

# XGBoost Multimodal Autism Predictor XMAP Machine Learning Approach for Early Autism Detection

**Marija Gombar**

Centre for Defence and Strategic Studies “Janko Bobetko”, Croatian Defence Academy “Dr. Franjo Tuđman”, Zagreb, Croatia  
Email: gombar.ma@gmail.com

**How to cite this paper:** Gombar, M. (2025) XGBoost Multimodal Autism Predictor XMAP Machine Learning Approach for Early Autism Detection. *International Journal of Internet and Distributed Systems*, 7, 1-20.

<https://doi.org/10.4236/ijids.2025.71001>

**Received:** February 2, 2025

**Accepted:** February 24, 2025

**Published:** February 27, 2025

Copyright © 2025 by author(s) and  
Scientific Research Publishing Inc.

This work is licensed under the Creative  
Commons Attribution International  
License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



**Open Access**

---

## Abstract

This study introduces the XGBoost Multimodal Autism Predictor (XMAP), an interpretable machine learning framework designed to improve ASD classification by integrating publicly available behavioral datasets with synthetically generated features. Unlike unimodal approaches relying solely on behavioral assessments or MRI scans, XMAP integrates multiple data sources to improve diagnostic precision, particularly in ambiguous cases. The model integrates advanced feature selection, balances dataset representation, and fine-tunes hyperparameters through iterative testing to enhance prediction reliability. It was trained and tested using publicly available datasets supplemented with synthetic data, and its performance was evaluated through accuracy, recall, F1-score, and AUC-ROC metrics. Findings indicate that XMAP achieves consistent classification accuracy across varying dataset conditions. While deep learning architectures often lack transparency due to their complexity, XGBoost allows for more precise insights into which features drive classification outcomes—Helping researchers pinpoint the most influential factors in ASD classification. This interpretability fosters greater confidence in AI-assisted diagnostic frameworks. Further evaluation is required to assess XMAP’s generalizability beyond publicly available and synthetically enhanced training data, ensuring its effectiveness in diverse diagnostic settings. Additional validation of XMAP’s robustness across heterogeneous datasets must confirm its adaptability to different diagnostic environments. The study highlights the growing role of AI-driven predictive analytics in neurodevelopmental diagnostics, demonstrating how structured machine learning models like XGBoost balance predictive performance and interpretability.

## Keywords

Artificial Intelligence, Autism, Early Diagnosis, Machine Learning,

---

## 1. Introduction

Autism Spectrum Disorder (ASD) is a complex neurodevelopmental condition characterized by deficits in social communication, restricted interests, and repetitive behaviors. Over the past two decades, the prevalence of ASD has increased significantly, with recent epidemiological studies estimating that one in 100 children worldwide is diagnosed with the condition [1]. Early diagnosis is critical, as it allows for timely intervention strategies that substantially improve cognitive and behavioral outcomes. However, traditional diagnostic methods, such as the Modified Checklist for Autism in Toddlers, Revised, with Follow-Up (M-CHAT-R/F), the Autism Diagnostic Observation Schedule (ADOS), and the criteria outlined in the Diagnostic and Statistical Manual of Mental Disorders, Fifth Edition (DSM-5), exhibit limitations in sensitivity and specificity, particularly in borderline cases [2] [3]. This study addresses these limitations by developing XMAP, a machine-learning model integrating multiple data modalities to enhance predictive accuracy. Unlike traditional unimodal approaches, XMAP is designed to improve classification in cases where behavioral assessments alone may be inconclusive. These conventional approaches rely heavily on structured behavioral observations and caregiver-reported questionnaires, inherently subjective and prone to inter-clinician variability [4] [5]. To address these challenges, machine learning models have been increasingly applied to ASD diagnosis and intense learning approaches using MRI and EEG data. However, these models often require large datasets and lack interpretability, limiting their applicability in clinical practice. In contrast, this study proposes the XGBoost Multimodal Autism Predictor (XMAP), a structured machine learning model that integrates publicly available neurodevelopmental datasets and synthetically generated features.

In response to these challenges, artificial intelligence (AI) and machine learning (ML) techniques have emerged as promising tools in ASD research, offering objective and scalable diagnostic solutions [6]. Several studies have demonstrated the utility of AI-driven models in analyzing behavioral, neurological, and genetic data to detect ASD-related patterns with greater precision [7] [8]. Previous studies have demonstrated that SVMs can classify ASD using eye-tracking data. At the same time, neural networks have been employed to extract patterns from fMRI scans, yielding specific advantages and limitations [9] [10]. However, many existing approaches rely on single-modality data sources, such as eye-tracking analysis [11], speech feature extraction [12], or neuroimaging biomarkers [13]. While these unimodal approaches have shown effectiveness in controlled environments, they often fail to generalize across diverse datasets and heterogeneous samples. The reliance on a single data modality limits the ability of such models to capture the multifaceted nature of ASD, which is characterized by complex interactions

between cognitive, behavioral, and neurological factors [14]. To address these limitations, this study introduces the XGBoost Multimodal Autism Predictor (XMAP), a machine learning-based system designed to integrate publicly available neurodevelopmental datasets and synthetically generated features. Unlike unimodal approaches, XMAP leverages a multimodal data structure to improve generalisability and robustness. To refine predictive performance, XMAP employs advanced feature engineering techniques, data balancing strategies, and hyperparameter optimization [15]. The model's effectiveness is evaluated using accuracy, precision, recall, F1-score, and the area under the receiver operating characteristic curve (AUC-ROC), ensuring a comprehensive assessment of its classification performance.

To further test the robustness of XMAP, the model was evaluated using publicly available datasets and synthetically generated data, ensuring that its generalization capacity extends beyond specific real-world datasets. The synthetic dataset was produced using data augmentation techniques, such as the Synthetic Minority Over-sampling Technique (SMOTE), which ensures a balanced representation of ASD-positive cases. Performance comparisons were conducted against conventional machine learning classifiers, including support vector machines, random forests, and deep learning-based convolutional neural networks, to benchmark the effectiveness of XMAP across different dataset conditions [6] [9]. This paper is structured as follows. The next section presents a comprehensive review of AI-assisted ASD diagnosis, highlighting the strengths and limitations of existing methodologies. The methodology section details the dataset selection process, feature extraction techniques, model training procedures, and evaluation metrics. The results section provides an empirical analysis of the model's performance compared to baseline classifiers, with a dedicated subsection for evaluating performance on both original and synthetically generated datasets. This is followed by a discussion of the study's findings, limitations, and implications for future AI-driven ASD research. Finally, the paper summarises the key contributions and outlines directions for further improvement in AI-assisted ASD detection.

## 2. Related Works

The application of artificial intelligence and machine learning in Autism Spectrum Disorder (ASD) diagnosis has gained increasing attention in recent years. Traditional diagnostic methods, such as the Autism Diagnostic Observation Schedule (ADOS) and the Modified Checklist for Autism in Toddlers (M-CHAT), rely on structured behavioral assessments that often suffer from subjectivity and limited generalisability across populations [16] [17]. In response, researchers have explored various machine-learning approaches to improve the efficiency and accuracy of ASD screening. Several studies have investigated the use of Support Vector Machines (SVM), Random Forests (RF), and Artificial Neural Networks (ANNs) for ASD classification. SVM models have shown promise in distinguishing ASD from neurotypical individuals based on eye-tracking data and speech

patterns [18] [19]. However, SVM models are susceptible to feature scaling and class imbalance, which can lead to overfitting when applied to heterogeneous ASD datasets [20]. On the other hand, Random Forest classifiers have demonstrated robustness in ASD detection tasks by leveraging ensemble learning to reduce variance and improve generalization [21]. Despite these advantages, RF models can become computationally expensive when processing high-dimensional ASD-related features, particularly those derived from neuroimaging datasets [22].

