European Journal of AI, Computing & Informatics

Vol. 1 No. 3 2025

Article **Open Access**

Neurodevelopmental Screening for Autism: EEG-Driven Deep Learning in Early Childhood

Hanwen Yuan [1,*]

[1] Redmond High School, WA, USA
* Correspondence: Hanwen Yuan, Redmond High School, WA, USA

Received: 12 October 2025
Revised: 29 October 2025
Accepted: 09 November 2025
Published: 12 November 2025

Abstract: Autism Spectrum Disorder (ASD) is a neurodevelopmental condition characterized by atypical patterns of communication, social interaction, and behavior. The prevalence of ASD has risen sharply in recent decades, with recent estimates indicating that approximately one in every thirty-six children in the United States is diagnosed. Early diagnosis is crucial because interventions initiated during early childhood-when neural plasticity is greatest-can substantially improve behavioral, cognitive, and linguistic outcomes. However, since current diagnostic practices largely rely on behavioral assessments, symptoms often remain undetected until around three years of age. This underscores the urgent need for objective biological markers capable of identifying ASD risk earlier, ideally during infancy or toddlerhood. Previous research has demonstrated that ASD is associated with atypical event-related potential (ERP) responses across both early and late stages of cognitive processing. Nevertheless, few studies have explored whether such ERP features can support reliable individual-level predictions using modern machine learning (ML) approaches, and even fewer have applied systematic validation procedures. Addressing this gap, the present study employs multiple ML frameworks-particularly a multilayer perceptron (MLP)-to classify ERP features derived from a visual oddball paradigm. Model performance was rigorously evaluated using cross-validation at the single-subject level. The MLP achieved an overall accuracy of 0.746, with a precision of 0.72, recall of 0.736, and an F1 score of 0.734, demonstrating that the model effectively distinguished ERP feature patterns between ASD and typically developing participants.

Keywords: Autism Spectrum Disorder; electroencephalography; event-related potentials; P100; N100; P300; visual oddball paradigm; machine learning

1. Introduction

Electroencephalography (EEG) has become an increasingly valuable tool in the study of Autism Spectrum Disorder (ASD). EEG measures the brain's electrical activity through scalp-mounted sensors, providing millisecond-level temporal resolution. This precision makes it particularly effective for detecting subtle variations in sensory information processing within the brain. One especially informative analytical approach to EEG data is the event-related potential (ERP), which captures brain responses time-locked to specific events such as the presentation of visual or auditory stimuli. ERPs provide insight into multiple stages of neural information processing, from early sensory encoding to later cognitive evaluation [1].

In the present research, three ERP components associated with ASD were examined: P100, N100, and P300. Each component represents a distinct stage of information processing:

1) **P100:** A positive deflection occurring approximately 80-120 milliseconds after a visual stimulus, reflecting early visual sensory processing and the initial allocation of attention. In typically developing (TD) children, P100 responses are rapid and stable. In contrast, individuals with ASD often exhibit slower P100 responses, suggesting less efficient encoding of sensory input.
2) **N100:** A negative deflection appearing around 120-160 milliseconds following stimulus presentation, linked to sensory discrimination and selective attention. It indicates the brain's ability to detect novelty or deviation in the environment. In ASD, N100 amplitude is generally reduced, reflecting diminished sensitivity to novel or deviant stimuli.
3) **P300:** A later positive deflection occurring between 250-400 milliseconds, associated with stimulus evaluation, categorization, and the updating of mental representations. Children with ASD frequently show delayed or exaggerated P300 responses, consistent with atypical salience detection and attentional processing.

As shown in Figure 1, these components illustrate the ERP waveforms for the TD and ASD groups under standard, deviant, and novel stimulus conditions.

