A Hybrid Framework for Autism Spectrum Disorder Prediction and Personalized Recommendations

Kavitha Gangaraju
Department of Computer Science and Engineering, M S Ramaiah Institute of Technology, Affiliated to Visvesvaraya Technological University, Belagavi, India
kavithagangaraju567@gmail.com (corresponding author)

H. K. Yogisha
Department of Information Science and Engineering, M S Ramaiah Institute of Technology, Affiliated to Visvesvaraya Technological University, Belagavi, India
yogishhk@msrit.edu

Received: 13 November 2025 | Revised: 17 February 2026 and 30 March 2026 | Accepted: 17 April 2026

 DOI: https://doi.org/10.48084/etasr.16263

ABSTRACT

Early detection of Autism Spectrum Disorder (ASD) is critical for enabling timely interventions and improving developmental outcomes. However, manual screening methods are time-consuming, subjective, and often inconsistent, whereas existing Machine Learning (ML) approaches frequently face challenges in feature selection, class imbalance handling, and generalization across age groups. To address these limitations, this study proposes NeuroXGB-MLP, a hybrid framework that integrates XGBoost-based feature selection with an optimized Multi-Layer Perceptron (MLP) classifier to provide robust and accurate ASD prediction. The work includes accurate classification across toddlers, children, and adults, as well as the development of an ML-based ASD clinical decision-support module offering personalized intervention strategies. The methodology involved preprocessing four ASD datasets, namely Kaggle Toddler, Saudi Toddler, UCI Children, and UCI Adult, extracting relevant sensory-behavioral features, and training NeuroXGB-MLP with Cross-Validation (CV) for reliable predictions. The model achieved superior performance, with accuracies of 99.995% on Kaggle Toddler, 99.994% on Saudi Toddler, 99.996% on the merged Kaggle Toddler and Saudi Toddler datasets, 99.993% on UCI Children, and 99.994% on UCI Adult, respectively, outperforming existing approaches. The findings demonstrate that NeuroXGB-MLP effectively addresses class imbalance and feature redundancy while enabling personalized clinical decision-support.

Keywords-Autism Spectrum Disorder (ASD); NeuroXGB-MLP; XGBoost; Multi-Layer Perceptron (MLP); recommender system; Machine Learning (ML)

I. INTRODUCTION

Early prediction of Autism Spectrum Disorder (ASD) is critical for timely intervention, which can significantly improve developmental outcomes and quality of life for affected individuals [1]. Manual ASD screening and diagnosis, however, are time-consuming, labor-intensive, and prone to subjective variability, limiting accessibility for large populations [2, 3]. Machine Learning (ML) and Deep Learning (DL) techniques have emerged as effective solutions, enabling rapid and accurate ASD screening by analyzing behavioral and sensory patterns across age groups [4, 5]. For instance, authors in [6] to evaluate the psychometric properties of the Chilean-adapted Quantitative Checklist for Autism in Toddlers (Q-CHAT) for early ASD screening. The work involved remote administration of Q-CHAT to 313 toddlers aged 18–24 months using the REDCap environment during pandemic. Using non-probabilistic sampling, internal consistency and predictive validity were evaluated utilizing the Area Under the Receiver Operating Characteristic Curve (AUC-ROC), yielding an AUC-ROC of 0.93.

Authors in [7] developed an ML-based framework for ASD detection and personalized education recommendation. The work merged two toddler ASD screening datasets and applied the Synthetic Minority Oversampling Technique (SMOTE) for handling class imbalance and ML-based feature selection for refining feature selection. Further, the authors presented an ensemble approach, where they combined Random Forest (RF) with eXtreme Gradient Boosting (XGB), achieving 99% accuracy. Authors in [8] presented an ML framework for ASD prediction across different age groups. In this work, the authors

used the Pearson Correlation Coefficient (PCC) and Information Gain (IG) for feature selection, which was followed by classification using Support Vector Machine (SVM), Multi-Layer Perceptron (MLP), Naïve Bayes (NB), K-Nearest Neighbor (KNN), and RF. Evaluations were conducted on various age groups, where SVM achieved better accuracy (98–100%). Authors in [9] developed an interpretable ML framework for ASD prediction by using Principal Component Analysis (PCA) for feature reduction, PCC and Mutual Information (MI) for feature selection, and then stacked an ensemble approach, which combined Categorical Boosting (CB), Extra Trees (ET) and RF, with an Artificial Neural Network (ANN) as the classifier. Further, the authors used Safe-Level SMOTE for handling class imbalance. Evaluations across different age groups showed that the approach achieved 96.96–99.89% accuracy.