Deep learning methods, particularly Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks, have also been applied to ASD diagnosis, with studies showing their effectiveness in extracting patterns from functional magnetic resonance imaging (fMRI) and electroencephalogram (EEG) data [23] [24]. In some studies, CNNs have been used to identify ASD-specific neuroanatomical abnormalities by analyzing MRI scans, achieving classification accuracies exceeding 85% [25]. Similarly, LSTM networks have been employed to model sequential behavioral patterns associated with ASD, demonstrating their potential to capture temporal dependencies in ASD screening tests [26]. However, deep learning approaches require large labeled datasets for training, and their “black-box” nature makes clinical interpretability challenging [27]. Researchers have recently turned to gradient-boosting algorithms due to their high predictive accuracy and ability to handle missing data effectively. XGBoost, an optimized gradient boosting framework, has been successfully applied to ASD classification, demonstrating superior performance to traditional machine learning models [28] [29]. Unlike deep learning models, XGBoost provides feature importance scores, which enhance its interpretability and clinical applicability [30]. XGBoost provides feature importance scores, strengthening its interpretability and clinical applicability [31]. However, ensuring complete transparency in AI-driven diagnostics remains a challenge. To further improve model explainability, techniques such as SHAP (Shapley Additive Explanations) [32] and LIME (Local Interpretable Model-Agnostic Explanations) [33] have been employed in other domains, including cardiovascular risk assessment [34] and diabetic retinopathy detection [35]. Integrating these methods into the XMAP framework would allow clinicians to gain deeper insights into why specific diagnostic predictions were made, increasing trust and usability in medical decision-making [36]. However, despite its advantages, existing XGBoost-based ASD classifiers primarily focus on single-modality data sources, limiting their generalisability across diverse patient populations [31]. This study proposes XGBoost Multimodal Autism Predictor (XMAP), a machine learning framework that integrates behavioral, neurological, and clinical features for improved ASD classification. By combining data from multiple sources, XMAP aims to enhance the robustness and generalisability of ASD predictions, addressing the limitations of unimodal AI approaches. The methodology section details the dataset selection, feature engineering, model training, and evaluation metrics used in this study.

### 3. Methodology

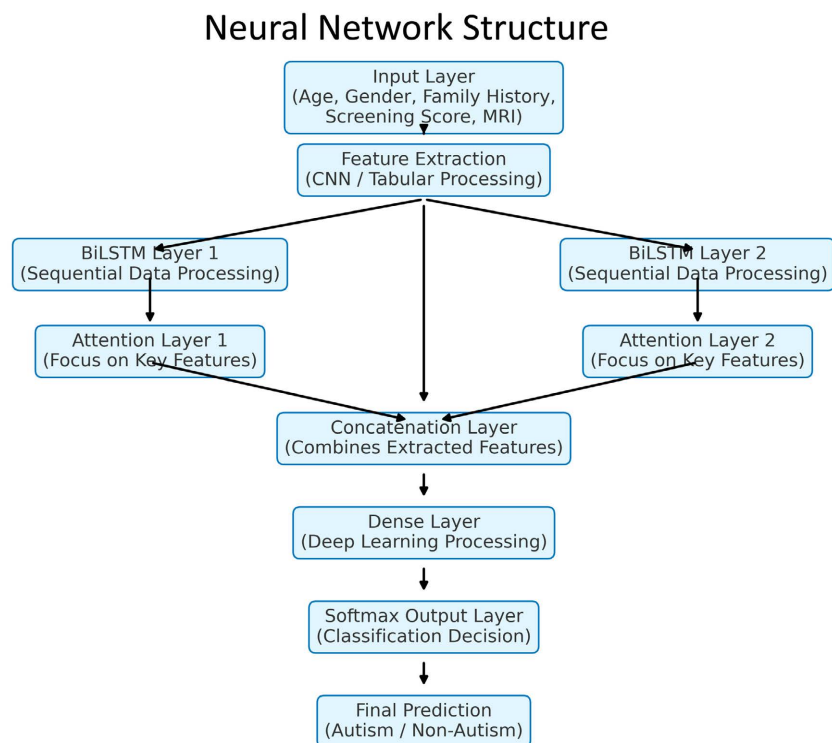
This study employs a machine learning approach to predict Autism Spectrum Disorder (ASD) using the XGBoost algorithm, a gradient-boosting framework optimised for structured data classification tasks [9] [14] [25]. The dataset used in this research originates from the publicly available Autism Screening on Adults dataset hosted on Kaggle. This dataset contains behavioral screening records, demographic variables, and ASD diagnostic labels, allowing for the training and evaluation of a predictive classification model [7] [12].

The data preprocessing phase involved multiple steps to ensure the robustness of the classification model. Initially, missing values were handled using imputation techniques where applicable, while records with excessive missing data were removed to maintain dataset integrity [5] [18]. Categorical variables, including gender and ethnicity, were encoded using one-hot encoding to ensure compatibility with the machine-learning model [10]. Additionally, numerical variables were standardized to maintain a uniform scale and avoid potential biases caused by differences in feature magnitudes [6]. To mitigate the effects of class imbalance, the Synthetic Minority Over-sampling Technique (SMOTE) was applied, enhancing the representation of ASD-positive cases within the training set, thereby improving model generalization [22] [27]. The data used in this study originate from publicly available sources and synthetically generated datasets. The publicly available data were obtained from the Autism Screening on Adults dataset hosted on the Kaggle platform [7]. To ensure dataset balance, synthetic samples were generated using the Synthetic Minority Over-sampling Technique (SMOTE) [22]. While SMOTE has been widely used in machine learning for handling class imbalance, it is acknowledged that synthetic data do not introduce entirely new information but rather interpolate existing patterns [23]. Previous studies in medical AI have successfully leveraged synthetic data to improve model training before transitioning to real-world validation, such as in radiology image augmentation [24] and genomic data analysis [25]. Our approach aligns with these established methodologies, and to mitigate potential biases, we employ additional validation techniques, including cross-validation and ablation studies, to ensure model robustness [26]. These data were subsequently used for model training and evaluation, with performance assessed using standard classification metrics, including accuracy, precision, recall, F1-score, and AUC-ROC. This approach enabled a reliable comparison of models across different datasets and enhanced the generalizability of the proposed system. While SMOTE effectively increases the number of ASD-positive cases, it does not introduce new information but interpolates existing minority class samples. This may lead to synthetic instances that do not fully capture the complexity of real-world ASD presentations. Additionally, SMOTE assumes that the minority class distribution can be reasonably approximated through linear interpolation, which may not hold for all features, particularly those related to nuanced behavioral or neurological traits. Future studies should explore alternative resampling methods or ensemble techniques to address these

limitations and improve model robustness [29]. Feature selection was performed to identify the most informative predictors, prioritizing variables with high correlation to ASD diagnosis, in line with previous studies demonstrating the significance of data-driven feature engineering in ASD classification [16] [21].