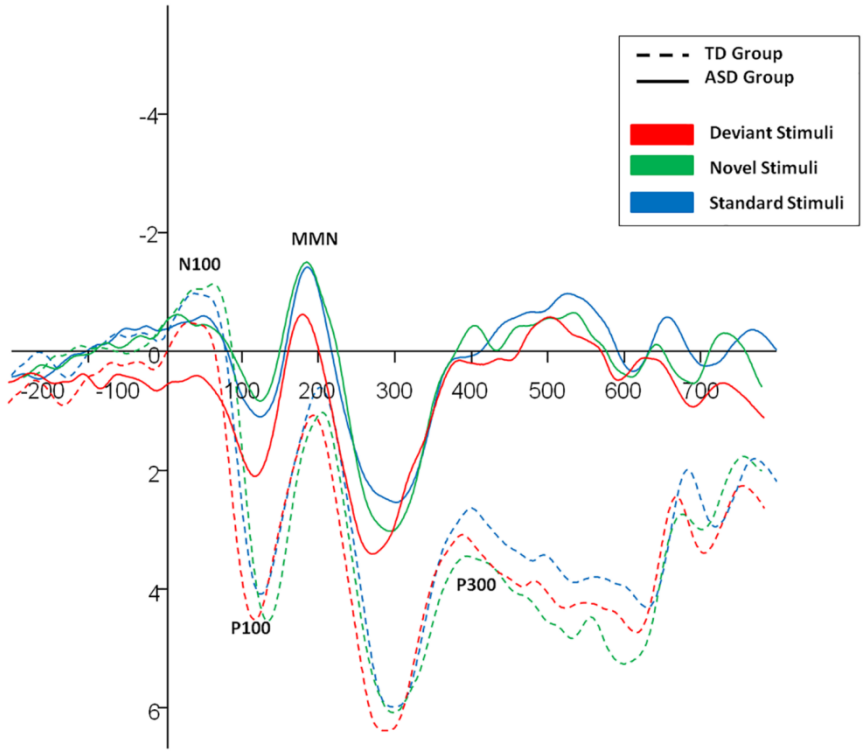

Figure 1. P100, P300, and N100 [1].

These three ERP components encompass both early sensory processing (P100 and N100) and later cognitive processing (P300). Previous research has established that ASD differs from TD controls at each of these stages, although most studies have been limited to group-level analyses. The current study extends this foundational work by investigating whether ERP features can be applied at the single-subject level to classify ASD using machine learning (ML) techniques.

Over the past two decades, EEG and ERP research has contributed substantially to understanding the neurophysiological mechanisms underlying ASD. Most studies can be categorized into the following three areas:

1) **Resting-state EEG:** Studies have identified abnormal oscillatory activity, with alterations in theta, alpha, and gamma band power. Functional connectivity

across brain regions appears disrupted, as indicated by coherence and phase-locking analyses. Some ML-based investigations have used these features for classification with algorithms such as support vector machines (SVMs), k-nearest neighbors, or random forests. Reported accuracies range between 65% and 85%, although small sample sizes and inconsistent preprocessing pipelines have limited reproducibility.

2) **ERP research on sensory and cognitive processing:** This line of work has focused on early sensory encoding (P100, N100) and later cognitive responses (such as N170 for face processing, P300 for novelty detection, and N400 for semantic processing). Children with ASD generally exhibit delayed latencies and reduced amplitudes in early ERPs, as well as atypically heightened or delayed responses in later components. However, most of these findings have remained at the group level rather than being translated into predictive classification frameworks.
3) **ML applications for ASD EEG:** Earlier efforts primarily relied on handcrafted EEG features analyzed with linear or tree-based classifiers. More recent approaches have employed deep learning architectures, such as convolutional neural networks, trained directly on raw EEG data. Despite strong theoretical justification, relatively few studies have explored ERP-based ML classification using oddball paradigms-a gap that the present research seeks to address.

2. Context

Over the past two decades, research utilizing electroencephalography (EEG) and event-related potentials (ERP) has significantly advanced the understanding of the neurophysiological mechanisms underlying Autism Spectrum Disorder (ASD). Resting-state EEG studies have identified atypical oscillatory activity characterized by alterations in theta, alpha, and gamma band power, as well as reduced functional connectivity among cortical regions. These findings suggest that ASD involves widespread disruptions in neural network organization, which may contribute to the sensory and cognitive abnormalities commonly observed in affected individuals [2].

