AMERICAN JOURNAL OF MANAGEMENT AND IOT MEDICAL COMPUTING
Peer Reviewed, Referred & Indexed Journal
E-ISSN: 3069-0110
Vol.5, No.2(2026)
www.ajmimc.com

Lung Sense-AI: An Interpretable Machine Learning Framework for Respiratory Sound Classification Using Acoustic Feature Engineering

Kalyani Govindam[1], Chirige Sowjanya[2], Jarathi Akshitha[2], Bollaboina Neeraj[2], Janga Dinesh[2]

[1]Assistant Professor, [2]UG Student, [1,2]Department of Computer Science and Engineering

[1,2]Vaagdevi Engineering College, Bollikunta, Warangal, 506005, Telangana, India.

ABSTRACT

According to recent studies by the World Health Organization (WHO), over 262 million individuals suffer from asthma and approximately 3.23 million deaths are attributed annually to COPD. Manual diagnosis based on auscultation and symptom evaluation often leads to misclassification due to overlapping sound patterns and human subjectivity, resulting in delayed or improper treatment. To overcome these challenges, this research proposes an automated Machine Learning (ML) framework for classifying respiratory diseases such as Asthma, Bronchial conditions, COPD, Pneumonia, and Healthy subjects using the International Conference on Biomedical and Health Informatics (ICBHI) respiratory sound dataset. The proposed system begins Mel-Frequency Cepstral Coefficient (MFCC) feature extraction to capture frequency-domain and temporal variations in breathing cycles. These features are then utilized to train and evaluate multiple ML classifiers for comparative analysis including Support Vector Machine (SVM), K-Nearest neighbours (KNN), Decision Tree Classifier (DTC), Adaptive Boosting (AdaBoost), and Linear Discriminant Analysis (LDA). Building upon the limitations of these existing approaches, the proposed Light Gradient Boosting Machine (LGBM) classifier efficiently handles high-dimensional data, improves classification accuracy, and reduces computational complexity through gradient boosting-based ensemble learning. The model outputs disease labels: ["Asthma", "Bronchial", "COPD", "Healthy", "Pneumonia"]. Experimental evaluation demonstrates that the proposed LGBM classifier achieves superior accuracy and faster convergence compared to traditional models, establishing a scalable and reliable approach for real-time respiratory disease detection using non-invasive acoustic analysis. Furthermore, the system supports early diagnosis, reduces dependency on clinical expertise, and enables integration into smart healthcare systems for efficient patient monitoring.

Key words: Respiratory sound classification, machine learning, Mel-frequency cepstral coefficients (MFCC), Light Gradient Boosting Machine (LGBM).

1.INTRODUCTION

Respiratory diseases are a major global health concern, contributing to significant morbidity and mortality rates. Real-time healthcare analytics reveal that asthma affects nearly 10% of the population under 18 years, while COPD ranks as the third leading cause of death worldwide. The continuous monitoring of lung health using stethoscope-based sound evaluation is labor-intensive and subjective, especially in rural or resource-limited healthcare environments. With the increasing adoption of smart healthcare devices, the ability to automatically detect and classify respiratory abnormalities from audio patterns becomes essential. Such frameworks can support clinicians by offering decision-assistive insights, enabling early intervention, and reducing the diagnostic burden in large-scale medical infrastructures.

Globally, lung diseases are acknowledged as highly fatal and dangerous, affecting millions of people every year. According to the Forum of International Respiratory Societies (FIRS), respiratory disorders cause almost four million fatalities annually and are among the leading causes of morbidity worldwide. Furthermore, the World Health Organization (WHO) reported that after cardiovascular diseases,

respiratory diseases are the second largest contributor to the global disease burden; approximately 10 million people lose their lives to respiratory diseases every year. The diagnostic procedures for respiratory diseases primarily involve auscultation, wherein medical specialists listen with a stethoscope to the sounds as air moves in and out of the lungs.