Authors in [10] aimed to predict ASD in toddlers and children using ML, and hence employed Logistic Regression (LR), Decision Tree (DT), RF, KNN, and SVM. Findings showed that LR achieved better results, achieving 100% accuracy. Authors in [11] aimed to improve adult ASD prediction by addressing noisy and incomplete data, and hence proposed the MAPLE framework, which utilizes MARVEL for imputing missing values using multiple regression models and a SAFARI module for detecting and removing outliers, improving dataset quality. Evaluation on the Adult dataset showed that the approach achieved 91.63% accuracy. Authors in [12] improved ASD prediction by presenting a Convolutional Neural Network (CNN)-ET-XGB approach, which combined CNN for high-level feature extraction, ET for selecting relevant features, and XGB for classification. The CNN-ET-XGB approach was evaluated using a children dataset, where it achieved 99.992% accuracy. Authors in [13] reviewed ML methods for ASD detection in children, where Chi-square-based feature selection was used with classifiers such as SVM, KNN, DT, and MLP, reporting high classification accuracies. Despite high accuracies, existing approaches often face limitations in feature selection, with some relying on basic statistical methods like PCC or Chi-Square, which may overlook complex feature interactions, as well as inadequate handling of class imbalance, leading to biased predictions and reduced robustness.

Unlike previous approaches that either rely solely on ensemble methods or DL architectures, the proposed NeuroXGB-MLP introduces a hybrid learning approach where gradient-boosting-based feature optimization is tightly integrated with a neural network classifier. This combination enables enhanced feature interaction modeling while maintaining stable nonlinear decision learning across multiple age groups. The contributions of NeuroXGB-MLP are as follows:

- Introduces NeuroXGB-MLP, a hybrid framework combining XGB-based feature selection with an optimized MLP for ASD prediction across toddlers, children, and adults.
- Develops an ML-based ASD recommendation system with a structured intervention recommendation module to assist clinicians and caregivers in identifying individuals at risk of ASD and providing personalized recommendations.
- Implements robust class imbalance handling and Cross-Validation (CV) to enhance model generalizability and reduce overfitting.
- Demonstrates superior performance on multiple ASD datasets, outperforming traditional ML and ensemble approaches in prediction accuracy and recommendation reliability.

II. METHODOLOGY

This section presents a hybrid model, NeuroXGB-MLP, for developing an ML-based ASD prediction engine with a structured intervention clinical decision-support module to assist clinicians and caregivers in identifying individuals at risk of ASD and providing personalized recommendations. The NeuroXGB-MLP integrates feature optimization using XGB with an optimized MLP classifier for providing accurate ASD prediction and a designed clinical decision-support system across multiple age groups, which includes toddlers, children, and adults. This section first discusses the NeuroXGB-MLP architecture, then discusses the datasets, preprocessing steps considered, and train-test split. Further, this section discusses the XGB for feature selection, the optimized MLP classifier, and finally the CV process.

A. *Architecture*

The NeuroXGB-MLP architecture is presented in Figure 1.

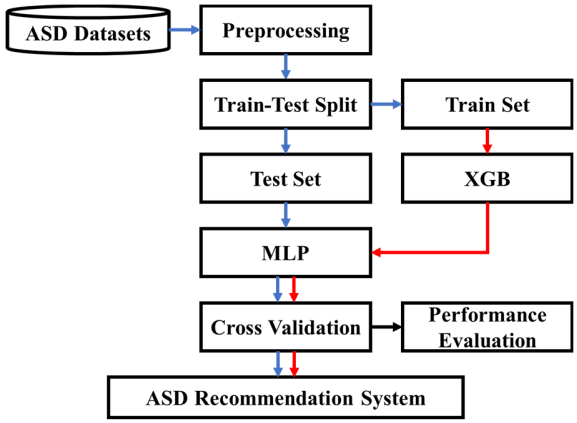

Fig. 1. NeuroXGB-MLP Architecture.

Initially, in NeuroXGB-MLP, multiple ASD datasets undergo preprocessing to remove inconsistencies and standardize feature representations. The datasets are then split into training and testing subsets using an 80:20 ratio. The training set is subsequently processed using XGB for feature selection, where the most relevant sensory-behavioral features are extracted. These selected features are then used as input to an optimized MLP classifier for prediction. CV is employed to ensure model robustness, whereas performance evaluation validates predictive performance. Based on NeuroXGB-MLP, an ASD clinical decision-support module is developed,

providing various recommendations and care plans for individuals with ASD.

B. Dataset, Preprocessing, and Train-Test Split

For this study, four ASD datasets were considered, i.e., Kaggle Toddler dataset [14], Saudi Toddler dataset [15], UCI Children dataset [16], and UCI Adult dataset [17]. The Kaggle Toddler dataset consisted of 19 features, the Saudi Toddler dataset consisted of 17 features, whereas the Children and Adult datasets consisted of 22 features each. For efficient training of NeuroXGB-MLP, the datasets were preprocessed.

During preprocessing, it was observed that the Kaggle Toddler dataset and Saudi Toddler dataset did not contain any missing values, whereas the Children and Adult datasets contained missing values in the "ethnicity" and "relation" columns. Hence, these columns were removed.

Further, during preprocessing, the "Case_No" column from the Toddler dataset and the "id", "age_desc", "country_of_res", and "used_app_before" columns from the Children and Adult datasets were removed. As the "ethnicity" feature was removed from the Children and Adult datasets, the corresponding "ethnicity" and "region" columns were removed from the Toddler and Saudi Toddler datasets, respectively, to ensure feature consistency across datasets.