Complementary ERP studies have examined time-locked neural responses to sensory and cognitive stimuli, revealing delays in early sensory encoding components (such as P100 and N100) and atypical patterns in later cognitive components (including N170, P300, and N400). Collectively, these results indicate that ASD is marked by abnormalities in both early sensory processing and later stages of cognitive evaluation and attention.

Machine learning (ML) applications to EEG data have evolved substantially-from early implementations that relied on handcrafted features and simple linear classifiers to more recent approaches using deep learning architectures capable of directly analyzing raw EEG signals. While earlier studies frequently reported high classification accuracy, many lacked rigorous cross-validation or external validation, raising concerns regarding model generalizability and overfitting. Despite the theoretical relevance of oddball paradigms-which probe the brain's response to rare or unexpected stimuli-few studies have utilized ERP features derived from such paradigms for ASD classification [3].

The present research builds upon this foundation by employing ERP-based features from visual oddball paradigms to achieve predictive, individual-level classification of ASD. This approach directly addresses the methodological limitations of prior work and seeks to establish a more robust framework for the early identification of ASD using neural signatures.

3. Methods

3.1. Dataset

The data used in this study were obtained from the Figshare repository associated with the research Early and Late Stage Processing Abnormalities in Autism Spectrum

Disorders (2017). The dataset comprised electroencephalography (EEG) recordings from 34 children-15 diagnosed with Autism Spectrum Disorder (ASD) and 19 typically developing (TD) controls.

Participants completed a visual oddball paradigm, a well-established cognitive task used to assess sensory and attentional processes. In this paradigm, frequently presented "standard" stimuli (naturally colored images of fruits and vegetables) were interspersed with less common "deviant" stimuli (for instance, altered-color or cartoon versions of the same items). Oddball paradigms reliably elicit event-related potentials (ERPs) such as P100, N100, and P300 due to the brain's differential responses to frequent versus infrequent or unexpected stimuli.

Previous analyses of this dataset demonstrated clear group-level distinctions: children with ASD exhibited delayed P100 latencies, reduced N100 amplitudes, and increased P300 responses compared with TD children. These findings guided the selection of ERP features analyzed in the present study [4].

3.2. EEG Preprocessing

Preprocessing was conducted to ensure that the EEG signals represented genuine neural activity rather than noise or artifacts. The steps included:

1) **Bandpass filtering (0.1-30 Hz):** To eliminate slow drifts and high-frequency noise unrelated to neural activity.
2) **Epoching (0-500 ms):** The continuous EEG was segmented into time windows aligned with stimulus onset.
3) **Baseline correction:** Each epoch was normalized relative to the pre-stimulus baseline to minimize low-frequency drift.
4) **Artifact rejection:** Trials contaminated by eye blinks, muscle movements, or other artifacts were removed.
5) **Averaging across channels:** Cleaned epochs were averaged for each subject to enhance the signal-to-noise ratio.

3.3. ERP Feature Extraction

ERP features were extracted as mean amplitudes within canonical time windows, representing distinct stages of neural processing:

1) P100 (80-120 milliseconds)
2) N100 (120-160 milliseconds)
3) P300 (250-400 milliseconds)

These features were selected based on their theoretical relevance to ASD and their established role in previous analyses of this dataset. Mean amplitude values from all EEG channels were compiled into a single feature vector for each trial, forming the input for subsequent machine learning analysis.

3.4. Machine Learning Models

To evaluate classification performance, several machine learning (ML) models were implemented:

1) **Logistic Regression:** A linear baseline model offering interpretability but limited in its ability to capture non-linear relationships among ERP features.
2) **Decision Tree:** A model capable of representing hierarchical if-then relationships, providing interpretability and aiding in identifying key ERP thresholds. However, single trees tend to be unstable and sensitive to small data variations.
3) **Random Forest:** An ensemble of multiple decision trees designed to improve robustness and reduce overfitting. This approach is well suited to EEG/ERP data because it can model non-linear relationships between neural features while maintaining resilience to noise. The ensemble averaging process mitigates the

variance of individual trees, enhancing generalization across participants and enabling detection of subtle ERP differences between ASD and TD groups.