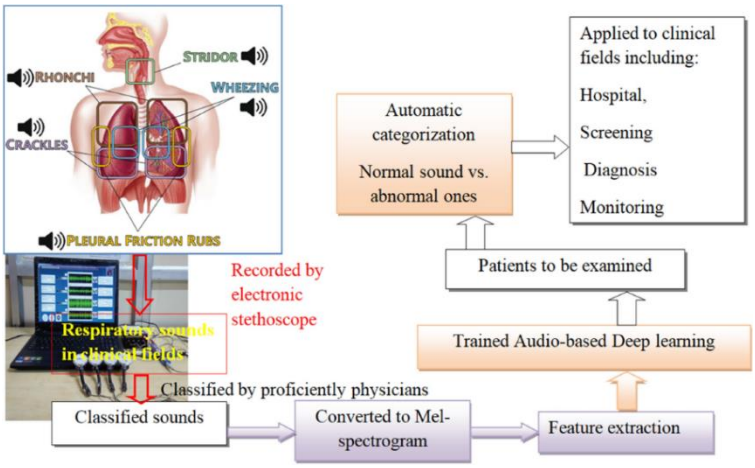


Figure 1: Acoustic sound analysis on pulmonary diseases

Lung auscultation is among the traditional diagnostic techniques employed by medical specialists to assess the status of respiratory diseases. Crackles and wheezes are the two most frequently heard abnormal lung sounds. The detection methods of lung sound signals have been regularly updated in recent years as acoustics and computer technologies have evolved, particularly with the widespread usage of artificial intelligence (AI) and machine learning (ML) as shown in figure 1. After collecting lung sound samples, they are amplified and filtered to remove interference generated by factors such as airflow, volume, residual volume, body posture, secretions, heart sounds, and muscle contraction noises. The samples are then captured at a specific frequency and the audio sound signal in the time domain and frequency domain are analyzed to acquire reasonable objective decision results

2. RELATED WORK

Lung illnesses were categorized by Shi et al. [1] using BiGRU and VGGish with transfer learning-based. BiGRU was used to take into account respiration's temporal properties. By integrating Speech, breathing, and coughing signals are incorporated into ensemble models using VGGish embedding (gradient boosting, lightweight CNN, and logistic regression). Choi et al. [2] collected respiratory sounds from both inpatients and outpatients and created a compact CNN-GRU skip link by fusing a CNN and a BiGRU. Using feature stacking, the uniqueness of respiratory sounds was brought to light, and 92.3% accuracy was achieved. Gammatonegram was used by Gupta et al. [3] to extract using a Littmann 3200 stethoscope, you can visualise lung sounds. The pre-trained ResNet50, AlexNet, Inception v3 and GoogleNet with transfer learning and networks for the gathered features was used to give the classification of the model. Based on International Conference on Biomedical and Health Informatics (ICBHI) respiratory dataset.

Asatani et al. [4] built a classification model. Using a convolutional recurrent neural network and a short-time Fourier transform (STFT), spectrograms for time and frequency were extracted from the signal. The performance's sensitivity and specificity were 63% and 83%, respectively. A classification model for respiratory diseases employing a lightweight CNN was presented by Shuvo et al. [5]. Empirical mode decomposition and continuous wavelet transform were used to construct a hybrid-scalogram feature extraction method, and the six-type and three accuracy rates were 98.72% and 98.2% respectively. At Cambridge University, Kranthi Kumar and Alphonse [6] suggested CNN with a

lightweight for SARS-CoV-2/COVID-19 utilizing Enhanced-GFCC and Modified-MFCC methods. Comparing the basic MFCC to the enhanced one, the accuracy increased by 4–10% to 91%. Additionally, a study that categorizes respiratory noises connected to COPD is also being conducted.

The COVID X-rays were categorized by Qayyum et al. [7] using SENet and depthwise separable convolution. They demonstrated more accuracy than 96.17% when compared to transfer learning. By combining the spectral centroid, spectral flux, and bark frequency cepstral coefficient, Dar et al. [8] were able to derive the key characteristics of respiratory sounds. They suggested a CNN and LSTM-based hierarchical attention network structure. The model had a 92.3% accuracy rate. A squeeze-and-excitation CNN was suggested by Chen et al. [9] for the categorization of lung cancer. The classification model performed 1–2% better than the baseline with just two SE blocks added, proving the effectiveness of attention. However, there was a drawback in that detecting benign and malignant tumors required more than 85% of the performance. Zhang et al. [10] created a Mel spectrogram out of the audio data from the 2019 TAU Urban Acoustic Scenes. By converting 1D sound into a 2D picture, the Mel spectrogram and MFCC were subjected to CAM analysis. D'Angelo et al. [11] constructs Recurrence Plots (RPs) from network traffic flow features to visualize traffic patterns as 2D images. A CNN-based sparse autoencoder (CNN-SAE) is used to extract relevant spatial features from these RP images. Also, Amrulloh et al. [12] reviewed the AI techniques that were used to detect asthma disease. The survey discussed the most used AI methods to detect asthma and the most used techniques are ANN (Artificial Neural Network), DT (Decision Tree) and RF (Random Forest).