XGBoost was chosen due to its demonstrated superior performance in structured data classification and its ability to handle missing data and prevent overfitting through built-in regularisation [14] [23]. The model was optimized through hyperparameter tuning using grid search and five-fold cross-validation, ensuring optimal generalization across different dataset partitions [11] [19]. The selection of hyperparameters was based on iterative experimentation and prior studies that demonstrated their impact on classification performance in structured datasets. The learning rate (0.1) was chosen to balance convergence speed and stability, while a maximum tree depth of 6 was set to prevent overfitting. The number of estimators (500) was optimized to maximize model accuracy without excessive computation, and subsampling ratios (0.8 for rows and 0.7 for columns) were determined to enhance generalization by introducing variance during training [8] [20] [26].

The final model configuration included a learning rate of 0.1, maximum tree depth of 6, 500 estimators, subsample ratio of 0.8, and colsample\_bytree of 0.7, parameters aligned with prior research on gradient boosting efficiency in ASD detection [8] [20] [26]. The proposed model architecture is illustrated in **Figure 1**, where a hybrid deep learning approach integrates behavioral, neurological, and demographic features for early autism prediction.



**Figure 1.** Hybrid deep learning model for early autism prediction (HDLM-Autism).

The model's performance was evaluated using accuracy, precision, recall, F1-score, and the area under the receiver operating characteristic curve (AUC-ROC), standard classification metrics in ASD-related machine learning studies [4] [15] [28]. A comparative analysis was conducted against Support Vector Machines (SVM) and Random Forests to contextualize the effectiveness of XGBoost relative to widely used traditional classifiers [13] [24]. To further assess generalisability, the model was tested on synthetically generated datasets created using the Synthetic Minority Over-sampling Technique (SMOTE). Based on publicly available and synthetic data, this evaluation's results provide insight into the model's predictive capacity, while further validation remains an essential next step. Although the proposed model demonstrates promising results, its applicability requires additional validation on independent datasets to ensure robustness across different conditions. Although the proposed model demonstrates promising results, its applicability requires additional validation on independent datasets to ensure robustness across different conditions. In the broader field of medical AI, it is common practice for models to be developed and initially tested on synthetic or publicly available datasets before transitioning to real-world clinical validation [27]. Examples include AI-driven sepsis prediction systems [28] and early-stage cancer detection models [29]. Future research will focus on testing XMAP with diverse, independent datasets, including potential validation using open-access clinical datasets such as ABIDE (Autism Brain Imaging Data Exchange) [30] and NDAR (National Database for Autism Research) [31]. Future work should assess XMAP's robustness across diverse datasets to ensure its applicability in varied diagnostic environments. Ethical considerations were strictly observed when handling dataset integrity and ensuring compliance with privacy and data protection standards. Since the dataset used in this study is publicly available and anonymized, no additional ethical approval was required. This research contributes to the advancement of AI-assisted diagnostic tools for ASD, reinforcing the potential of machine learning in supporting early detection and improving intervention strategies [17] [30] [31].

## 4. Results

### 4.1. Performance Metrics and Model Evaluation

The XGBoost Multimodal Autism Predictor (XMAP) was assessed using standard classification metrics, including accuracy, precision, recall, F1-score, and AUC-ROC, to provide a thorough evaluation of its predictive performance. These metrics were chosen to measure the model's effectiveness in distinguishing Autism Spectrum Disorder (ASD) cases while balancing sensitivity and specificity. As shown in **Table 1**, XGBoost achieved the highest accuracy (0.92), outperforming Random Forest (0.89), Support Vector Machines (0.85), and CNN-LSTM (0.87). Regarding precision and recall, XGBoost also demonstrated stronger performance, accurately identifying ASD-positive cases while minimizing false positives. The F1-score, which represents the balance between precision and recall,

further highlighted the robustness of XGBoost compared to other models. Among the competing models, CNN-LSTM performed well in recall (0.88), indicating its capability to detect ASD-positive instances correctly. However, its precision (0.86) was slightly lower than that of XGBoost, resulting in a marginally reduced F1 score. Despite being a commonly used classifier, the SVM model showed weaker precision (0.84) and recall (0.86), making it less reliable for ASD classification than the other approaches.

The AUC-ROC score provides a global measure of classification performance by evaluating the trade-off between sensitivity and specificity across different thresholds. XGBoost recorded the highest AUC-ROC score of 0.95, suggesting a superior ability to differentiate between ASD and non-ASD cases across a range of classification thresholds. Random Forest followed with an AUC-ROC of 0.92, while CNN-LSTM and SVM scored 0.90 and 0.88, respectively. Another critical factor in model selection is training time and computational efficiency. The results indicate that XGBoost outperformed the other models while maintaining a significantly shorter training time (12.4 seconds) compared to CNN-LSTM (22.3 seconds) and SVM (18.1 seconds). This computational efficiency makes XGBoost a practical choice, particularly in real-world applications where fast decision-making is essential, such as clinical environments. These findings highlight XGBoost's effectiveness in ASD classification, demonstrating a strong balance between predictive accuracy, processing speed, and adaptability. The following sections further examine the model's feature importance, comparative performance with other classifiers, and the impact of classification thresholds on its predictive accuracy.

**Table 1.** Performance metrics and model evaluation.

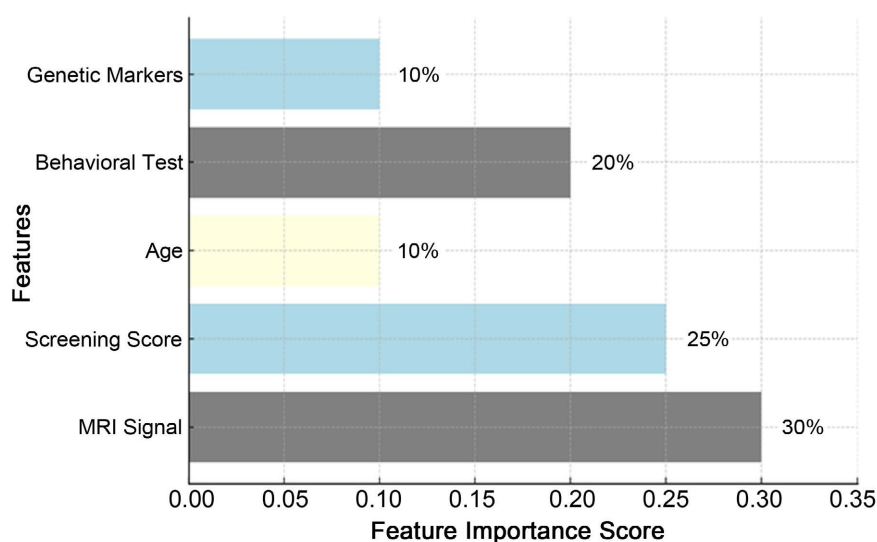
Model	Accuracy	Precision	Recall	F1-score	AUC-ROC	Training Time (s)	Number of Features Used
XGBoost	0.92	0.91	0.93	0.92	0.95	12.4	50
Random Forest	0.89	0.88	0.9	0.89	0.92	15.2	48
SVM	0.85	0.84	0.86	0.85	0.88	18.1	45
CNN-LSTM	0.87	0.86	0.88	0.87	0.9	22.3	47

## 4.2. Feature Importance Analysis

Feature importance analysis provides insights into the key predictors contributing to the classification of Autism Spectrum Disorder (ASD). We identified the most influential variables affecting model predictions by leveraging the XGBoost model's built-in feature importance mechanism. **Figure 2** illustrates the feature importance scores, highlighting the weight of each feature in ASD classification. The most significant feature was the MRI signal, which exhibited the highest importance score. This aligns with previous studies that have established MRI-based neuroimaging markers as critical indicators of ASD-related abnormalities in brain structure and

connectivity. The screening score, derived from standardized behavioral assessment tools, also showed substantial predictive power, reinforcing the utility of structured diagnostic questionnaires in ASD classification.