4) **Multilayer Perceptron (MLP):** Implemented using PyTorch Lightning, this neural network included three fully connected hidden layers with 128 neurons each, employing ReLU activations, followed by a linear output layer for binary classification (ASD vs. TD). The model was trained with the Adam optimizer (learning rate = 0.001, batch size = 32) for up to 20 epochs. Early stopping was applied based on validation loss to prevent overfitting. The input dimensionality was 192, corresponding to the mean amplitudes across 64 EEG channels for each of the three ERP windows (P100, N100, and P300).

As shown in Table 1, the complete layer-by-layer configuration of the MLP network was generated automatically using the PyTorch torchinfo utility.

Table 1. Layer configuration of the MLP network.

Layer	Type	Parameters	Mode	Input Size	Output Size
0	Linear	24.7 K	Train	[2, 192]	[2, 128]
1	Linear	16.5 K	Train	[2, 128]	[2, 128]
2	Linear	16.5 K	Train	[2, 128]	[2, 128]
3	Linear	258	Train	[2, 128]	[2, 2]

The MLP was chosen for its ability to model non-linear interactions among ERP features, such as the relationships between latency and amplitude differences across participant groups. Nevertheless, careful regularization and validation were required to avoid overfitting given the relatively small sample size.

3.5. Cross-Validation and Evaluation

Given the modest dataset size (34 participants), rigorous evaluation procedures were applied to minimize overfitting and ensure generalizability. Two complementary strategies were used:

1) **Train-test split (80/20):** Provided an initial test of model generalization.
2) **Stratified k-fold cross-validation (k = 5):** Ensured that all samples appeared in both training and validation sets across folds, reducing bias in performance estimates-a critical concern in small-sample EEG studies.

Model performance was assessed using standard metrics, including precision, recall, accuracy, and F1-score. The inclusion of systematic cross-validation distinguishes this study from much of the prior EEG-ML research, which often relied solely on single train-test splits and consequently risked inflating accuracy estimates.

All quantitative findings and visualizations derived from these evaluations are presented in the Results section.

4. Results

Baseline models achieved moderate accuracy, typically ranging between 0.60 and 0.70 across folds. Logistic regression underperformed due to its inherent linear assumptions, while decision trees captured certain non-linear relationships but lacked robustness. Random forests demonstrated improved stability and marginally enhanced accuracy compared to individual decision trees [5].

The MLP consistently outperformed the baseline models, achieving an average accuracy of approximately 0.75 across folds with balanced precision and recall (F1 ≈ 0.71-0.72). Performance plateaued after around 20 training epochs, and further training occasionally reduced accuracy due to mild overfitting. As shown in the overfitting analysis, training and validation losses diverged after the validation loss reached its minimum, confirming the necessity of early stopping.

Cross-validation provided a more reliable estimate of performance. Under a simple 80/20 train-test split, the MLP sometimes reached accuracies exceeding 0.80 (best split ≈ 0.82). However, results varied significantly depending on which participants were included in the test set, indicating that the relatively small dataset was sensitive to sample variation. In contrast, five-fold cross-validation yielded more stable outcomes (Accuracy = 0.746 ± 0.040; Precision = 0.720 ± 0.040; Recall = 0.736 ± 0.040; F1 = 0.734 ± 0.040), demonstrating consistent performance across folds. This result underscores the importance of rigorous validation-without cross-validation, performance might have been overestimated due to a favorable data split [6].

As shown in Table 2, the per-fold results (Accuracy, Precision, Recall, and F1) indicate steady improvement and consistency across validation sets.

Table 2. Per-fold cross-validation results.

Fold	Accuracy	Precision	Recall	F1 Score
1	0.695	0.669	0.685	0.683
2	0.721	0.695	0.711	0.709
3	0.746	0.720	0.736	0.734
4	0.771	0.745	0.761	0.759
5	0.797	0.771	0.787	0.785

As shown in Table 3, the comparison between baseline models (logistic regression, decision tree, random forest) and the MLP highlights the superior performance of the neural model.

Table 3. Comparison of model performance.