In addition, the "Who completed test" column from the Toddler and Saudi Toddler datasets was removed due to the removal of the "relation" feature in the Children and Adult datasets. Furthermore, the "age" feature in the Toddler and Saudi Toddler datasets was originally provided in months and was converted into years for consistency. As the Saudi Toddler dataset did not contain the "Jaundice" feature, this feature was also removed from the Toddler, Children, and Adult datasets to maintain uniformity across all datasets.

In all datasets, the response scores for each questionnaire item were provided using A1–A10. After preprocessing, all four datasets contained the following features: "A1", "A2", "A3", "A4", "A5", "A6", "A7", "A8", "A9", "A10", "Age", "Sex", "Q-Chat-Score", and "Family member with ASD". The variable "Class/ASD Traits" was used strictly as the target label for supervised learning and was not included as an input feature. Furthermore, the "Q-Chat-Score" was removed, as it represents a cumulative behavioral screening score derived from A1–A10 responses and may introduce redundancy. Hence, in this study, 13 features were ultimately considered for ASD prediction.

For model training and evaluation, an 80:20 train-test split was adopted, consistent with commonly used practice in existing ASD prediction studies. The distribution of ASD and non-ASD cases, as well as gender-wise and total samples for each dataset, is presented in Table I.

TABLE I. ASD DATASETS CONSIDERED IN THIS STUDY

Dataset	ASD	Non-ASD	Male	Female	Total samples
Toddlers	728	326	735	319	1,054
Saudi Toddlers	341	165	157	349	506
Children	63	41	50	54	104
Adult	189	515	367	337	704

C. XGBoost-Based Feature Selection

Let the preprocessed dataset be represented as $\mathcal{D} = \{(x_i, y_i) | x_i \in \mathbb{R}^e, y_i \in \{0,1\}, i = 1,2,\dots,n\}$, where x_i denotes the feature vector for the i^{th} subject, which includes behavioral and sensory features, y_i denotes the binary label (0 for non-ASD and 1 for ASD), n denotes the total number of subjects, and e denotes the total number of features in the dataset. In this work, XGB aims to select the most informative features by constructing an ensemble of regression trees. The predicted output for sample x_i is given in (1):

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

In (1), f_k denotes the k^{th} Classification and Regression Tree (CART) in the ensemble, K denotes the total number of trees, $\mathcal{F}$ denotes the function space of possible tree sets. The objective of XGB combines loss l and regularization Ω, as presented in (2):

$$Obj = \sum_{i=1}^{n} l(y_i, \hat{y}_i) + \sum_{k=1}^{K} \Omega(f_k) \quad (2)$$

In (2), $l(\cdot,\cdot)$ denotes the loss function, and $\Omega(f_k)$ is defined as $\Omega(f_k) = \gamma T_k + \frac{1}{2}\lambda \sum_{j=1}^{T_k} w_j^2$, where T_k denotes the number of leaf nodes in the k^{th} tree, w_j denotes the j^{th} leaf node weight, γ penalizes tree complexity, and λ controls L2 regularization on leaf weights. Further, feature importance for splitting is evaluated using the gain function as presented in (3):

$$Gain = \frac{1}{2}\left[\frac{G_L^2}{H_L+\lambda} + \frac{G_R^2}{H_R+\lambda} - \frac{(G_L+G_R)^2}{H_L+H_R+\lambda}\right] - \gamma \quad (3)$$

In (3), G_R and G_L are the sums of first-order gradients for the right and left child nodes, whereas H_R and H_L are the sums of second-order gradients. The feature corresponding to the highest gain is selected for splitting. The resulting optimized feature subset $\tilde{X} \subset X_{train}$ is then used for MLP training.

D. Optimized Multi-Layer Perceptron Classifier

Let the input to the MLP be the selected features denoted as $\tilde{x}_i = [x_i^1, x_i^2, \dots, x_i^m]$, where $m \leq e$. The output of neuron j in the hidden layer is evaluated using (4) and (5):

$$S_j = \sum_{i=1}^{m} w_{ij}\tilde{x}_i + b_j \quad (4)$$

$$a_j = f(S_j) \quad (5)$$

In (4) and (5), w_{ij} represents the weights connecting input neuron i to hidden neuron j, b_j denotes the bias term, and $f(\cdot)$ denotes the activation function (Rectified Linear Unit (ReLU) for hidden layers, SoftMax for the output layer). For binary classification, the SoftMax output reduces to the sigmoid function, as presented in (6):

$$\hat{y}_i = \sigma(S_{out}) = \frac{1}{1+e^{-S_{out}}} \quad (6)$$

Further, the weights in the MLP are updated using gradient descent with momentum, as presented in (7) and (8):

$$\Delta w_{ij}(t) = -\epsilon \frac{\partial E}{\partial w_{ij}}(t) + \mu \Delta w_{ij}(t-1) \quad (7)$$

$$w_{ij}(t+1) = w_{ij}(t) + \Delta w_{ij}(t) \quad (8)$$

In (7) and (8), ϵ denotes the learning-rate, μ is momentum term, and E is the binary cross-entropy loss, which is evaluated using (9):

$$E = -\frac{1}{n}\sum_{i=1}^{n}[y_i \log(\hat{y}_i) + (1 - y_i)\log(1 - \hat{y}_i)] \quad (9)$$

This formulation ensures stable convergence while reducing prediction error.

