

Autism Spectrum Disorder (ASD) Risk Prediction Using Adaptive Risk Fusion Network with Deep Neural Network (ARFN-DNN)

Selva Kavın S¹, Dr.Thara L², Dr.Nithyanandh Selvam³

¹II-MCA,

²Associate Professor and Head of Department,

³Associate Professor

DOI: <https://doi.org/10.51583/IJLTEMAS.2026.150900103>

Received: 03 October 2026; Accepted: 08 October 2026; Published: 19 October 2026

ABSTRACT

There is a need for computational screening techniques to supplement traditional assessment since delayed diagnosis of Autism Spectrum Disorder (ASD) can limit the possibility for prompt developmental support. This study suggests an Adaptive likelihood Fusion Network with Deep Neural Network (ARFN-DNN) that uses clinical-profile features and structured screening to predict the likelihood of ASD in children. Through probability-level fusion, the suggested architecture integrates Random Forest (RF), XGBoost (XGB), Logistic Regression (LR), and a regularized Deep Neural Network (DNN).. In order to illustrate behavioral, developmental, clinical, social-communication, and behavioral-communication link the dataset is printed six domain-inspired risk features. For the neural meta-learner, unbiased base-model predictions were produced using a five-fold stratified out-of-fold process. The suggested ARFN-DNN obtained 91.88% accuracy, 91.67% precision, 89.53% recall, 90.59% F1-score, 93.69% specificity, and 0.9424 ROC-AUC on the frozen unseen test set with 197 samples. With a train-test accuracy gap of just 1.89%, the model's generalization was comparatively steady. The results imply that while formal clinical evaluation and external validation are still necessary, probability-level fusion can offer a helpful additional screening layer.

Keywords: Autism Spectrum Disorder, ASD Screening, Machine Learning, Deep Neural Network, Random Forest, XGBoost, Logistic Regression, SMOTE, Ensemble Learning, Risk Prediction, ARFN-DNN.

INTRODUCTION

The neurodevelopmental disorder known as autism spectrum disorder (ASD) can impact behavior, social interaction, communication, and developmental patterns. Because prompt screening and intervention can promote better developmental outcomes, early identification of ASD risk is crucial. However, traditional screening methods could rely on a small number of behavioral or developmental characteristics and might not successfully integrate data from other domains. Recent developments in machine learning offer the chance to examine several screening characteristics and spot trends linked to the risk of ASD.

For the purpose of predicting the risk of ASD, this study suggests an Adaptive Risk Fusion Network with Deep Neural Network (ARFN-DNN). The suggested method integrates the predictions of a regularized Deep Neural Network, Random Forest, XGBoost, and Logistic Regression. A final ASD risk prediction is produced by combining the probability outputs of these models' analysis of the available screening data from various learning perspectives via a fusion method. Data preparation, domain-based feature engineering, training using the training dataset, five-fold cross-validation with out-of-fold predictions, and final evaluation on a different test dataset are all included in the study.

This study focuses on combining behavioral, developmental, and associated screening variables to screen for ASD risk using machine learning. Instead than taking the place of expert clinical diagnosis, the method is

meant to offer a risk assessment that can assist initial screening. The goal is to create and assess an integrated ARFN-DNN model that maintains strong predictive performance and generalization on untested data while combining several machine-learning predictions to provide dependable ASD risk classification.

LITERATURE REVIEW

Computational Approaches for ASD Screening

Because behavioral and developmental data contain patterns that can be challenging to reliably assess through manual examination alone, machine learning has been studied more and more for ASD screening. Classical machine learning, deep learning, behavioral analysis, neuroimaging, eye tracking, wearable sensors, and other data sources have all been discussed in reviews of ASD screening methods [3], [7]. These studies show that computational methods can find correlations between observable traits and outcomes associated to ASD; nevertheless, a model's effectiveness is heavily influenced by the population, feature representation, validation strategy, and dataset. The variety of AI techniques employed for early ASD identification has also been emphasized in recent systematic studies. Researchers have used behavioural observations, demographic factors, fMRI, EEG, eye-tracking data, voice, and multimodal data to represent classical machine learning, deep learning, and hybrid approaches in the literature.

Nevertheless, clinical reliability cannot be established solely by high reported accuracy. Confidence in reported generalization might be lowered by small sample sizes, frequent use of certain benchmark datasets, a lack of external validation, and methodological variations among studies [4].

