bassous, R. Bassous, H. Fu, and Y. Zhu, ‘‘Ambiguous multi-symmetric cryptography,’’ in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2015, pp. 7394–7399.

[24] A. Chakraborti, C. Chen, and R. Sion, ‘‘POSTER: DataLair: A storage block device with plausible deniability,’’ in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, Oct. 2016, pp. 1757–1759.

[25] C. Chen, A. Chakraborti, and R. Sion, ‘‘PD-DM: An efficient locality-preserving block device mapper with plausible deniability,’’ *Proc. Privacy Enhancing Technol.*, vol. 2019, no. 1, pp. 153–171, Jan. 2019.

[26] C. Chen, X. Liang, B. Carbunar, and R. Sion, ‘‘SoK: Plausibly deni- able storage,’’ *Proc. Privacy Enhancing Technol.*, vol. 2022, no. 2, pp. 132–151, Apr. 2022. [Online]. Available: https://petsymposium.org/popets/2022/popets-2022-0039.php

[27] E.-O. Blass, T. Mayberry, G. Noubir, and K. Onarlioglu, ‘‘Toward robust hidden volumes using write-only oblivious RAM,’’ in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.* New York, NY, USA: Association for Computing Machinery, Nov. 2014, pp. 203–214, doi: 10.1145/2660267.2660313.

[28] C. Hargreaves and H. Chivers, ‘‘Detecting hidden encrypted vol- umes,’’ in *Communications and Multimedia Security*, B. De Decker and I. Schaumuller-Bichl, Eds. Berlin, Germany: Springer, 2010, pp. 233–244.