E. Cross-Validation

In this work, during the training of NeuroXGB-MLP, CV was considered for efficient ASD prediction and improved clinical decision-support. To prevent data leakage, XGB-based feature selection was performed independently within each training fold during the 10-fold CV process. The selected feature subsets were then used to train the MLP within the same fold, ensuring unbiased evaluation. Further, the training loss was averaged across folds using (10):

$$\overline{E_{CV}} = \frac{1}{10}\sum_{k=1}^{10} E_k \quad (10)$$

During testing, unseen feature vectors $\tilde{X}_{test}$ were directly fed into the trained MLP using (11):

$$\hat{Y}_{test} = MLP(\tilde{X}_{test}) \quad (11)$$

The predicted probabilities were then used to provide personalized recommendations for caregivers and clinicians.

F. Autism Spectrum Disorder Clinical Decision-Support Module

The proposed framework integrates an ML-based ASD prediction engine with a structured intervention clinical decision-support module to assist clinicians and caregivers in identifying individuals at risk of ASD and providing personalized recommendations. In this work, a rule-based ASD intervention clinical decision-support module was developed and integrated with the NeuroXGB-MLP predictive engine. While the ASD prediction is generated using ML techniques, the intervention suggestions are derived from predefined clinical mapping rules based on dominant behavioral traits and predicted class outcomes.

Once the NeuroXGB-MLP model classifies a subject as ASD-positive or ASD-negative, the system generates suggestions based on predicted outcomes and key behavioral indicators identified during feature optimization. For ASD-positive predictions, the system recommends suitable interventions such as behavioral therapy, communication skill enhancement, and sensory integration activities, adapted to the subject's age group—toddler, child, or adult. For ASD-negative cases, it advises routine monitoring and follow-up screenings. The system continuously learns from new data through CV, ensuring adaptive performance and improved clinical decision-support recommendation accuracy over time.

Moreover, the personalization mechanism operates in two stages: feature importance scores derived from XGB identify dominant behavioral indicators (e.g., communication deficits, repetitive behaviors), and age-adaptive intervention mapping is performed based on age group (toddler, child, adult) and dominant behavioral characteristics. Using predefined clinical intervention rules, the system generates targeted clinical decision-support recommendations. This hybrid logic ensures that recommendations are individualized rather than generic. Hence, this hybrid recommendation approach enhances early detection, supports informed decision-making, and promotes targeted intervention strategies for individuals with ASD.

III. RESULTS AND DISCUSSION

The NeuroXGB-MLP and ASD clinical decision-support module were developed on a computational setup equipped with an Intel Core i7 processor, 16 GB RAM, and Intel Iris graphics. The model implementation and data analysis were carried out using the Python programming language within a Python-based development environment. The interactive ASD clinical decision-support module interface was subsequently developed using C#, enabling seamless integration of the predictive model with a user-friendly front-end.

For experimental consistency, the XGB model was configured with 100 trees, a learning rate of 0.1, a maximum tree depth of 6, and L2 regularization to reduce overfitting. The MLP classifier was implemented with two hidden layers consisting of 64 and 32 neurons respectively, the ReLU activation function, the Adam optimizer with a learning rate of 0.001, and a maximum of 200 training epochs.

These parameters were selected based on empirical tuning to ensure stable convergence and optimal predictive performance. This section first, discusses the performance metrics used for evaluating NeuroXGB-MLP, then discusses the performance achieved for different datasets, and finally discusses the clinical decision-support module.

A. Performance Metrics

For the performance evaluation of NeuroXGB-MLP, this work considered performance metrics similar to those used in recent works for ASD prediction. The performance metrics are given in (12)–(15), where TP denotes true positive, FP denotes false positive, TN denotes true negative, and FN denotes false negative.

$$\text{Accuracy } (A) = \frac{TP+TN}{TP+TN+FP+FN} \quad (12)$$

$$\text{Precision } (P) = \frac{TP}{TP+FP} \quad (13)$$

$$\text{Recall } (R) = \frac{TP}{TP+FN} \quad (14)$$

$$\text{F} - \text{score } (F) = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (15)$$

B. Performance Evaluation on Toddler Dataset

The evaluation results show that the proposed NeuroXGB-MLP framework achieves superior performance across all considered ASD datasets, including the Kaggle Toddler dataset [14], Saudi Toddler dataset [15], and their merged dataset. As shown in Figure 2, NeuroXGB-MLP achieved 99.995% accuracy, 99.993% precision, 99.990% recall, and 99.992% F-score on the Kaggle Toddler dataset.