Conventional Machine Learning for ASD Prediction

Because of its comparatively straightforward decision structure and regularization capabilities, logistic regression is widely utilized as a baseline classification technique. It is a helpful benchmark for assessing whether more intricate algorithms yield appreciable gains over a linear decision boundary [8]. By building several decision trees and aggregating their forecasts, Random Forest expands the modeling potential. Breiman showed that the strength of individual trees and the correlation between them determine Random Forest's generalization behavior [5]. A gradient-boosting framework called XGBoost was created to effectively build an ensemble of decision trees. XGBoost, a scalable tree boosting method featuring regularization and computational techniques intended for large-scale learning, was first presented by Chen and Guestrin [6].

Boosting can simulate nonlinear interactions and correlations in structured tabular situations that a strictly linear classifier might not be able to fully capture. Complementary qualities are offered by these models. Random Forest uses randomized tree ensembles to represent nonlinear feature interactions, XGBoost conducts sequential error-focused boosting, and Logistic Regression provides a relatively straightforward baseline. Instead of depending on a single representation of the data, integrating these heterogeneous models allows for the option of merging several decision processes.

Time-Series Based Waiting-Time Prediction

Additionally, deep learning has been investigated for the identification of ASD in several modalities. Neural architectures have been used for behavioral, neuroimaging, EEG, eye-tracking, and other ASD-related data, according to recent reviews [4], [9]. A DNN has a different modeling capability than linear and tree-based classifiers since it can depict nonlinear correlations between several input variables.

However, there is a significant risk of overfitting when using neural networks on very tiny structured datasets. Instead of learning patterns that apply to unknown people, a model might pick up traits unique to the training data. This risk can be decreased by regularization strategies such as weight decay, batch normalization, dropout, and early termination. Because AdamW isolates weight decay from the gradient-based optimization process, it is especially pertinent [10].

Class Imbalance and Data Leakage

Another significant problem with medical and screening datasets is class imbalance. A classifier may discover a decision boundary that benefits the majority class if one class is underrepresented. Instead of just replicating preexisting cases, SMOTE creates synthetic minority-class observations to solve this issue [11]. SMOTE is used within the cross-validation training folds in this study and is limited to training data. The test data that hasn't been seen is left intact.

Data leakage is just as significant. The evaluation may become unduly optimistic if data from the test set affects feature selection, scaling, oversampling, or meta-learner creation. As a result, the suggested workflow only uses training data to fit feature-selection and scaling processes, and the meta-learner is trained using out-of-fold predictions.

Research Gap

Although there are still a number of methodological issues, the literature now in publication shows the potential of deep learning and machine learning for ASD screening. First, a lot of research concentrates on a specific data modality or individual classifiers. Second, learnt probability-level fusion and simple voting are not always distinguished by ensemble techniques. Third, stated performance may be challenging to interpret due to methodological issues with class imbalance, information leakage, and validation design. Small sample numbers and a lack of external validation are particularly noted in recent studies as ongoing issues for ASD AI research [4].

Therefore, the creation of a structured hybrid framework that explicitly controls the training, validation, and test borders while combining heterogeneous classifiers with a regularized DNN and neural probability-level fusion fills the research need. The employment of numerous algorithms is not the only contribution; an out-of-fold meta-learning mechanism and a weighted fusion component are used to integrate the probability outputs of these algorithms.

Comparison Of The Proposed Approach With Existing Methodological Directions

Table. 2.6.1. Comparison of proposed work with existing work


Methodological Direction	Main Strength	Limitation Addressed by ARFN-DNN
Single Classical Classifier[8]	Simple and computationally efficient	May capture only one type of decision structure
Tree Ensemble[5]	Captures nonlinear relationships	Does not directly learn complementary probability streams
Boosting[6]	Strong nonlinear predictive capability	Sensitive to training configuration and data characteristics
Standalone DNN[4],[9]	Learns complex nonlinear relationships	Can overfit small tabular datasets
Simple Ensemble Voting[3],[7]	Combines predictions from several models	Fixed voting may not learn relationships between model probabilities
Proposed ARFN-DNN	Combines heterogeneous probability outputs using a neural meta-learner and weighted fusion	Designed to integrate complementary model behavior while reducing reliance on a single decision structure through OOF predictions

RESEARCH DESIGN AND METHODOLOGY

The study employs an experimental design for supervised machine learning. Data cleaning, categorical encoding, stratified train-test partitioning, domain feature engineering, mutual-information analysis, normalization, class balancing, base-model training, out-of-fold prediction generation, neural meta-learning,

probability-level fusion, and final evaluation are all steps in the overall methodology, which starts with the pediatric ASD screening dataset.