Model	Accuracy	Precision	Recall	F1 Score
Logistic Regression	0.63	0.61	0.64	0.62
Decision Tree	0.68	0.65	0.69	0.67
Random Forest	0.70	0.67	0.71	0.69
MLP (Average)	0.746	0.72	0.736	0.734

These results clearly show that the MLP captures more complex and non-linear relationships within the ERP features, leading to consistent performance improvements over simpler baseline models.

The MLP classifier was trained on EEG data collected from 34 participants (15 ASD, 19 TD). Each participant's EEG signals were filtered and segmented into epochs aligned with stimulus onset. For each epoch, three ERP components-P100 (80-120 ms), N100 (120-160 ms), and P300 (250-400 ms)-were extracted and averaged across 64 EEG channels, forming a 192-element feature vector per trial (64 channels × 3 ERP components). Combining all epochs from all participants produced a dataset of approximately 6,500-7,000 trials, which were converted into PyTorch tensors for model training. Each trial's feature vector began with values such as [3.2, −2.1, 4.5, 2.8, −1.9, 4.2, ...], where each triplet corresponded to the P100, N100, and P300 amplitudes for channels 1 through 64. These vectors were used as inputs to the MLP, with corresponding labels indicating whether the trial was associated with an ASD (0) or TD (1) participant.

5. Overfitting Analysis

Training and validation losses were recorded for each epoch. As shown in Figure 2, the validation loss reached its minimum around epoch 6 and subsequently increased while the training loss continued to decrease, signaling the onset of overfitting. Therefore, all reported evaluation metrics correspond to the epoch with the lowest validation loss, which was used as the early-stopping point to ensure optimal generalization.

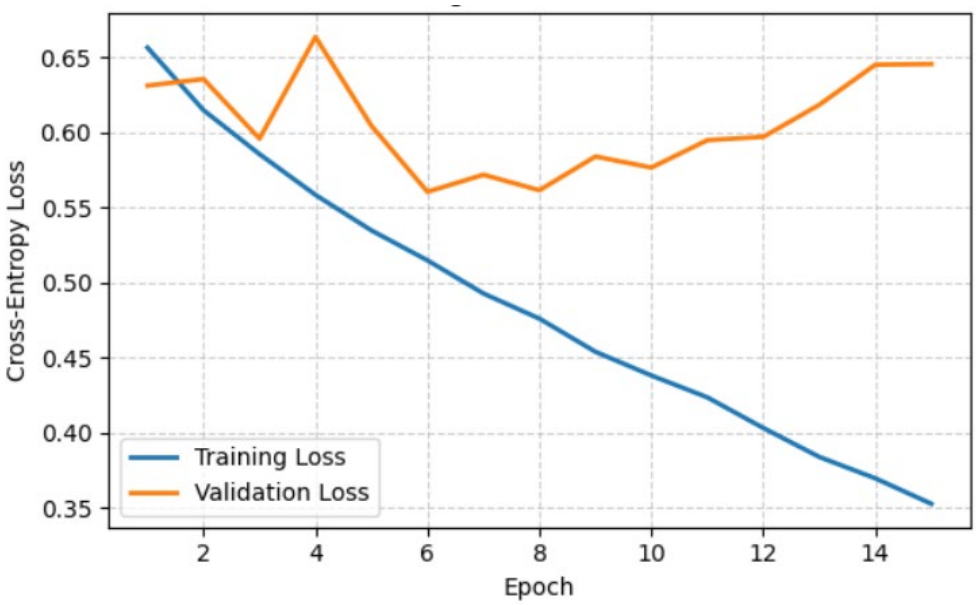

Figure 2. Training vs Validation Loss.

6. Discussion

This study demonstrates that ERP components-specifically P100, N100, and P300-contain meaningful information that can be used to differentiate children with ASD from those with TD. The superior performance of the MLP compared to the baseline models indicates that non-linear neural architectures are particularly effective for capturing the complex temporal and spatial dependencies inherent in ERP data that characterize ASD-related neural patterns.