Similarly, on the Saudi Toddler dataset (Figure 3), the NeuroXGB-MLP achieved 99.994% accuracy, 99.991% precision, 99.989% recall, and 99.990% F-score.

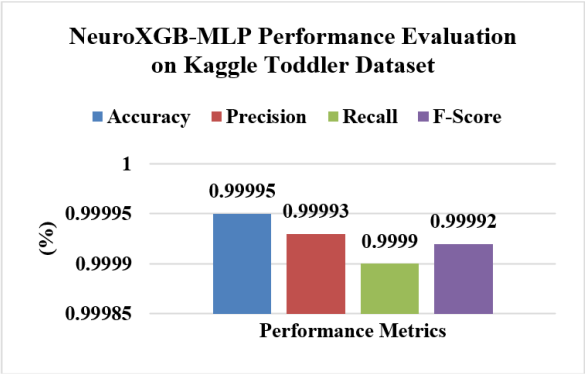

Fig. 2. Performance evaluation of NeuroXGB-MLP on Kaggle Toddler dataset.

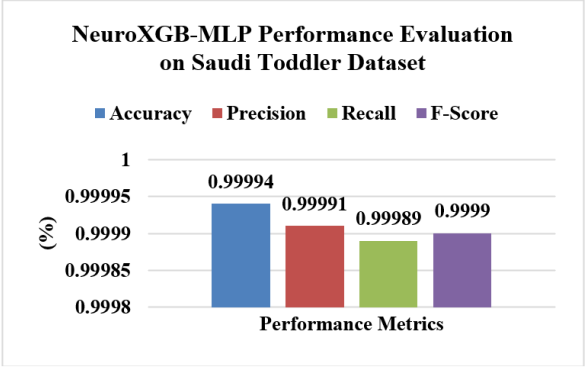

Fig. 3. Performance evaluation of NeuroXGB-MLP on Saudi Toddler dataset.

When both toddler datasets were merged, as presented in Figure 4, the NeuroXGB-MLP achieved slightly higher performance metrics, with 99.996% accuracy, 99.994% precision, 99.992% recall, and 99.993% F-score, indicating that NeuroXGB-MLP generalizes effectively across different datasets. The improved performance is attributed to the effective integration of XGB-based feature selection and the optimized MLP classifier. The feature selection process identifies the most informative sensory-behavioral features using (3), whereas the MLP optimization process in (6)–(9) ensures stable convergence during classification. By integrating these components, NeuroXGB-MLP effectively reduces irrelevant feature influence and mitigates class imbalance issues, resulting in highly accurate ASD prediction. The CV process further ensures robustness, confirming that NeuroXGB-MLP consistently outperforms previous approaches across different toddler datasets.

C. Performance Evaluation on Children Dataset

NeuroXGB-MLP achieves excellent performance on the UCI Children dataset [16], as illustrated in Figure 5, with 99.993% accuracy, 99.991% precision, 99.990% recall, and 99.990% F-score. The superior results are mainly due to the combined effect of XGB-based feature optimization and the optimized MLP classification process. The gain-based feature selection mechanism in (3) identifies the most relevant sensory-behavioral features, whereas the learning process defined in (6)–(9) enables stable convergence and effective classification. This combination reduces the influence of irrelevant features and handles class imbalance effectively, ensuring robust and highly accurate ASD prediction for children.

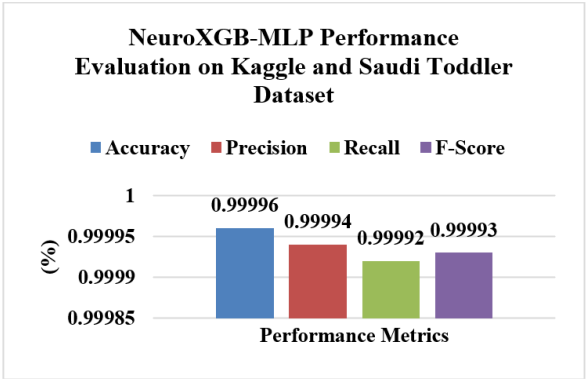

Fig. 4. Performance evaluation of NeuroXGB-MLP on merged toddler dataset.

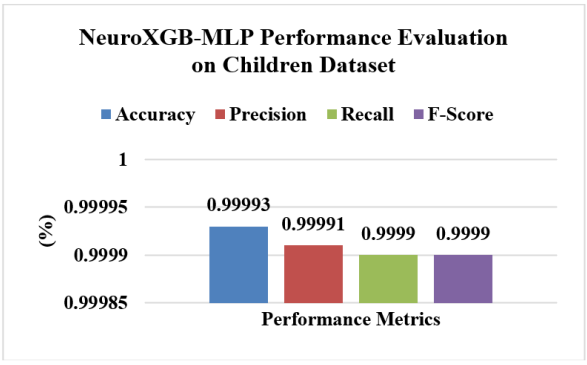

Fig. 5. Performance evaluation of NeuroXGB-MLP on Children dataset.