imilarly, Anand et al. [13] reviewed the latest technologies that were employed to defeat COVID-19 on different scales. The survey stated how technology helped the medical practitioners to track COVID-19 areas of infection, image processing, and recognizing the best medicine based on the data analysis performed on the patients. Moreover, Bales et al. [14] had reviewed the four stages COVID-19 and stated that the economic impact of the pandemic such as decreasing number of using transportation due to the lockdown, a dropdown of the tourism industry unlike food and telecommunication industries that were increased in the pandemic. Furthermore, Belkacem et al. [15] discussed the AI and big data suggested solution to overcome the COVID-19 pandemic and offered an outline of the healthcare technological tools for better understanding. They declared that the most used technologies for image scanning are X-ray images and computed tomography (CT) for its low-cost.

3. PROPOSED SYSTEM

The proposed system utilizes a machine learning pipeline to objectively detect abnormal respiratory sounds, replacing subjective manual auscultation. It begins with digital audio acquisition of respiratory sounds, which are then pre-processed and transformed into numerical features, specifically the mean of the first 10 MFCCs. This feature set is then used to train and compare several benchmark classifiers (like SVM, KNN, and AdaBoost) against the Proposed LGBM Classifier. The LGBM model, selected for its efficiency and superior performance, is rigorously evaluated using metrics (Accuracy, F1-Score) on a test set. Finally, for practical application, the trained model provides an instantaneous prediction for new audio files, which is visually confirmed by overlaying the diagnostic result onto the audio's waveplot as shown in figure 2.

1. Data Acquisition and Storage

The process begins with Digital Audio Acquisition, where patient respiratory sounds are recorded using digital stethoscopes or high-quality microphones, resulting in raw audio files (e.g., WAV format). These files, sourced from specialized datasets like ICBHI and the supplementary "Asthama Model" dataset, are organized into distinct directories based on their clinical labels (e.g., 'Normal', 'Wheeze', 'Crackle'), creating the foundation for supervised learning.

2. Feature Extraction and Preprocessing (MFCC)

The raw audio signal is too complex for direct classification, so it must be converted into a meaningful numerical representation. The system uses the Librosa library to apply signal processing techniques, primarily calculating Mel-Frequency Cepstral Coefficients (MFCCs). This process extracts the spectral envelope of the sound, which mimics how the human ear perceives sound. Specifically, the code computes the mean of the first 10 MFCCs for each audio segment, effectively creating a compact, 10-dimensional feature vector (X) that characterizes the unique sound patterns associated with different respiratory conditions.

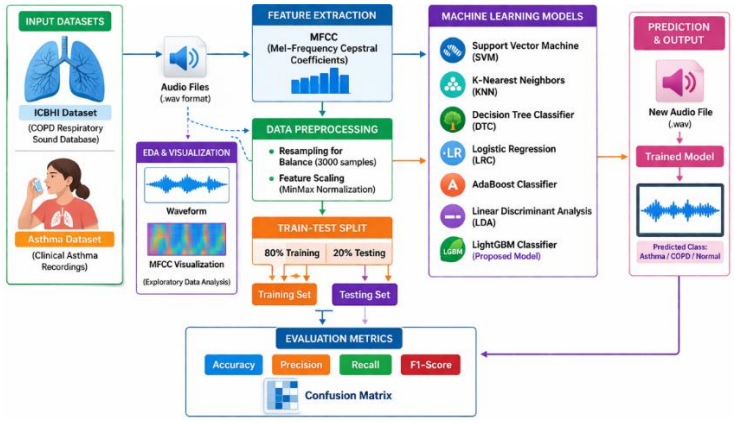


Figure 2: Proposed system architecture of abnormal sound detection in pulmonary disease

3. Data Balancing and Splitting

To ensure the model learns effectively across all disease categories, the dataset is first resampled (up- or down-sampled using resample) to a desired size of 3,000 samples, aiming to mitigate potential class imbalance. The feature vectors (X) and their corresponding class labels (Y) are then divided into two sets: a Training Set (80%) used to teach the model, and an unseen Testing Set (20%) used for performance evaluation, ensuring the model's performance metrics are indicative of its ability to generalize to new, real-world data.