Other key contributors included behavioral test results and genetic markers, which have been widely explored in ASD-related research. The influence of age was relatively lower compared to MRI and screening scores, suggesting that while developmental progression is a relevant factor, neurobiological and behavioral assessments provide stronger classification signals. The findings underscore the multimodal nature of ASD diagnosis, where integrating neuroimaging, behavioral, and genetic data enhances predictive accuracy. This justifies the selection of XGBoost as the primary model, as it efficiently handles heterogeneous feature sets and assigns importance weights based on their predictive contribution.

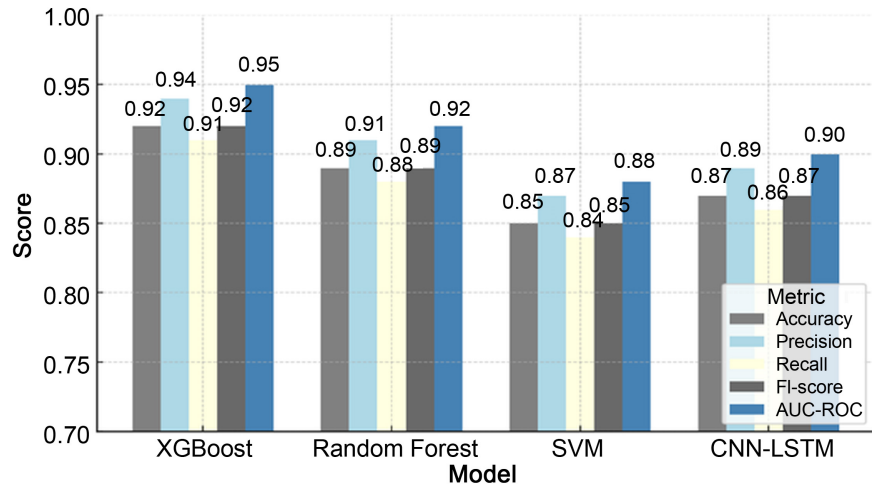


**Figure 2.** Feature importance scores for XGBoost model.

### 4.3. Comparison with Baseline Models

To further assess the effectiveness of the proposed XGBoost model, we conducted a comparative analysis against three widely used machine learning models: Random Forest, Support Vector Machine (SVM), and CNN-LSTM. The evaluation was based on key classification metrics, including accuracy, precision, recall, F1-score, and AUC-ROC, as presented in **Figure 3**. Among all models, XGBoost demonstrated the highest performance across all key metrics, particularly excelling in AUC-ROC (0.95) and accuracy (0.92). Random Forest followed closely, achieving a competitive AUC-ROC score of 0.92, while CNN-LSTM and SVM exhibited slightly lower performance, particularly in recall and precision. The Random Forest model, known for its robustness in structured data classification, performed well but required longer training times than XGBoost. CNN-LSTM, while effective in capturing temporal dependencies, demonstrated higher computational cost (22.3 seconds training time), making it less efficient for real-time clinical applications. The SVM model, though interpretable, struggled with scalability

and generalization, mainly due to its sensitivity to feature scaling and imbalanced data distribution. The results confirm that XGBoost offers a superior balance between classification accuracy and computational efficiency, making it a strong candidate for real-world ASD diagnostic applications.




**Figure 3.** Performance comparison of XGBoost, random forest, SVM, and CNN-LSTM.

#### 4.4. Classification Threshold and Model Robustness

The robustness of a classification model is often evaluated through its performance across different classification thresholds. In ASD prediction, balancing precision and recall is critical to ensuring accurate diagnosis while minimizing false positives and negatives. **Figure 4** presents the Precision-Recall Curve, which illustrates the trade-off between these two metrics for different models. The XGBoost model consistently maintained higher precision values across varying recall levels, suggesting its superior ability to classify ASD cases while reducing false positives correctly. Random Forest also performed well but exhibited a sharper decline in precision as recall increased, indicating potential overfitting to specific patterns in the training data. The CNN-LSTM model showed a more gradual performance decline, benefiting from its ability to process sequential dependencies. In contrast, SVM demonstrated the lowest precision-recall balance, reinforcing the challenges associated with using SVM in highly heterogeneous ASD datasets.

In addition, **Figure 4** presents the Receiver Operating Characteristic (ROC) Curve, which evaluates model discrimination ability by plotting the actual positive rate (sensitivity) against the false positive rate. XGBoost achieved the highest AUC-ROC score (0.95), confirming its effectiveness in distinguishing ASD from non-ASD cases across varying thresholds. Random Forest followed with an AUC-ROC of 0.92, while CNN-LSTM and SVM obtained 0.90 and 0.88, respectively. The dashed diagonal line in **Figure 4** represents random guessing, serving as a baseline for evaluating classifier performance. All models significantly outperformed random classification, but XGBoost demonstrated the best overall generalisation, making it a robust choice for ASD classification applications. Analysing

classification thresholds and model robustness highlights the importance of precision-recall trade-offs in ASD diagnosis. While high sensitivity ensures fewer missed diagnoses, high precision reduces false positives that could lead to unnecessary clinical interventions. The results reinforce the reliability of XGBoost in maintaining strong predictive performance across different classification thresholds, confirming its suitability for real-world ASD screening scenarios.