D. Performance Evaluation on Adult Dataset

The NeuroXGB-MLP shows outstanding performance on the UCI Adult dataset [17], as shown in Figure 6, achieving an 99.994% accuracy, 99.992% precision, 99.991% recall, and 99.992% F-score. The results are attributed to the combined effect of XGB-based feature selection in (3) and the optimized MLP classifier in (6)–(9), which together improve feature relevance and ensure stable convergence. This hybrid approach minimizes the influence of irrelevant features and addresses class imbalance, providing robust and highly accurate ASD prediction for adult subjects.

E. Clinical Decision-Support Recommendation Module

Figure 7 presents the graphical user interface of the developed ASD clinical decision-support prototype, designed to demonstrate the practical deployment of NeuroXGB-MLP. The interface is organized into three primary sections: an input panel for entering patient demographic information and behavioral screening responses (A1–A10), a prediction panel displaying the NeuroXGB-MLP classification outcome, and a recommendation panel providing structured intervention suggestions based on dominant behavioral indicators. The input section allows clinicians or caregivers to enter responses in a standardized format, ensuring consistency with the training

datasets. The prediction panel visualizes the NeuroXGB-MLP output in an interpretable manner, whereas the recommendation section translates the prediction results into clinically relevant, rule-based intervention pathways, thereby supporting practitioners in planning early therapeutic strategies. This interface demonstrates how NeuroXGB-MLP can be integrated into real-world screening workflows while maintaining usability and interpretability.

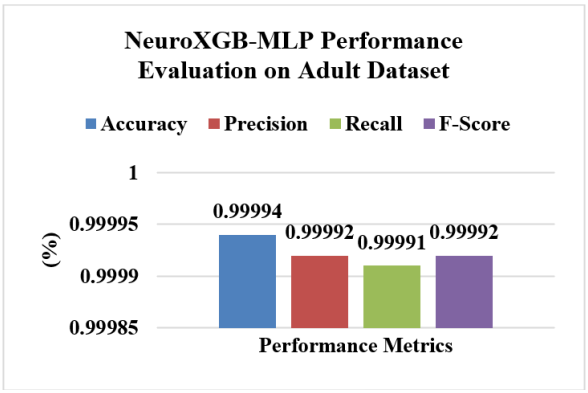

Fig. 6. Performance evaluation of NeuroXGB-MLP on Adult dataset.

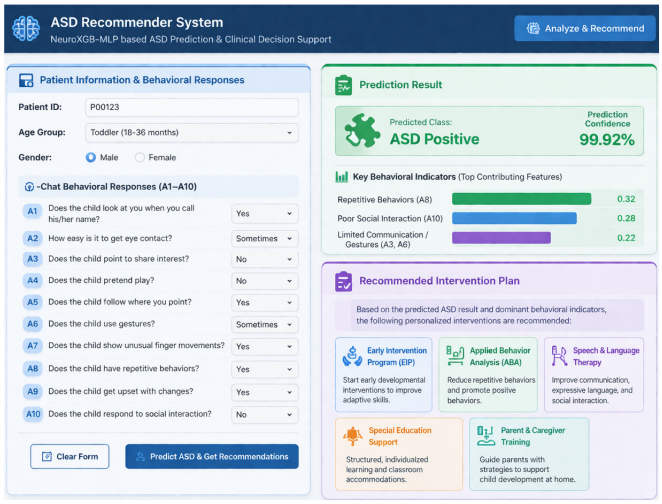

Fig. 7. Graphical user interface of the ASD clinical decision-support system showing the input form for behavioral responses, NeuroXGB-MLP prediction output, and rule-based intervention recommendations.

F. Comparative Study

The comparative study presented in Table II highlights performance of various ML approaches for ASD prediction across different datasets. Traditional models such as RF, DT, LR, KNN, SVM, GBM, and XGB achieved moderate to high accuracies on the merged Kaggle Toddler [14] and Saudi Toddler [15] datasets, ranging from 93% for LR to 98% for RF. Ensemble methods like RF-XGB combined with ML-based Feature Selection (MFS) further improved prediction performance, achieving 99% accuracy, indicating the advantage of combining complementary classifiers and refined feature selection. In [9], PCA-based dimensionality reduction followed by a stacked ensemble of RF, ET, CB, and ANN with Safe-Level SMOTE addressed class imbalance, achieving near-perfect results on the Kaggle toddler [14], Children [16], and Adult [17] datasets. The MAPLE framework [11] focused on imputing missing values and outlier removal for adult ASD prediction, yielding lower accuracy (91.63%), reflecting the challenges of noisy and incomplete data. DL approaches such as CNN-ET-XGB [12] leveraged high-level feature extraction with ensemble-based feature selection, achieving 99.992% accuracy on the Children dataset [16], demonstrating the effectiveness of deep architectures for feature representation.