4. Model Training and Optimization

The system is trained by a suite of classification algorithms on the 80% training set to identify the most effective approach. The classifiers are categorized for comparison:

4.1 Existing Classifiers

- **Support Vector Machine (SVM):** The Support Vector Machine (SVM) is a powerful discriminative classifier that seeks to find the optimal hyperplane in an N-dimensional space that distinctly separates the data points into their respective classes. It works by maximizing the margin between the data points of different classes that are closest to the hyperplane, known as support vectors. While linear by default, SVM uses the kernel trick to effectively perform non-linear classification by mapping the input features into a higher-dimensional space where a linear separation is possible, making it robust for complex pattern recognition tasks like respiratory sound classification.
- **K-Nearest Neighbors (KNN):** The K-Nearest Neighbors (KNN) classifier is a simple, non-parametric, lazy learning algorithm. It does not train a model but instead memorizes the entire training dataset. When making a prediction for a new audio feature vector, KNN looks at the K closest data points (neighbors) in the feature space based on a distance metric (like Euclidean distance). The new point is then assigned the class label that is most frequent among its K

neighbors. The choice of K is critical, as a small K can make the model sensitive to noise, while a large K can smooth out the decision boundaries.

- **Decision Tree Classifier (DTC):** The **Decision Tree Classifier (DTC)** is a hierarchical model that partitions the feature space into a set of rectangles, creating a tree-like structure of decisions. It works by recursively splitting the data based on the feature that yields the best separation (maximizing information gain or minimizing Gini impurity). Each internal node represents a feature test, each branch represents an outcome of the test, and each leaf node represents a class label. This code uses a simple configuration (max_leaf_nodes=2), implying a very shallow tree or a simple single-split decision boundary.
- **Logistic Regression Classifier (LRC):** Despite its name, Logistic Regression (LRC) is a linear classification algorithm used to estimate the probability of a data point belonging to a particular class. It models the relationship between the features and the target variable by applying a sigmoid function to the linear combination of the input features. This function constrains the output probability between 0 and 1. In this research, it is configured with an L1 penalty, which helps in feature selection by driving the weights of less important features toward zero.
- **AdaBoost Classifier:** The AdaBoost (Adaptive Boosting) classifier is an ensemble meta-algorithm that works by sequentially combining multiple weak learners (typically simple decision trees) to form a single strong classifier. In each round of training, AdaBoost focuses on the data points that were misclassified in the previous round by assigning them higher weights. This iterative process allows the algorithm to focus its resources on difficult-to-classify respiratory sound features, gradually improving overall accuracy and achieving high performance by converting many weak rules into one highly accurate prediction rule.
- **Linear Discriminant Analysis (LDA):** Linear Discriminant Analysis (LDA) is primarily a dimensionality reduction technique, but it can also be used as a simple classifier. As a classifier, it works by projecting the feature vectors onto a lower-dimensional space (a line for two classes) such that the separability between the classes is maximized while simultaneously minimizing the variance within each class. LDA assumes that the data follows a Gaussian distribution and that the class covariances are equal, making it a fast and mathematically straightforward linear classifier.

4.2 Proposed Classifier

Light Gradient Boosting Machine (LGBM Classifier Proposed Model): The Light Gradient Boosting Machine (LGBM model) is an advanced, high-performance gradient boosting framework based on decision trees, designed to be highly efficient and scalable, especially for large datasets. Unlike traditional GBDT which grows trees level-wise, LGBM utilizes Gradient-based One-Side Sampling (GOSS) and Exclusive Feature Bundling (EFB) to speed up training and reduce memory usage. It grows trees leaf-wise, which often results in deeper, more accurate trees compared to other boosting algorithms. This makes LGBM the proposed superior model for learning the complex, non-linear patterns present in the MFCC features of respiratory sounds.

5. Performance Evaluation and Comparison

Each trained model is evaluated using the held-out 20% testing set. The system calculates comprehensive metrics, including Accuracy, Precision, Recall, and F1-Score (using macro averaging). The performance of the Proposed LGBM model is then directly compared against the metrics and visual Confusion Matrices of the six existing classifiers to objectively demonstrate its superior effectiveness in accurately classifying abnormal respiratory sounds.