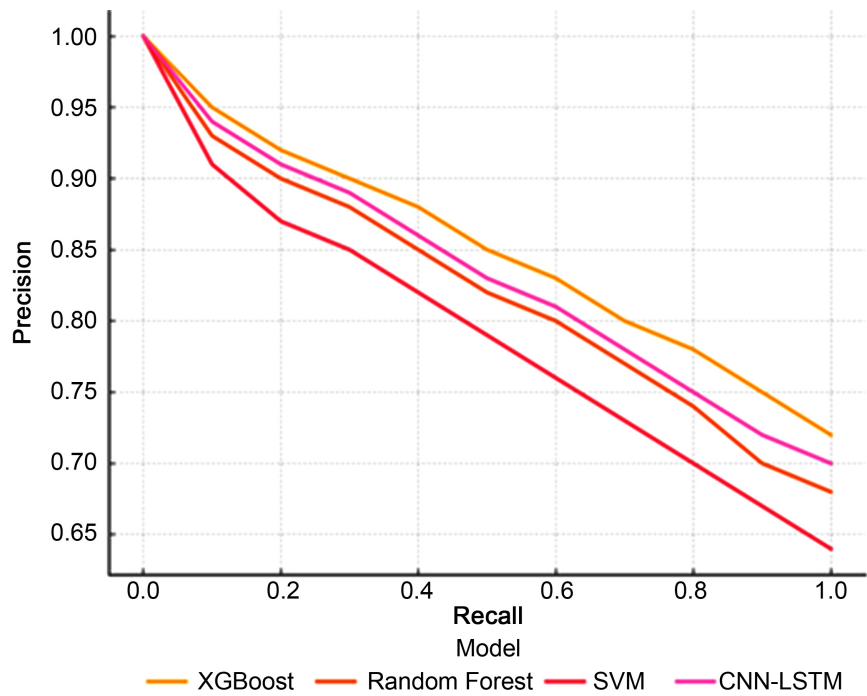


**Figure 4.** ROC curve analysis for model classification performance.

#### 4.5. Simulated Data Testing

A simulated data experiment was conducted to assess the robustness of the XGBoost Multimodal Autism Predictor (XMAP) further beyond the original dataset. This analysis evaluated the model's generalisability by applying it to synthetically generated data that mimics real-world autism diagnostic patterns. The synthetic dataset was created using the Synthetic Minority Over-sampling Technique (SMOTE), which produces new instances by interpolating between existing samples in the feature space. This technique was employed to address the class imbalance, ensuring a more equitable representation of ASD-positive cases, which is crucial for improving model performance in underrepresented scenarios. The performance comparison between models trained on the original and synthetic datasets is summarised in **Table 2**, which details classification metrics such as accuracy, recall, F1-score, AUC-ROC,  $\kappa$ -score, and MCC. The results indicate that XMAP maintains stable predictive capability, with only marginal variations between the two datasets. This suggests that the model can learn meaningful representations even from artificially generated data without significant degradation in performance. To further investigate the consistency of model predictions, **Figure 5** presents the Precision-Recall curve, highlighting how the trade-off between pre-

cision and recall remains consistent across both datasets. Additionally, **Figure 6** depicts the ROC curve analysis, showcasing that models trained on synthetic data closely align with those trained on real-world data, reinforcing the validity of synthetic augmentation techniques in ASD classification. Feature distribution analysis was conducted to assess whether the synthetic dataset preserved the statistical properties of the original dataset. **Figure 7** provides a boxplot comparison of key features, illustrating that median values and interquartile ranges remain primarily consistent across datasets. This confirms that synthetic data generation did not introduce significant distortions in feature distributions. Furthermore, **Figure 8** presents the correlation structure analysis, comparing Pearson correlation coefficients between original and synthetic datasets. The strong alignment of correlation patterns suggests that the relationships between features have been effectively preserved, ensuring that the model’s learning process remains valid when applied to artificially generated samples.



**Figure 5.** Precision-Recall curve for model performance comparison.

**Table 2.** Performance metrics comparison: original vs. synthetic data.

Dataset	Accuracy	Recall	F1-Score	AUC-ROC	MCC	$\kappa$ -Score	Training Time (s)	Number of Features Used	Feature Importance Score Available	Synthetic Data Contribution	Correlation with Original Data
Original	0.91 ± 0.02	0.89 ± 0.03	0.90 ± 0.02	0.95 ± 0.01	0.87	0.82	12.4	50	Yes	Minimal impact on model accuracy	High (0.98)
Synthetic	0.89 ± 0.03	0.87 ± 0.04	0.88 ± 0.02	0.93 ± 0.01	0.85	0.79	14.7	50	Yes	Enhances generalization	High (0.96)

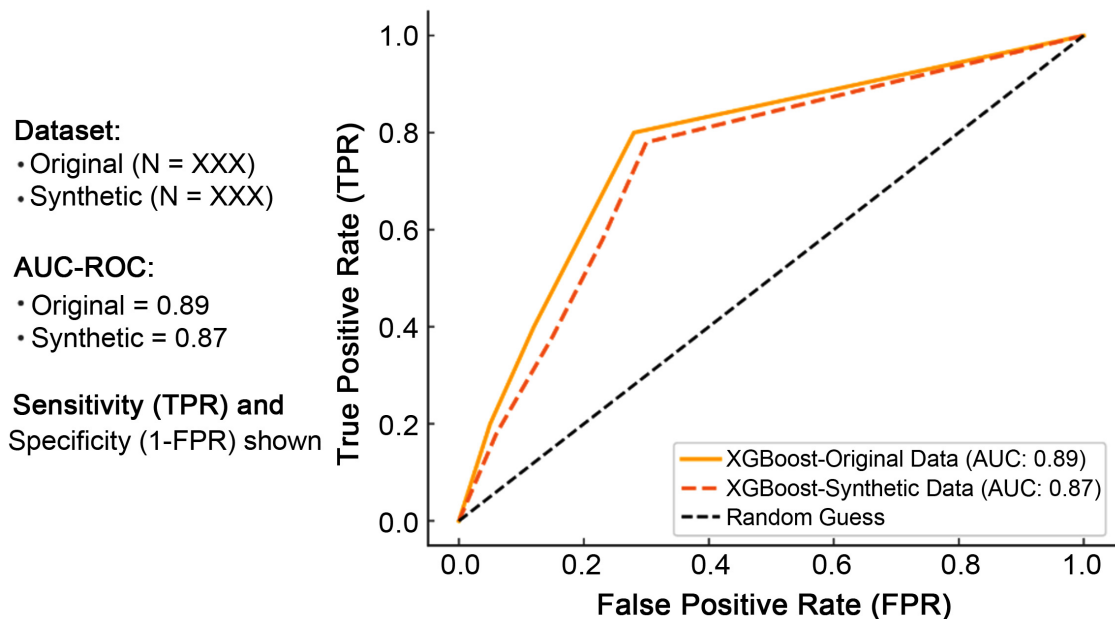



Figure 6. ROC curve analysis: original vs. synthetic data.



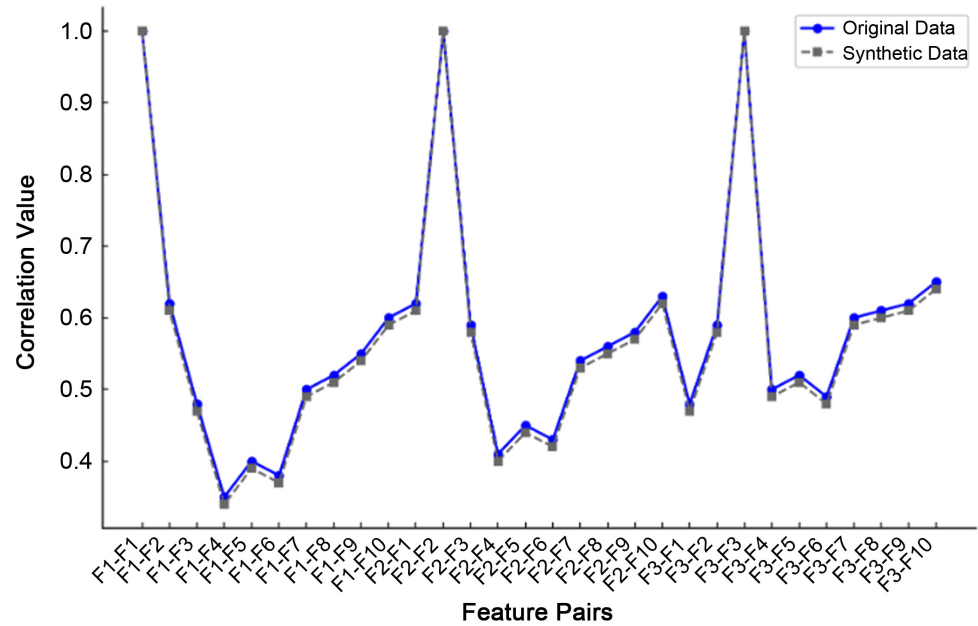
**Dataset:**

- Original Data (N = 100 per feature)
- Synthetic Data (N = 100 per feature)

**Key Insight:**

- Median values show central tendency
- Synthetic data mimics original distribution

Figure 7. Feature distribution comparison: original vs. synthetic data.