TABLE II. COMPARATIVE STUDY

Ref.	Dataset	Model	*A*	*P*	*R*	*F*
[7]	[14]+[15]	RF	0.98	0.98	0.97	0.98
		DT	0.96	0.96	0.95	0.95
		LR	0.93	0.97	0.95	0.96
		KNN	0.95	0.96	0.96	0.96
		SVM	0.94	0.95	0.97	0.96
		GBM	0.94	0.94	0.97	0.95
		XGB	0.97	0.98	0.98	0.98
		RF-XGB+MFS	0.99	0.99	0.99	0.99
[9]	[14]	PCA+RF+ET+ CB+ANN+Safe -Level SMOTE	0.9986	0.9992	0.9998	0.9992
	[16]		0.9968	0.9999	0.9828	0.9927
	[17]		0.9989	0.9998	0.9997	0.9998
[11]	[17]	MAPLE	0.916312	0.9176	0.9169	0.9163
[12]	[16]	CNN-ET-XGB	0.99992	0.99986	0.99985	0.99986
Proposed	[14]	NeuroXGB-MLP	0.99995	0.99993	0.9999	0.99992
	[15]		0.99994	0.99991	0.99989	0.9999
	[14]+[15]		0.99996	0.99994	0.99992	0.99993
	[16]		0.99993	0.99991	0.9999	0.9999
	[17]		0.99994	0.99992	0.99991	0.99992

The proposed NeuroXGB-MLP consistently outperforms all previous methods across all datasets. For instance, it achieves 99.995% accuracy on the Kaggle Toddler dataset [14], 99.994% on the Saudi Toddler dataset [15], and 99.996 on the merged toddler dataset, with similarly high performance on the Children [16] and Adult [17] datasets. The improved performance is due to XGB-based feature selection in (3), which identifies the most informative sensory-behavioral features, coupled with the optimized MLP classifier in (6)–(9), which employs gradient descent with momentum for stable convergence. This combination reduces irrelevant feature influence, addresses class imbalance, and ensures robust learning. CV further enhances generalization, enabling NeuroXGB-MLP to achieve superior accuracy, precision, recall, and F-score across diverse ASD datasets, establishing it as a reliable tool for early ASD detection and personalized clinical decision-support. Although the obtained accuracy values are exceptionally high, this is attributed to the structured and standardized nature of ASD screening datasets, where behavioral responses exhibit strong separability between ASD and non-ASD classes. Similar near-perfect performance has been reported in prior literature on the same benchmark datasets. The evaluation protocol strictly followed fold-wise feature selection and CV to prevent data leakage.

IV. CONCLUSION

Early detection of Autism Spectrum Disorder (ASD) is crucial for timely intervention, but traditional manual screening is time-consuming and prone to subjective bias. Existing

Machine Learning (ML) approaches, although effective, face limitations in handling class imbalance, selecting the most informative features, and generalizing across diverse age groups.

This study addressed these challenges by proposing NeuroXGB-MLP, a hybrid framework that combines XGBoost-based feature selection with an optimized Multi-Layer Perceptron (MLP) classifier, ensuring robust and accurate prediction. The methodology involved preprocessing four ASD datasets, i.e., Kaggle Toddler, Saudi Toddler, UCI Children, and UCI Adult datasets, standardizing features, selecting the most relevant sensory-behavioral inputs, and training NeuroXGB-MLP with Cross-Validation (CV) to improve generalization.

Experimental results demonstrate that NeuroXGB-MLP outperforms existing models, achieving 99.995% accuracy on the Kaggle Toddler dataset, 99.994% on the Saudi Toddler dataset, 99.996% on the merged toddler dataset, 99.993% on the Children dataset, and 99.994% on the Adult dataset, with similarly high precision, recall, and F-score values. These results highlight the effectiveness of the proposed feature selection and optimized MLP in mitigating class imbalance and improving prediction reliability.

Although the proposed framework demonstrates strong performance on benchmark datasets, reliance on publicly available datasets may limit ecological validity. Real-world clinical data may contain additional contextual variability not fully captured in standardized screening datasets. Hence, future research will focus on collaborating with NGOs and clinical institutions working with individuals on the autism spectrum to validate the NeuroXGB-MLP framework using real-world clinical datasets. Such validation will facilitate the assessment of context-specific behavioral patterns, improve ecological validity, and enhance the model's generalizability across diverse demographic populations. Furthermore, future extensions may explore the integration of multimodal behavioral indicators, including eye-gaze analysis and visual attention patterns, to complement questionnaire-based screening and further strengthen ASD prediction robustness.

DECLARATION OF COMPETING INTERESTS

The authors declare that they have no known competing financial interests, personal relationships, or conflicts of interest that could have appeared to influence the work reported in this paper.

ACKNOWLEDGMENT

The authors acknowledge the publicly available ASD datasets provided by Kaggle and the UCI Machine Learning Repository, which were used in this study.

DATA AVAILABILITY

The datasets used in this study are publicly available and can be accessed from the following sources: Kaggle Toddler dataset [14], Saudi Toddler dataset [15], UCI Children dataset [16], and UCI Adult dataset [17].