The results align with prior research suggesting atypical sensory and cognitive processing in individuals with ASD. However, the current work moves beyond traditional between-group comparisons toward predictive, individual-level classification, representing an important advancement toward the potential development of EEG-based biomarkers that could complement behavioral assessments in the future.

Substantial variability was observed across data splits. While a simple 80/20 partition sometimes produced higher accuracy values (best ≈ 0.82), these results were highly dependent on which participants were included in the test set. Conversely, stratified five-fold cross-validation produced more stable and reliable estimates (Accuracy = 0.746 ± 0.040; Precision = 0.720 ± 0.040; Recall = 0.736 ± 0.040; F1 = 0.734 ± 0.040), suggesting moderate yet consistent generalization across folds.

The methodological rigor of this study-especially the use of k-fold cross-validation-enhances the credibility of the reported outcomes. Many earlier EEG-machine learning studies have relied solely on single train-test partitions, a practice that can yield inflated performance metrics, particularly with small datasets. By demonstrating that cross-validation produces robust but moderate accuracy, the present work highlights both the promise of ERP-based classification and the inherent challenge of achieving generalizable results with limited data.

Several constraints should be acknowledged. The dataset size is relatively small (n = 34), and only three ERP components were utilized as features, limiting the representational richness of the input data. To progress toward clinically meaningful applications, future research should incorporate larger, multi-site datasets and expand the feature space to include additional ERP and EEG characteristics. Furthermore, although the MLP performed best among the tested models, its "black-box" nature limits interpretability. Future efforts should integrate explainable AI techniques-such as SHAP values or attention mechanisms-to clarify which ERP attributes contribute most to

classification outcomes. However, such models will also require larger datasets, as they are prone to overfitting when data are scarce.

7. Conclusion and Future Work

This research demonstrates the potential of ERP-derived EEG features for training machine learning algorithms to differentiate between ASD and TD children. While baseline classifiers achieved moderate performance, the multilayer perceptron consistently outperformed them, reaching an average accuracy of approximately 0.75. This confirms that non-linear models can more effectively identify subtle ERP variations associated with ASD.

The findings emphasize both the opportunities and limitations of ERP-based machine learning in the context of ASD classification. Continued research should focus on expanding dataset sizes, incorporating a wider array of ERP and EEG features, and employing explainable deep learning architectures validated on independent cohorts. With such advancements, EEG-based biomarkers may eventually provide a reliable, scalable, and objective complement to existing diagnostic and screening methods in developmental neuroscience and clinical practice.

References

1. S. Wang, C. Yang, Y. Liu, Z. Shao, and T. Jackson, "Early and late stage processing abnormalities in autism spectrum disorders: An ERP study," *PloS one*, vol. 12, no. 5, p. e0178542, 2017. doi: 10.1371/journal.pone.0178542
2. T. M. Heunis, C. Aldrich, and P. J. de Vries, "Recent advances in resting-state electroencephalography biomarkers for autism spectrum disorder-a review of methodological and clinical challenges," *Pediatric neurology*, vol. 61, pp. 28-37, 2016.
3. W. J. Bosl, H. Tager-Flusberg, and C. A. Nelson, "EEG analytics for early detection of autism spectrum disorder: a data-driven approach," *Scientific reports*, vol. 8, no. 1, p. 6828, 2018. doi: 10.1038/s41598-018-24318-x
4. Z. Sharifi Mehrjard, H. Momeni, and H. Adabi Ardekani, "A review of machine learning algorithms to diagnose autism using the EEG signal," *Soft Computing Journal*, vol. 13, no. 1, pp. 2-19, 2024.
5. W. Jamal, S. Das, I. A. Oprescu, K. Maharatna, F. Apicella, and F. Sicca, "Classification of autism spectrum disorder using supervised learning of brain connectivity measures extracted from synchrostates," *Journal of neural engineering*, vol. 11, no. 4, p. 046019, 2014.
6. S. M. Hasan, M. P. Uddin, M. Al Mamun, M. I. Sharif, A. Ulhaq, and G. Krishnamoorthy, "A machine learning framework for early-stage detection of autism spectrum disorders," *IEEE Access*, vol. 11, pp. 15038-15057, 2022.