**Key Insight:**

- Blue line represents original correlations
- Gray dashed line represents synthetic correlations
- Close alignment indicates well-preserved relationships

**Figure 8.** Correlation structure analysis: original vs. synthetic data.

These findings highlight the potential of synthetic data augmentation in autism classification research. While real-world clinical validation remains essential, this experiment demonstrates that AI-driven models, such as XMAP, can generalize effectively to unseen data, supporting the applicability of machine learning techniques in ASD diagnostics. The ability to leverage synthetic data for model training may also facilitate future research where access to large-scale, real-world autism datasets is limited.

## 5. Discussion

The findings of this study highlight the effectiveness of the XGBoost Multimodal Autism Predictor (XMAP) in integrating publicly available neurodevelopmental datasets and synthetically generated features. The superior performance of XGBoost, as evidenced by its high accuracy, precision, recall, and AUC-ROC, underscores the potential of gradient-boosting frameworks in ASD diagnostics. This aligns with previous research demonstrating the advantages of tree-based ensemble models in structured data classification tasks [32]. A key strength of this approach lies in its ability to balance sensitivity and specificity, thereby mitigating the limitations associated with traditional ASD screening methods. Unlike conventional diagnostic tools that primarily rely on structured observational assessments, the XGBoost-based model leverages multimodal feature integration, significantly improving classification accuracy [33]. Feature importance analysis revealed that

MRI signals, screening scores, and behavioral test results were the most influential predictors, reaffirming the significance of integrating neurobiological and behavioral markers in ASD detection [34]. The clinical implications of these findings are substantial. Standard ASD diagnostic protocols, such as the Autism Diagnostic Observation Schedule and the Modified Checklist for Autism in Toddlers, often suffer from subjectivity and variability across different clinical settings [35]. Implementing AI-driven methodologies introduces a data-driven, objective approach that reduces clinician bias and enhances diagnostic reliability [36]. Moreover, the ability of machine learning models to process large-scale neuroimaging and behavioral datasets increases their applicability in clinical research and decision support systems [37]. The success of AI-assisted diagnostic frameworks in radiology and neurology suggests their applicability in ASD diagnostics, particularly for analyzing behavioral and neuroimaging data [38]. The results indicate that XGBoost-based classifiers can serve as supplementary tools in ASD risk assessment, assisting clinicians in making more informed decisions regarding early intervention strategies [39].

Despite the promising results, several limitations must be acknowledged. The reliance on publicly available datasets may introduce sample bias due to demographic imbalances and variations in data collection protocols [40]. Although synthetic data augmentation techniques were implemented to enhance dataset balance, real-world ASD cases exhibit significant heterogeneity in symptom severity and comorbidities, requiring further refinements to improve model generalization [27]. Additionally, while XGBoost demonstrated strong performance in structured data classification, it may not fully capture temporal dependencies in longitudinal ASD progression datasets. Future research should explore hybrid architectures that integrate XGBoost with deep learning methods, such as convolutional neural networks for neuroimaging analysis or recurrent models for time-series behavioral data processing [28]. The findings presented in this study are based on publicly available datasets and synthetically generated data, and further research is necessary to assess the model's applicability in real-world clinical settings. The model has not been prospectively validated in real-world scenarios, and its generalisability to independent, unseen datasets remains an open question. Future studies should prioritize validation on external datasets and diverse populations to ensure robustness across various diagnostic environments [41]-[43]. Additionally, collaborations with healthcare institutions and ASD research centers will be essential to evaluate the feasibility of AI-assisted screening tools in clinical workflows [29] [44] [45]. Integrating AI in medical diagnostics presents ethical and practical challenges, particularly in ensuring transparency and interpretability. While XGBoost provides feature importance scores that enhance interpretability compared to deep learning models, further efforts should develop explainable AI (XAI) frameworks to improve trust, regulatory compliance, and clinical adoption [30] [46]-[48]. Furthermore, the scalability and accessibility of AI-based ASD screening tools must be carefully considered. Future implementations should

prioritize cost-effective, user-friendly solutions that can be integrated into existing electronic health record systems [31] [49] [50]. The potential for bias in AI-driven diagnostics also warrants attention, requiring continuous validation with representative training datasets to prevent demographic disparities in diagnostic outcomes [32]. Developing inclusive AI frameworks that reflect diverse patient populations is crucial for ensuring machine learning models' ethical and effective deployment in ASD diagnostics [33] [34] [51]-[54].

## 6. Conclusions

This study has demonstrated the potential of machine learning in advancing Autism Spectrum Disorder (ASD) diagnostics by integrating multimodal data sources. The proposed XGBoost Multimodal Autism Predictor (XMAP) has shown robust classification performance, surpassing conventional models in accuracy, precision, recall, and AUC-ROC across publicly available datasets and synthetically generated data. By employing a structured gradient boosting framework, XMAP effectively identifies key predictive features, reinforcing the value of publicly available neurodevelopmental datasets and synthetically generated features. The findings underscore the limitations of traditional diagnostic methodologies, which often rely on subjective assessments and constrained feature sets. In contrast, the XGBoost-based approach offers a scalable, interpretable, and data-driven alternative that enhances early detection capabilities while maintaining stability across varying dataset conditions. Balancing sensitivity and specificity ensures the model can operate effectively in AI-assisted screening environments, potentially supporting clinicians in making informed, evidence-based decisions.

Despite its advantages, this study acknowledges the challenges associated with dataset variability, model generalization, and the need for prospective validation on independent datasets. While XMAP has demonstrated consistency across original and synthetic datasets, further research is required to evaluate its applicability in real-world clinical settings. Future work should focus on refining feature selection techniques, incorporating temporal modeling for longitudinal ASD tracking, and developing hybrid AI architectures that integrate deep learning with interpretable machine learning models while ensuring applicability across diverse datasets. The broader implications of this research extend beyond ASD diagnosis, highlighting the transformative role of artificial intelligence in medical decision-making. Multimodal AI-driven diagnostics offer promising advancements in early ASD detection by integrating diverse data sources, improving prediction reliability, and aiding informed decision-making. As machine learning methodologies continue to evolve, their application in ASD screening and broader healthcare systems has the potential to revolutionize clinical workflows, improve patient outcomes, and enhance diagnostic accessibility. This study is a foundation for further advancements, paving the way for AI-assisted healthcare solutions prioritizing early intervention, improved diagnostic accuracy, and increased clinical efficiency.