Overall Research Workflow



Proposed Architecture

In order to stop the model from learning identity-specific artifacts, the preprocessing pipeline started with data sanitization, which involved eliminating duplicate entries and record identifiers. The target variable was translated to binary labels (0 for No ASD, 1 for ASD) after categorical variables were transformed into binary numerical encodings. In order to preserve class distributions across both subsets and isolate the test partition as a fully frozen evaluation benchmark, the cleaned cohort of 984 records was divided into 787 training instances and 197 unseen test examples using a stratified 80:20 train–test split.

Six domain-guided composite features were created from the original attributes in order to enhance predictive capacity: Behavioral Risk, which combines communication, social interaction, and repetitive behaviors; Developmental Risk, which includes language delay, familial ASD history, and jaundice; Clinical Risk, which standardizes developmental metrics by age; Comprehensive Risk Score, which aggregates core diagnostic dimensions; and two interaction terms, Social Communication Risk and Behavior Communication Risk. These designed indices capture intricate clinical interdependencies that are directly related to screening outcomes, expanding the dataset to 14 total variables.

Following the evaluation of feature relevance using Mutual Information (MI) on the training set, the top target-dependent predictors were found to be Behavioral Risk (0.4046), Social Communication Risk (0.3669), and the Comprehensive Risk Score (0.3661). Using the MI analysis solely for interpretability and ranking rather than aggressive dimensionality reduction, all 14 features were kept in the final pipeline. In order to ensure that

there was no data leakage throughout the model assessment process, a standard scaler was fitted solely on the training distribution and then applied downstream to the test set.

Data Collection

The file `asd.csv`, which is described in the technical package as a synthesized pediatric ASD screening clinical-profile dataset that matches the structure of an AQ-10 child screening benchmark, served as the experimental dataset for this investigation. Eight predictive feature columns, an identification column, and 1,000 entries make up the original dataset. There are 984 records left after duplicates are eliminated. After preprocessing, there are no missing values in the dataset.

Table. 4.1.1. Dataset Summary Table

Dataset	Samples	Negative	Positive
Original	1,000	Not applicable	Not applicable
Cleaned	984	552	432
Training	787	441	346
Unseen Test	197	111	86

Diagnosed_ASD is the target variable, with 0 denoting no ASD and 1 denoting ASD. There are 432 positive samples and 552 negative samples in the cleaned dataset. The obtained class proportions are roughly 43.90% positive and 56.10% negative. There are 441 negative and 346 positive observations in the training partition and 111 negative and 86 positive observations in the frozen test partition.

The results should be taken as an experimental assessment of the suggested machine learning framework rather than as clinical evidence since the dataset is described as synthesized rather than independently gathered clinical data. When evaluating whether the reported performance may be applied to different pediatric populations, this distinction is crucial.

Data Analysis

Base Model Development

The suggested framework makes use of four base classifiers. $C=0.8$ and a maximum of 1,000 iterations are the settings for logistic regression. With 250 trees, a maximum depth of five, a minimum split size of six, a minimum leaf size of three, and square-root feature selection, Random Forest is set up. 120 estimators, a maximum depth of three, a learning rate of 0.025, subsampling and column subsampling of 0.8, L1 regularization of 0.1, L2 regularization of 1.5, and a minimum child weight of three are all used by XGBoost.

The 14 processed features are sent to the DNN. Batch normalization, ReLU activation, and a dropout of 0.35 come after the first fully connected layer, which maps 14 inputs to 32 neurons. The second layer uses batch normalization, ReLU, and dropout once more while mapping 32 neurons to 16 neurons. One sigmoid probability is generated by the last layer. AdamW is used to train the network using a weight decay of 0.001 and a learning rate of 0.002. The loss function is binary cross-entropy, and to reduce needless training, early termination based on validation loss is employed.

Training And Testing Procedure

Model development and final evaluation are kept apart throughout the entire experimental process. The dataset is first cleaned and split into subsets for training and untested. The training part is then used for all model-related actions. Five-fold stratified cross-validation separates the observations into training and validation subsets within the training data.

SMOTE is only used for the fold-training part of each fold. This balanced fold-training data is used to train the base classifiers, and predictions are made for the unaltered validation portion. Out-of-fold probabilities are used to hold the final forecasts. Each training observation has a probability prediction produced by a model that did not train on it once all five folds have been finished. As a result, an OOF matrix with four probability streams—Logistic Regression, Random Forest, XGBoost, and DNN is created. The neural meta-learner is then given the OOF matrix. The meta-learner cannot simply learn predictions made by models that were directly trained on the same observations because to its design. For the fusion stage, the process offers a more suitable training signal.