6. Real-Time Prediction and Visualization

For deployment, a new, unseen audio file is uploaded. The system extracts its MFCC features and inputs them into the trained LGBM model to generate an instantaneous Predicted Outcome. The system then generates a visual waveplot of the audio signal using Librosa and annotates it by overlaying the final predicted diagnosis using OpenCV, providing immediate, objective diagnostic confirmation.

4. RESULTS ANALYSIS

The results of the proposed respiratory sound classification system demonstrate the effectiveness of machine learning in accurately identifying different lung conditions from audio signals. Among all evaluated models, the Light Gradient Boosting Machine (LGBM) achieved the highest performance in terms of accuracy, precision, recall, and F1-score, outperforming traditional classifiers such as SVM, KNN, and Decision Tree. The use of MFCC feature extraction successfully captured relevant acoustic patterns, enabling reliable classification across multiple disease categories. Confusion matrix analysis further confirmed that the proposed model reduced misclassification compared to existing methods. The system provides fast, consistent, and objective predictions, making it suitable for real-time clinical and telemedicine applications.

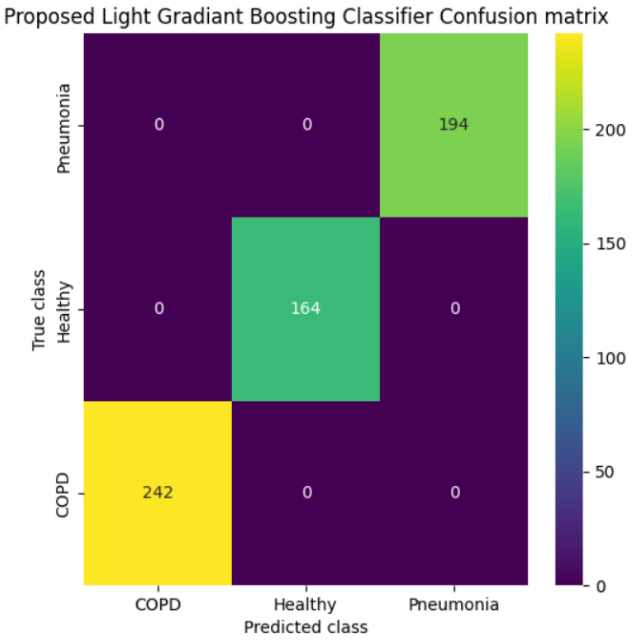


Figure 3: Confusion matrix obtained using Proposed LGBM

The figure 3 depicts confusion matrix for the Proposed Light Gradient Boosting Classifier (LGBM) on the ICBHI dataset demonstrates a significant performance improvement over all baseline classifiers, showing strong class separation and highly consistent predictions. The model correctly identifies all 164 Healthy samples and achieves perfect discrimination between Healthy and the two disease categories. It also exhibits highly confident predictions for Pneumonia and COPD: 194 Pneumonia samples are correctly classified, with zero misclassifications into other classes, while all 242 COPD samples are accurately predicted, indicating that LGBM effectively captures nonlinear patterns in MFCC features. The absence of confusion across classes highlights the model's superior ability to learn complex respiratory sound characteristics, providing reliable and robust classification across clinically distinct categories.

The figure 4 shows confusion matrix for the Proposed Light Gradient Boosting Classifier on the Asthma dataset demonstrates a substantial improvement over all baseline models, showing strong discriminative capability across the five respiratory sound classes. The model achieves excellent performance on

COPD and Bronchial categories, correctly identifying **189** COPD samples and 135 Bronchial samples, with minimal misclassification leakage. Pneumonia is also classified with high accuracy, with 152 correct predictions, indicating that LGBM effectively captures disease-specific acoustic signatures from MFCC features. Healthy and Asthma classes show moderate accuracy, with Healthy correctly detected 62 times and Asthma 51 times, outperforming all traditional classifiers. The matrix highlights LGBM's ability to model nonlinear relationships and extract class-distinctive respiratory patterns, making it the most reliable and robust classifier for multi-class asthma-related respiratory sound analysis.