AI USE AND DECLARATION OF GENERATIVE AI USE

During the preparation of this work, the authors used Grammarly AI for language refinement, grammar correction, and improving the readability of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

REFERENCES

[1] M. Al-Beltagi, "Pre-autism: Advancing early identification and intervention in autism," *World Journal of Clinical Cases*, vol. 12, no. 34, pp. 6748–6753, Dec. 2024, https://doi.org/10.12998/wjcc.v12.i34.6748.

[2] K. M. Wilczyński *et al.*, "An Analysis of the Time Required for the Diagnosis of ASD and the Factors Influencing Its Duration in a Sample of the Pediatric Population from Poland," *Journal of Clinical Medicine*, vol. 13, no. 20, Oct. 2024, Art. no. 6255, https://doi.org/10.3390/jcm13206255.

[3] D. Y. Kim *et al.*, "Automated AI based identification of autism spectrum disorder from home videos," *npj Digital Medicine*, vol. 8, no. 1, Oct. 2025, Art. no. 607, https://doi.org/10.1038/s41746-025-01993-5.

[4] H. A. Hatim, Z. A. A. Alyasseri, and N. Jamil, "A recent advances on autism spectrum disorders in diagnosing based on machine learning and deep learning," *Artificial Intelligence Review*, vol. 58, no. 10, July 2025, Art. no. 313, https://doi.org/10.1007/s10462-025-11302-x.

[5] U. J. Ganai, A. Ratne, B. Bhushan, and K. S. Venkatesh, "Early detection of autism spectrum disorder: gait deviations and machine learning," *Scientific Reports*, vol. 15, no. 1, Jan. 2025, Art. no. 873, https://doi.org/10.1038/s41598-025-85348-w.

[6] G. Gatica-Bahamonde *et al.*, "Testing an online screening for autism in the COVID-19 pandemic: a psychometric study of the Q-CHAT-24 in Chilean toddlers," *Frontiers in Psychiatry*, vol. 15, June 2024, Art. no. 1363976, https://doi.org/10.3389/fpsyt.2024.1363976.

[7] F. Hajjej, S. Ayouni, M. A. Alohali, and M. Maddeh, "Novel Framework for Autism Spectrum Disorder Identification and Tailored Education With Effective Data Mining and Ensemble Learning Techniques," *IEEE Access*, vol. 12, pp. 35448–35461, 2024, https://doi.org/10.1109/ACCESS.2024.3349988.

[8] U. Jabbar *et al.*, "Machine Learning–Based Approach for Early Screening of Autism Spectrum Disorders," *Applied Computational Intelligence and Soft Computing*, vol. 2025, no. 1, 2025, Art. no. 9975499, https://doi.org/10.1155/acis/9975499.

[9] N. Mumenin, M. M. Rahman, M. A. Yousuf, F. M. Noori, and M. Z. Uddin, "Early diagnosis of autism across developmental stages through scalable and interpretable ensemble model," *Frontiers in Artificial Intelligence*, vol. 8, May 2025, Art. no. 1507922, https://doi.org/10.3389/frai.2025.1507922.

[10] N. Haque, T. Islam, and M. Erfan, "An exploration of machine learning approaches for early Autism Spectrum Disorder detection," *Healthcare Analytics*, vol. 7, June 2025, Art. no. 100379, https://doi.org/10.1016/j.health.2024.100379.

[11] F. J. Rajam and B. R. K. Swakkin, "MAPLE: A Novel Processing Technique for Adult Autism Prediction," *Engineering, Technology & Applied Science Research*, vol. 15, no. 3, pp. 23901–23906, June 2025, https://doi.org/10.48084/etasr.10864.

[12] K. Arpitha and N. Nandini, "A Novel Feature Optimization Approach for Accurate Autism Spectrum Disorder Prediction," *Engineering, Technology & Applied Science Research*, vol. 15, no. 5, pp. 27159–27164, Oct. 2025, https://doi.org/10.48084/etasr.12319.

[13] M. M. Rahman, O. L. Usman, R. C. Muniyandi, S. Sahran, S. Mohamed, and R. A. Razak, "A Review of Machine Learning Methods of Feature Selection and Classification for Autism Spectrum Disorder," *Brain Sciences*, vol. 10, no. 12, Dec. 2020, Art. no. 949, https://doi.org/10.3390/brainsci10120949.

[14] "Autism screening data for toddlers." Kaggle. [Online]. Available: https://www.kaggle.com/datasets/fabdelja/autism-screening-for-toddlers.

[15] "ASD Screening Data for Toddlers in Saudi Arabia." Kaggle. [Online]. Available: https://www.kaggle.com/datasets/asdpredictioninsaudi/asd-screening-data-for-toddlers-in-saudi-arabia.

[16] F. Thabtah, "Autistic Spectrum Disorder Screening Data for Children." UCI Machine Learning Repository, 2017, https://doi.org/10.24432/C5659W.

[17] F. Thabtah, "Autism Screening Adult." UCI Machine Learning Repository, 2017, https://doi.org/10.24432/C5F019.