The frozen test set is run through the learned preprocessing pipeline and base classifiers following the development of the base learners and meta-learner. The meta-learner is given four test probabilities. The weighted fusion component also makes advantage of these probabilities. The meta-learner probability and the weighted fusion probability are combined to get the final ARFN-DNN probability.

Cross Validation

The training process's stability is evaluated using five-fold stratified cross-validation. Logistic regression has an average validation accuracy of 90.72% and a standard deviation of 0.56%. XGBoost reaches 91.87% with a standard deviation of 0.81%, but Random Forest obtains 92.25% with a standard deviation of 0.82%. With a standard deviation of 0.61%, the DNN achieves 90.73%. The mean accuracy of the ARFN-DNN fusion is 92.88%, with a standard deviation of 0.51%.

Table. 5.3.1.Five Fold Cross Validation Performance

Model	Mean Accuracy	Standard Deviation
Logistic Regression	90.72%	±0.56%
Random Forest	92.25%	±0.82%
XGBoost	91.87%	±0.81%
DNN	90.73%	±0.61%
ARFN-DNN Fusion	92.88%	±0.51%

The ARFN-DNN configuration's comparatively low standard deviation shows consistent behavior over the five validation folds. However, as all folds come from the same dataset, cross-validation performance shouldn't be taken as independent external validation.

Unseen Test Set Performance

The frozen test set of 197 samples is used for the final assessment. 89.85% accuracy is attained by Logistic Regression, 91.37% by Random Forest and XGBoost, 91.88% by DNN, and 91.88% by ARFN-DNN. ARFN-DNN generates a different performance profile even though it does not outperform DNN in test accuracy. In contrast to the DNN's 90.70% precision and 0.9070 F1-score, ARFN-DNN obtains 91.67% precision and an F1-score of 0.9059. While ARFN-DNN records 0.8953 recall, the DNN earns a little higher recall of 0.9070. This suggests that the suggested fusion mechanism is not always better than all base models. Its contribution is better understood as offering a balanced and competitive fusion technique that maintains the same overall test accuracy while increasing precision in comparison to the DNN.

Fig 5.4.1. Final Test Performance

Model	Accuracy	Precision	Recall	F1-Score	Specificity	ROC-AUC
Logistic Regression	89.85%	0.9125	0.8488	0.8795	0.9369	0.9387
Random Forest	91.37%	0.9157	0.8837	0.8994	0.9369	0.9471

XGBoost	91.37%	0.9059	0.8953	0.9006	0.9279	0.9443
DNN	91.88%	0.9070	0.9070	0.9070	0.9279	0.9453
ARFN-DNN	91.88%	0.9167	0.8953	0.9059	0.9369	0.9424

The suggested system performs competitively across all reported parameters, according to the test findings. A comparatively minor generalization gap is indicated by the 1.89% difference between training accuracy of 93.77% and test accuracy of 91.88%. The outcome is in line with the DNN component's application of regularization and validation procedures.

Confusion Matrix ANALYSIS

There are 104 true negatives, seven false positives, nine false negatives, and 77 real positives in the ARFN-DNN confusion matrix. As a result, 181 of the 197 test observations are properly classified by the model, yielding an accuracy of 91.88%.

Table 5.5.1. ARFN-DNN Confusion Matrix

Actual / Predicted	Predicted No ASD	Predicted ASD
Actual No ASD	104 (TN)	7 (FP)
Actual ASD	9 (FN)	77 (TP)
Total	113	84

The model successfully identifies a significant percentage of negative test observations, as indicated by the specificity of 93.69%. The model detects the majority of positive data, but it still generates false-negative cases, as evidenced by the recall of 89.53%. False negatives are crucial in the context of ASD screening since a missed risk indication could postpone additional evaluation. Therefore, rather than depending just on accuracy, future optimization should take into account the clinical cost associated with various error kinds.

Feature Analysis

An extra perspective of the data in the engineered feature space is offered by mutual information analysis. The highest MI score is 0.4046 for behavioral risk, 0.3669 for social communication risk, and 0.3661 for comprehensive risk. These findings suggest that significant information related to the target in the training dataset can be found in aggregated behavioral and social-communication representations.