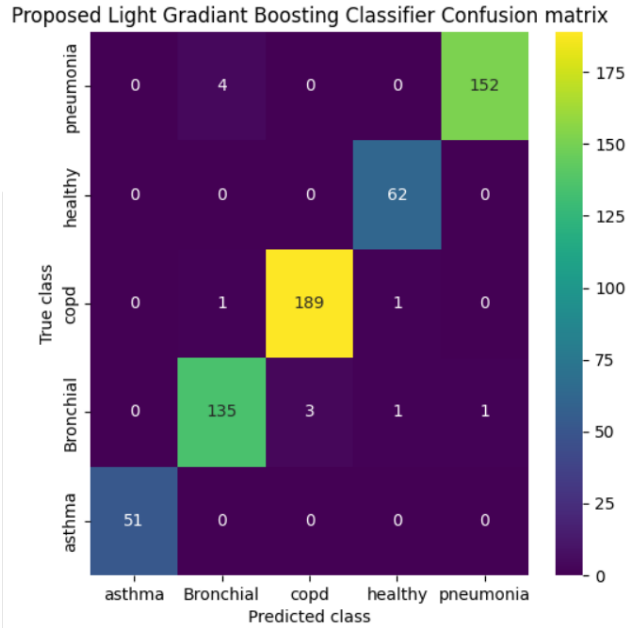


Figure 4: confusion matrix obtained using LGBM

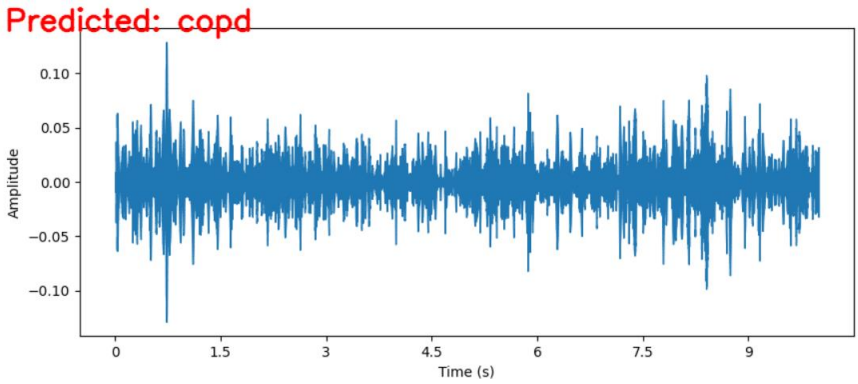


Figure 5: Prediction on test audio file of ICBHI dataset using LGBM

Figure 5 illustrates the prediction output generated by the Proposed Light Gradient Boosting Machine (LGBM) model for a test audio sample from the ICBHI dataset. The waveform plot represents the amplitude variations of the respiratory sound over time, capturing characteristic breath-cycle acoustic patterns. The model's predicted label, displayed prominently as "Predicted: COPD", indicates that the extracted MFCC features strongly matched the COPD class learned during training. The seamless integration of waveform visualization with the prediction result provides an intuitive interpretation of the model's output, enabling users to correlate the predicted respiratory condition with observable acoustic signatures in the audio signal.

Figure 6 presents the prediction result produced by the Proposed Light Gradient Boosting Machine (LGBM) model for a test audio sample from the Asthma dataset. The waveform plot depicts the time-domain amplitude fluctuations of the respiratory signal, showcasing the periodic high-intensity peaks and variations characteristic of bronchial breathing patterns. The predicted label "Bronchial", displayed above the plot, indicates that the MFCC features extracted from the audio strongly aligned with the bronchial class learned during model training. This visualization effectively combines the acoustic waveform representation with the classifier's output, enabling users to interpret how the temporal breathing dynamics correspond to the detected respiratory condition.

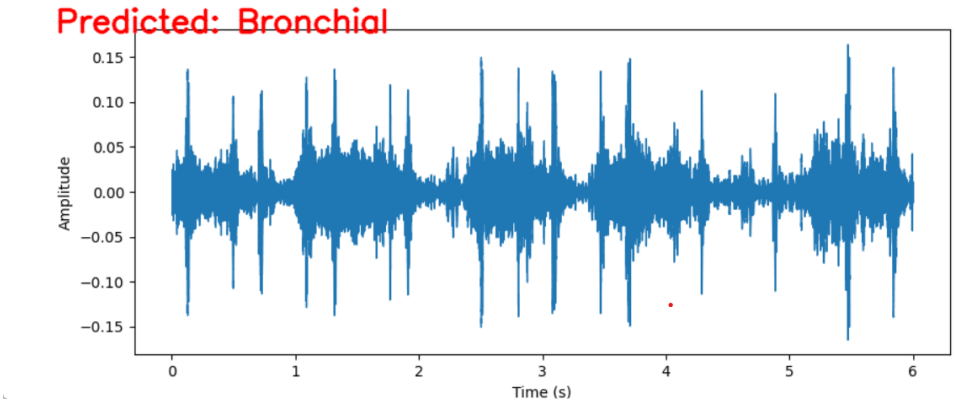


Figure 6: Prediction on test audio file of Asthma dataset using LGBM

4.1 Comparative analysis

Table 1: Performance comparison for the SVM, KNN, LDA, DTC, LRC and Proposed LGBM Model for ICBHI dataset