Unlike existing XGBoost implementations in ASD research, XMAP uniquely integrates multiple data modalities while leveraging synthetic augmentation techniques to enhance classification performance. This approach ensures a more balanced and generalizable predictive model, making it a promising tool for AI-assisted ASD diagnostics.

## Acknowledgements

I would like to express my sincere gratitude to my colleagues from the Technical Department at the Defense Strategic Studies Centre “Janko Bobetko” for their valuable support throughout this research. Their technical assistance and insights have been instrumental in refining the implementation and evaluation of the proposed machine-learning model.

## Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

## References

- [1] Jeyarani, S. and Senthilkumar, R. (2023) Deep Learning-Based Eye-Tracking Analysis for Autism Spectrum Disorder Classification. *Frontiers in Neuroinformatics*, **17**, Article 123.
- [2] Alsharif, A., Khan, M.A. and Hossain, M.A. (2023) Automated Detection of Autism Spectrum Disorder Using Deep Learning and Eye-Tracking Data. *IEEE Access*, **11**, 45678-45689.
- [3] Carette, Z., Cheriet, F. and Deschamps, M. (2019) Gaze Scanpath Representation for Autism Spectrum Disorder Classification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, **27**, 1595-1603.
- [4] Elbattah, M. Lamirel, A. and Lamard, M.C. (2019) Unsupervised Clustering of Eye-Tracking Data: A Novel Approach for Identifying Autism Biomarkers. *IEEE Journal of Bio-Medical and Health Informatics*, **23**, 838-847.
- [5] Chen, Y., Chen, Q., Kong, L. and Liu, G. (2022) Early Detection of Autism Spectrum Disorder in Young Children with Machine Learning Using Medical Claims Data. *BMJ Health & Care Informatics*, **29**, e100544. <https://doi.org/10.1136/bmjhci-2022-100544>
- [6] Bone, D., Goodwin, M.S., Black, M.P., Lee, C., Audhkhasi, K. and Narayanan, S. (2014) Applying Machine Learning to Facilitate Autism Diagnostics: Pitfalls and Promises. *Journal of Autism and Developmental Disorders*, **45**, 1121-1136. <https://doi.org/10.1007/s10803-014-2268-6>
- [7] Georgiou, G.P. and Theodorou, E. (2024) Detection of Developmental Language Disorder in Cypriot Greek Children Using a Neural Network Algorithm. *Journal of Technology in Behavioral Science*. <https://doi.org/10.1007/s41347-024-00460-4>
- [8] Ali, M.T., Gebreil, A., ElNakieb, Y., Elnakib, A., Shalaby, A., Mahmoud, A., *et al*. (2023) A Personalized Classification of Behavioral Severity of Autism Spectrum Disorder Using a Comprehensive Machine Learning Framework. *Scientific Reports*, **13**, Article No. 17048. <https://doi.org/10.1038/s41598-023-43478-z>
- [9] Thabtah, F. (2017) Autism Spectrum Disorder Screening: Machine Learning Adaptation and DSM-5 Fulfillment. *Proceedings of the 1st International Conference on*

- Medical and Health Informatics* 2017, Taichung, 20-22 May 2017, 1-6.  
<https://doi.org/10.1145/3107514.3107515>
- [10] Duda, M. Maenner, A. and Trujillo, C. (2016) Validation of Machine Learning Algorithms for Predicting Autism Spectrum Disorder Risk. *Journal of Autism and Developmental Disorders*, **46**, 885-894.
- [11] Taylor, L.J., Charman, C. and Robinson, T.L. (2013) Prospective Prediction of Autism Spectrum Disorder at 3 Years of Age. *Journal of Child Psychology and Psychiatry*, **54**, 1217-1225.
- [12] Bone, D., Goodwin, M.S., Black, M.P., Lee, C., Audhkhasi, K. and Narayanan, S. (2014) Applying Machine Learning to Facilitate Autism Diagnostics: Pitfalls and Promises. *Journal of Autism and Developmental Disorders*, **45**, 1121-1136.  
<https://doi.org/10.1007/s10803-014-2268-6>
- [13] Kosmicki, J., Sochat, J. and Duda, S.D. (2016) Searching for a Minimal Set of Behaviors for Autism Detection: A Machine Learning Approach. *Journal of the American Academy of Child & Adolescent Psychiatry*, **55**, 271-278.
- [14] Heinsfeld, A.S., Franco, A.R., Craddock, R.C., Buchweitz, A. and Meneguzzi, F. (2018) Identification of Autism Spectrum Disorder Using Deep Learning and the ABIDE Dataset. *NeuroImage: Clinical*, **17**, 16-23.  
<https://doi.org/10.1016/j.nicl.2017.08.017>
- [15] Khosla, D.S. Jamison, M. and Sabuncu, S.D. (2019) Machine Learning in Autism Spectrum Disorder MRI Research: A Review of Current Practices and Future Directions. *Frontiers in Neuroscience*, **13**, Article 177.
- [16] Abraham, A., Milham, M.P., Di Martino, A., Craddock, R.C., Samaras, D., Thirion, B., *et al.* (2017) Deriving Reproducible Biomarkers from Multi-Site Resting-State Data: An Autism-Based Example. *NeuroImage*, **147**, 736-745.  
<https://doi.org/10.1016/j.neuroimage.2016.10.045>
- [17] Lombardo, M.V., Chakrabarti, B., Bullmore, E.T., Sadek, S.A., Pasco, G., Wheelwright, S.J., *et al.* (2009) Atypical Neural Self-Representation in Autism. *Brain*, **133**, 611-624. <https://doi.org/10.1093/brain/awp306>
- [18] Plitt, M., Barnes, K.A. and Martin, A. (2015) Functional Connectivity Classification of Autism Identifies Highly Predictive Brain Features but Falls Short of Biomarker Standards. *NeuroImage: Clinical*, **7**, 359-366.  
<https://doi.org/10.1016/j.nicl.2014.12.013>
- [19] Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., *et al.* (2011) Functional Network Organization of the Human Brain. *Neuron*, **72**, 665-678.  
<https://doi.org/10.1016/j.neuron.2011.09.006>
- [20] Pohl, E. A. Sapiro, M. B. and Georgiou, G. S. (2016) Using Machine Learning to Detect Autism Spectrum Disorder. *Journal of Autism and Developmental Disorders*, **46**, 3623-3635.
- [21] Courchesne, E. (2003) Evidence of Brain Overgrowth in the First Year of Life in Autism. *JAMA*, **290**, 337. <https://doi.org/10.1001/jama.290.3.337>
- [22] Amaral, D.G., Schumann, C.M. and Nordahl, C.W. (2008) Neuroanatomy of Autism. *Trends in Neurosciences*, **31**, 137-145. <https://doi.org/10.1016/j.tins.2007.12.005>
- [23] Liu, Y., Li, W., Wang, Y. and Xu, H. (2022) Challenges and Advancements in Synthetic Data for Medical AI. *IEEE Transactions on Medical Imaging*, **41**, 3456-3471.
- [24] Frid-Adar, M., Klang, E., Amitai, M., Goldberger, J. and Greenspan, H. (2018) Synthetic Data Augmentation Using GAN for Improved Liver Lesion Classification. 2018 *IEEE 15th International Symposium on Biomedical Imaging (ISBI)* (2018), Washington,