Because Behavioral_Risk directly integrates scores for social contact, communication, and repeated behavior, its high ranking is compatible with the input data's structure. Instead of treating two behavioral characteristics separately, Social Communication Risk reflects their relationship. Several risk characteristics are further integrated into a single representation by Comprehensive Risk Score. It is not appropriate to interpret the results as demonstrating that these variables represent diagnostic biomarkers or clinical causes. This dataset's statistical dependency is described by the MI scores, which should not be construed causally.

Ablation Analysis

The ablation experiment assesses if the suggested framework's elements influence its behavior. Test accuracy is 89.85% for Logistic Regression alone, 91.37% for Random Forest and XGBoost, and 91.88% for DNN. Weighted risk fusion by itself likewise gets 91.37%, as does a straightforward equal-weighted ensemble. 91.88% is attained by the whole ARFN-DNN model.

Fig 5.7.1. Ablation Study

Configuration	Accuracy	Precision	Recall	F1-Score	ROC-AUC
Logistic Regression	89.85%	0.9125	0.8488	0.8795	0.9387
Random Forest	91.37%	0.9157	0.8837	0.8994	0.9471
XGBoost	91.37%	0.9059	0.8953	0.9006	0.9443
DNN	91.88%	0.9070	0.9070	0.9070	0.9453
Equal-Weighted Ensemble	91.37%	0.9157	0.8837	0.8994	0.9431
Weighted Risk Fusion	91.37%	0.9059	0.8953	0.9006	0.9438
ARFN-DNN	91.88%	0.9167	0.8953	0.9059	0.9424

The ablation results show that an improvement in accuracy over the DNN is not the primary advantage of ARFN-DNN. Rather, despite preserving the best accuracy seen in the experiment, the meta-learning stage modifies the ratio of precision, recall, and F1-score. Therefore, rather than guaranteeing a numerical improvement across all metrics, the complete design offers a learning fusion technique.

FINDINGS AND DISCUSSION

The empirical results show that a strong framework for modeling structured pediatric screening data may be obtained by combining deep learning with traditional machine learning architectures. The suggested ARFN-DNN model successfully utilizes the complementary inductive biases of its base classifiers—linear calibration from logistic regression, nonlinear recursive partitioning from Random Forest, gradient-boosted decision boundaries from XGBoost, and dense feature representations from the baseline DNN—to achieve an overall test accuracy of 91.88%. This multi-model design is justified over a single classifier by subtle trade-offs revealed by multi-metric evaluation. Random Forest exhibits better threshold-independent discrimination with a peak ROC-AUC of 0.9471, while ARFN-DNN leads in precision (91.67%) with a competitive F1-score of 0.9059. The standalone DNN obtains slightly greater recall (90.70%).

The suggested system reduces the vulnerabilities of individual classifiers to provide a balanced, variance-reduced probability estimate by equally combining a trained neural meta-learner with a calibrated risk-weighted ensemble (assigning weights of 0.30 to XGBoost, 0.25 to Random Forest and DNN, and 0.20 to Logistic Regression). The system's potential as a decision-support tool is further supported by statistical stability and clinical efficacy. The training pipeline produced a mean cross-validation accuracy of 92.88% with little variance ($\sigma = 0.51\%$), indicating strong stability across folds. Additionally, the regularization strategy—which includes dropout, batch normalization, AdamW weight decay, and training-isolated SMOTE—effectively inhibited empirical memorization, as demonstrated by a narrow generalization gap of 1.89 percentage points between training (93.77%) and test (91.88%) accuracy. The incorporation of domain-specific feature engineering is validated by feature relevance analysis via mutual information, which finds composite clinical markers, particularly Behavioral Risk, Social Communication Risk, and the Comprehensive Risk Score, as the main drivers of model predictions. Importantly, manual case evaluations verify that the model generates continuous, monotonic risk probabilities instead of strict binary labels, allowing clinicians to more precisely stratify patient risk. Nevertheless, independent multisite clinical validation is still a necessary precondition for establishing true population-level generalization.

CONCLUSION

By combining Logistic Regression, Random Forest, XGBoost, and a regularized DNN via probability-level fusion, the study created an Adaptive Risk Fusion Network with Deep Neural Network for structured ASD risk prediction. The suggested methodology integrates a neural meta-learner, training-only normalization, SMOTE-

based class balancing, stratified five-fold out-of-fold prediction generation, and domain-inspired feature engineering. ARFN-DNN attained 91.88% accuracy, 91.67% precision, 89.53% recall, 90.59% F1-score, 93.69% specificity, and 0.9424 ROC-AUC on the frozen unseen test set of 197 observations. Relatively steady generalization within the assessed dataset is also indicated by the 1.89 percentage-point train-test gap.