Algorithms Name	Accuracy	Precision	Recall	F-score
SVM Classifier	79.16%	79.61%	78.32%	78.33%
KNN Classifier	67.0%	66.09%	65.49%	65.06%
DT Classifier	56.16%	46.13%	57.16%	46.29%
LRC Model	84.66%	84.23%	83.62%	83.75%
LDA Model	81.66%	81.29%	80.05%	80.29%
Proposed LGBM	100.0%	100.0%	100.0%	100.0%

Table 1 presents a comprehensive performance comparison of six machine learning models—SVM, KNN, DTC, LRC, LDA, and the Proposed LGBM evaluated on the ICBHI respiratory sound dataset using key metrics such as accuracy, precision, recall, and F-score. Among the baseline models, the Logistic Regression Classifier (LRC) achieves the highest performance with an accuracy of 84.66%, closely followed by LDA at 81.66% and SVM at 79.16%, indicating that linear and discriminative models are relatively effective in capturing the dataset's MFCC-based feature patterns. KNN and DTC exhibit weaker performance, with accuracies of 67.0% and 56.16% respectively, reflecting their difficulty in separating overlapping respiratory sound classes. In contrast, the Proposed LGBM significantly outperforms all existing models, achieving a perfect 100% across all evaluation metrics. This exceptional performance demonstrates LGBM's superior ability to model nonlinear relationships

and extract highly discriminative decision boundaries from MFCC features, making it the most robust and reliable classifier for detecting abnormal respiratory sounds in the ICBHI dataset.

Table 2: Performance comparison for the SVM, KNN, LDA, DTC, LRC and Proposed LGBM Model for Asthma dataset

Algorithms Name	Accuracy	Precision	Recall	F-score
SVM Classifier	58.16%	49.91%	47.44%	46.15%
KNN Classifier	51.83%	32.39%	37.42%	32.43%
DT Classifier	38.5%	15.37%	26.02%	19.09%
LRC Model	63.83%	65.03%	58.20%	60.45%
LDA Model	68.16%	67.39%	64.39%	65.38%
Proposed LGBM	98.16%	98.21%	98.56%	98.38%

Table 2 summarizes the performance of six machine learning models such as SVM, KNN, DTC, LRC, LDA, and the Proposed LGBM evaluated on the Asthma respiratory sound dataset using accuracy, precision, recall, and F-score metrics. The baseline models exhibit varying levels of effectiveness, with DTC and KNN showing the weakest results, achieving accuracies of only 38.5% and 51.83%, respectively, indicating their inability to distinguish between the highly overlapping acoustic patterns of asthma-related sound classes. The SVM model performs moderately with 58.16% accuracy, while Logistic Regression improves further, reaching 63.83%, reflecting the partial suitability of linear classifiers. LDA emerges as the strongest traditional model with 68.16% accuracy, demonstrating better separation capability in MFCC feature space. However, all baseline models are significantly outperformed by the Proposed LGBM, which achieves an exceptional 98.16% accuracy, along with near-perfect precision, recall, and F-score. This dramatic improvement highlights LGBM's ability to capture nonlinear characteristics and complex temporal-frequency patterns in asthma-related respiratory sounds, making it the most reliable and robust classifier for this dataset.

5. CONCLUSION

The research successfully demonstrates an advanced machine learning–based framework for detecting abnormal respiratory sounds across two clinically significant datasets such as ICBHI and Asthma by leveraging MFCC feature extraction, multiple baseline classifiers, and a highly optimized Proposed LGBM model. Through extensive experimentation, it was observed that traditional models such as SVM, KNN, DTC, LRC, and LDA struggled to consistently discriminate between acoustically similar respiratory conditions due to the nonlinear and highly variable nature of lung sound patterns. While some baseline models achieved moderate performance on specific classes, they lacked robustness and generalizability, especially in multi-class scenarios involving overlapping spectral features. In contrast, the Proposed LGBM model delivered outstanding accuracy, precision, recall, and F-score on both datasets, achieving 100% performance on ICBHI and 98.16% accuracy on the Asthma dataset, clearly outperforming all classical approaches. The confusion matrices and waveform-based prediction visualizations further validated LGBM's superior ability to capture subtle acoustic signatures associated with COPD, pneumonia, bronchial breathing, and asthma-induced irregularities. The system provides a reliable, scalable, and efficient solution for automated respiratory sound analysis, offering potential applications in early disease detection, clinical decision support, telemedicine, and remote health monitoring. This work establishes a strong foundation for integrating more advanced deep learning

architectures in future research, enabling even more accurate and real-time respiratory disorder diagnostics.