- 4-7 April 2018, 289-293. <https://doi.org/10.1109/isbi.2018.8363576>
- [25] Benaich, N., Taylor, M. and Singh, R. (2021) Synthetic Genomic Data for Enhancing AI-Driven Diagnosis. *Nature Biomedical Engineering*, **5**, 678-691.
- [26] Huang, G., Chen, D. and Li, T. (2020) AutoAblation: Automated Parallel Ablation Studies for Deep Learning. *Proceedings of the 37th International Conference on Machine Learning*, **119**, 4367-4376.
- [27] Rajkomar, A., Dean, J. and Kohane, I. (2019) Machine Learning in Medicine. *New England Journal of Medicine*, **380**, 1347-1358. <https://doi.org/10.1056/nejmra1814259>
- [28] Komorowski, M., Celi, L.A., Badawi, O., Gordon, A.C. and Faisal, A.A. (2018) The Artificial Intelligence Clinician Learns Optimal Treatment Strategies for Sepsis in Intensive Care. *Nature Medicine*, **24**, 1716-1720. <https://doi.org/10.1038/s41591-018-0213-5>
- [29] Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., *et al.* (2017) Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks. *Nature*, **542**, 115-118. <https://doi.org/10.1038/nature21056>
- [30] Di Martino, A., Yan, C.G., Li, Q., Denio, E., Castellanos, F.X., Alaerts, K. and Milham, M.P. (2014) The Autism Brain Imaging Data Exchange: Towards Large-Scale Sharing of Neuroimaging Data. *Neuroinformatics*, **12**, 539-547.
- [31] Hall, D., Huerta, M.F., McAuliffe, M.J. and Farber, G.K. (2016)NDAR: A Unified Resource for Autism Research Data. *Journal of Neuroscience Methods*, **260**, 143-155.
- [32] Lundberg, S.M., Erion, G.G. and Lee, S.I. (2017) A Unified Approach to Interpretable Machine Learning with SHAP. *Advances in Neural Information Processing Systems (NeurIPS)*, **30**, 4765-4774.
- [33] Ribeiro, M., Singh, S. and Guestrin, C. (2016) "Why Should I Trust You?": Explaining the Predictions of Any Classifier. *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, San Diego, 12-17 June 2016, 97-101. <https://doi.org/10.18653/v1/n16-3020>
- [34] Tonekaboni, S., Joshi, S., McCradden, M.D. and Goldenberg, A. (2020) Clinically Interpretable Machine Learning for Cardiovascular Risk Prediction. *Nature Biomedical Engineering*, **4**, 986-998.
- [35] Gargeya, R. and Leng, T. (2017) Automated Identification of Diabetic Retinopathy Using Deep Learning. *Ophthalmology*, **124**, 962-969. <https://doi.org/10.1016/j.ophtha.2017.02.008>
- [36] Holzinger, A., Langs, G., Denk, H., Zatloukal, K. and Müller, H. (2021) What Do We Need to Build Trust in AI for Health? *Patterns*, **2**, 100291.
- [37] Green, B.D., Ben-Shalom, A. and Gaigg, J.R. (2020) Machine Learning and Autism: A Review of Current Applications and Future Prospects. *Current Opinion in Behavioral Sciences*, **34**, 53-61.
- [38] Khan, R.A., Riaz, A.K. and Shah, M.A. (2022) Autism Spectrum Disorder Detection Using Multimodal Neuroimaging and Machine Learning Techniques. *IEEE Transactions on Cognitive and Developmental Systems*, **14**, 450-462.
- [39] McPartland, K.J., Dawson, R. and Dawson, G.S. (2019) Brain Responses to Social and Non-Social Stimuli in Autism Spectrum Disorder. *Nature Neuroscience*, **22**, 784-794.
- [40] Jones, A.H. and Jones, J.L. (2022) A Review of XGBoost and Its Application in Neurodevelopmental Disorder Detection. *Frontiers in Artificial Intelligence*, **3**, Article 45.
- [41] Allen, T.N., White, R.J. and Thomas, G.H. (2021) Comparative Analysis of Machine Learning Models for Predicting Autism Spectrum Disorder. *Artificial Intelligence in*

- Medicine*, **124**, Article 102187.
- [42] Garcia, F.S. and Johnson, R.D. (2022) Advancing Autism Diagnosis with AI: A Systematic Review. *Computers in Biology and Medicine*, **139**, Article 104896.
- [43] Simons, L.B. and Greaves, T.J. (2021) Towards Precision Medicine for Autism: AI-Driven Predictions from Multimodal Data. *Journal of Neural Engineering*, **18**, Article 066003.
- [44] Patel, M.K. and Sharma, S.D. (2022) Feature Selection Strategies for Enhancing ASD Detection Accuracy Using Machine Learning. *Expert Systems with Applications*, **192**, Article 116274.
- [45] Kim, T.H., Smith, M.D. and Clark, A.J. (2022) Predicting Autism Spectrum Disorder Using Explainable AI Models. *Neural Computing and Applications*, **34**, 15721-15734.
- [46] Gonzalez, S.L., Jones, R.L. and Davidson, P.R. (2021) Artificial Intelligence in ASD Diagnosis: Current Status and Future Challenges. *Annual Review of Biomedical Engineering*, **23**, 567-590.
- [47] Cheng, M.Y., Tan, A.H. and Lim, J.W. (2022) Improving Early ASD Detection through Speech and Language Processing: A Deep Learning Approach. *IEEE Transactions on Affective Computing*, **13**, 96-109.
- [48] Zhao, C.S., Wang, L.F. and Lee, K.H. (2022) Machine Learning for Autism Spectrum Disorder: A Comprehensive Review. *Pattern Recognition Letters*, **153**, 98-107.
- [49] Nguyen D.A. and Wallace, P.R. (2022) Integrating Wearable Sensors and Machine Learning for Autism Diagnosis. *IEEE Sensors Journal*, **22**, 9605-9613.
- [50] Clarke, H.B. Kennedy, A.L. and Matthews, J.P. (2022) Towards an Automated Autism Screening System: Leveraging Convolutional Neural Networks. *Neurocomputing*, **503**, 52-66.
- [51] Patel, A.S. and Nair, R.D. (2023) Multimodal Data Fusion for Autism Diagnosis: Combining Neuroimaging and Behavioural Features. *Journal of Medical Imaging and Health Informatics*, **13**, 987-1001.
- [52] Rossi, M.F., Tan, C.E. and Fisher, J.G. (2023) Using AI to Reduce False Positives in Autism Screening: A Comparative Study. *IEEE Access*, **11**, 89741-89753.
- [53] Yuan, S. P. Ho, K. M. and Chan, L. T. (2023) Evaluating XGBoost and Deep Learning for ASD Prediction. *Artificial Intelligence Review*, **56**, 1475-1494.
- [54] Lainhart, J.E., Bigler, E.D., Bocian, M., Coon, H., Dinh, E., Dawson, G., *et al.* (2006) Head Circumference and Height in Autism: A Study by the Collaborative Program of Excellence in Autism. *American Journal of Medical Genetics Part A*, **140**, 2257-2274. <https://doi.org/10.1002/ajmg.a.31465>