The findings indicate that rather than merely improving accuracy beyond each individual model, the main benefit of ARFN-DNN is in integrating complementary probability estimates. Random Forest generated the best ROC-AUC, while DNN obtained the same test accuracy and marginally higher recall. Nevertheless, ARFN-DNN demonstrated that learnt fusion can change the balance of predictive behavior by achieving the highest precision and a competitive F1-score. External validation on separate, clinically obtained kid cohorts is crucial because the dataset used in this work is synthetic and somewhat limited. Before contemplating real-world pediatric implementation, further research should concentrate on bigger multicenter datasets, calibration analysis, prospective validation, clinically appropriate threshold setting, and integration of additional behavioral or multimodal information.

Future Work

Validation using bigger, independent, clinically obtained pediatric datasets should be the primary focus of future ARFN-DNN framework development. Although the current work shows encouraging performance on the given dataset, external validation across various hospitals, geographical areas, age groups, and demographic populations is necessary to ascertain whether the model retains its predictive accuracy in practical settings. To enhance the representation of ASD-related risk factors, future research might also look into other clinical and behavioral factors, such as developmental milestones, speech and language evaluations, sensory traits, sleep patterns, and other standardized screening tools.

By substituting an adaptive learning process that assesses each base learner's contribution based on its prediction uncertainty and reliability for the manually supplied model weights, the fusion mechanism can be further enhanced. To enhance the quality of the probability-level fusion, more sophisticated brain architectures, attention processes, and calibrated probability models can be explored. Beyond feature importance, explainable AI methods may be expanded to offer comprehensible explanations for individual risk assessments, assisting medical professionals in deciphering the elements that contribute to a high-risk outcome.

Prospective clinical examination is another crucial path. The suggested system might be evolved into an easy-to-use screening application that takes screening data and generates a risk score that medical experts can understand. The system should go through thorough external validation, calibration analysis, fairness assessment, privacy and security review, and clinical testing prior to practical deployment. As a result, future research can advance ARFN-DNN from experimental validation to a trustworthy supplemental screening system that facilitates early detection and referral while retaining qualified clinical specialists to make the final diagnosis.

REFERENCES

1. L. Breiman, "Random forests," Springer, Dordrecht, Netherlands, 2001.
2. N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," J. Artif. Intell. Res., AI Access Foundation, Mountain View, CA, USA, 2002.
3. T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," Association for Computing Machinery, New York, NY, USA, 2016.
4. S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," Neural Information Processing Systems Foundation, Long Beach, CA, USA, 2017.
5. I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in Proc. Int. Conf. on Learning Representations, New Orleans, LA, USA, 2019.
6. D. Pavithra, A. N. Jayanthi, R. Nidhya, and S. Balamurugan, "Autism screening tools with machine learning and deep learning methods: A review," Wiley, Hoboken, NJ, USA, 2022.

7. R. Francese and X. Yang, "Supporting autism spectrum disorder screening and intervention with machine learning and wearables: A systematic literature review," Springer Nature, Cham, Switzerland, 2022.
8. M. E. Minissi, I. A. Chicchi Giglioli, F. Mantovani et al., "Assessment of the autism spectrum disorder based on machine learning and social visual attention: A systematic review," Springer Nature, Cham, Switzerland, 2022.
9. M. R. Lamani and J. B. Pernabas, "A thorough review of deep learning in autism spectrum disorder detection: From data to diagnosis," Bentham Science Publishers, Sharjah, UAE, 2024.
10. World Health Organization, "Autism," World Health Organization, Geneva, Switzerland, 2025.
11. R. Blitz and S. Swoboda, "How Pediatricians Screen for Autism," American Academy of Pediatrics, Itasca, IL, USA, 2025.
12. M. A. Zoromba et al., "Artificial intelligence methods in early detection of autism spectrum disorder: A DSM-5 criterion-based systematic review," Wiley, Hoboken, NJ, USA, 2026.
13. Scikit-learn Developers, "LogisticRegression," Scikit-learn Project, Online Documentation, USA, 2026.
14. Anupama N. and C. M. Patil, "Data-driven approaches for autism detection: A comprehensive review of machine learning algorithms and datasets," Wiley, Hoboken, NJ, USA, 2026.
15. K. Srinivasa Reddy and Y. Lakshmi Annapurna, "A scoping review of machine learning and deep learning methods for autism spectrum disorder diagnosis and analysis," J. Vis. Exp., Cambridge, MA, USA, 2026.