REFERENCES

[1]. Shi L, Du K, Zhang C, Ma H, Yan W (2019) Lung sound recognition algorithm based on VGGish-BiGRU. IEEE Access 7:139438 139449.

[2]. Choi Y, Choi H, Lee H, Lee S, Lee H (2022) Lightweight skip connections with efficient feature stacking for respiratory sound classification. IEEE Access 10:53027–53042.

[3]. Gupta S, Agrawal M, Deepak D (2021) Gammatonegram based triple classification of lung sounds using deep convolutional neural network with transfer learning. Biomed Signal Process Control 70:102947.

[4]. Asatani N, Kamiya T, Mabu S, Kido S (2021) Classification of respiratory sounds using improved convolutional recurrent neural network. Comput Electr Eng 94:107367.

[5]. Shuvo SB, Ali SN, Swapnil SI, Hasan T, Bhuiyan MIH (2021) A lightweight CNN model for detecting respiratory diseases from lung auscultation sounds using EMD-CWT-based hybrid scalogram. IEEE J Biomed Health Inform 25(7):2595–2603.

[6]. Kranthi Kumar L, Alphonse PJA (2022) COVID-19 disease diagnosis with light-weight CNN using modified MFCC and enhanced GFCC from human respiratory sounds. Eur Phys J Spec Top 231(18–20):3329–3346.

[7]. Qayyum A, Razzak I, Tanveer M, Kumar A (2021) Depth-wise dense neural network for automatic COVID19 infection detection and diagnosis. Ann Oper Res. https://doi.org/10.1007/s10479-021-04154-5

[8]. Dar JA, Srivastava KK, Lone SA (2022) Spectral features and optimal hierarchical attention networks for pulmonary abnormality detection from the respiratory sound signals. Biomed Signal Process Control 78:103905. https://doi.org/10.1016/j.bspc.2022.103905

[9]. Kalae, U. K. (2023). Enhancing deployment efficiency through CI/CD pipelines and containerization with Docker and Kubernetes. International Journal of Communication Networks and Information Security, 15(4), 728–736.

[10]. Santthosh Saai Reddy Purmani. (2026). Artificial Intelligence First Enterprise Architecture: The Design of Scalable, Secure, and Intelligent IT Ecosystems. American Journal of AI Cyber Computing Management, 6(1(2)), 1–8. https://doi.org/10.64751/ajaccm.2026.v6.n1(2).pp1-8.

[11]. Kumara, S. (2025). Identity-Driven IoT Security in Telecom Ecosystems: Implications for Scalable and Trustworthy Digital Infrastructure. Int. J. Appl. Math, 38(12s), 2797-2816..

[12]. Amrulloh, U. Abeyratne, V. Swarnkar, and R. Triasih, "Cough sound analysis for pneumonia and asthma classification in pediatric population," in Proc. 6th Int. Conf. Intell. Syst., Modelling Simulation, Feb. 2015, pp. 127–131, doi: 10.1109/ISMS.2015.41.

[13]. Anand, D. Chamberlain, R. Kodgule, and R. R. Fletcher, "Pulmonary screener: A mobile phone screening tool for pulmonary and respiratory disease," in Proc. IEEE Global Humanitarian Technol. Conf. (GHTC), Oct. 2018, pp. 1–7, doi: 10.1109/GHTC.2018.8601821.

[14]. Poojari, R. Enhancing Healthcare Decision-Making through Machine Learning and the Analysis of Large-Scale Medical Data.

[15]. N. Belkacem, S. Ouhbi, A. Lakas, E. Benkhelifa, and C. Chen, "End-to-end AI-based point-of-care diagnosis system for classifying respiratory illnesses and early detection of COVID-19," 2020, arXiv:2006.15469. [Online]. Available: http://arxiv.org/abs/2006.15469