IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

Retinal Photographs Improve the Diagnosis of Autism Spectrum Disorder

Seyoung Park

The Webb School, 1175 West Baseline Road, Claremont, CA, 91711, USA

ABSTRACT: Autism spectrum disorder (ASD) is a neurological and developmental condition that affects behavior, communication, and learning, often requiring lifelong support. Diagnosis by the age of two years can significantly reduce symptom severity and enhance cognitive, language, and social skills. However, current diagnostic methods rely heavily on subjective behavioral observations, rendering them prone to inaccuracies, stressful for caregivers, and time-consuming. To address this issue, this study introduces a novel and objective diagnostic system that utilizes retinal (fundus) photographs in conjunction with machine learning. The fast gradient sign method (FGSM), originally developed as an adversarial perturbation technique, was applied in this study to evaluate the robustness of convolutional neural networks in classifying ASD from retinal images. This robustness test also resulted in modest performance improvements across all tested models, surpassing baseline performances. These findings could aid the development of efficient, accurate, and non-invasive tools for early ASD detection and intervention, thereby significantly benefiting individuals with ASD and their families. Future studies should investigate additional adversarial methods and incorporate larger and more diverse datasets.

KEYWORDS: Behavioral and Social Sciences, Neuroscience, Autism Spectrum Disorder, Retinal Photographs, Fast Gradient Sign Method.

Introduction

Autism spectrum disorder (ASD) is a neurological and developmental condition that affects behavior, learning, and communication. It comprises a wide variety of types and severities among patients. ASD is a lifelong disorder that requires ongoing management, although medication and treatments can lessen its severity.[1] Affecting one in 36 children,[2] ASD occurs across all ages, sexes, and ethnicities, rendering early screening highly recommended. Early interventions for ASD, ideally at the age of two or younger, can significantly improve cognitive, language, and social interaction abilities.[3]

Traditional ASD diagnosis is a two-step process that involves healthcare providers, caregivers, and children. Wellness check-ups and visits to healthcare providers help identify symptoms of ASD. Children with abnormal birth conditions or a family history of ASD undergo more thorough screening. Although ASD diagnosis is considered an accurate process, it involves a long-term examination, assessment, and conversations that can distress both caregivers and children. As it is complex, these subjective evaluations are not suitable for all cases. To address this problem, a previous study examined the use of machine learning models with retinal or fundus photographs to screen for ASD and evaluate symptom severity, demonstrating a correlation between optic disc features and ASD diagnosis.[4] As shown in Figure 1 Kim *et al*. evaluated the model performance by progressively removing 5% of the fundus photographs that were considered the least important to observe the change in the area under the receiver operating characteristic curve (AUC-ROC). Surprisingly, even when 95% of the images were removed, no significant change was observed in the AUC -ROC. However, when the area with the optic disc was masked, the AUC -ROC decreased abruptly.

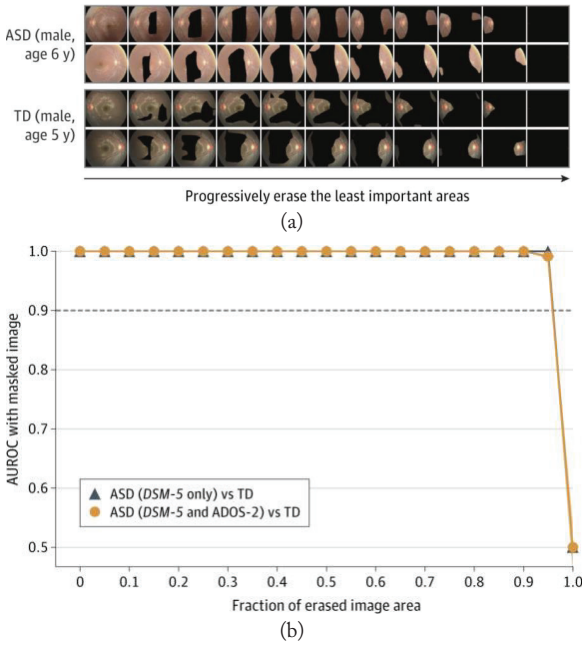

Figure 1: Previous studies on screening for ASD using fundus photographs. (a) Example of the data collection process; (b) Graph representing the AUC -ROC as the masked area of the image increases.[4]

Kim *et al* . demonstrated the feasibility of using machine learning and fundus photographs, particularly of the optic disc, to diagnose ASD. Nevertheless, their study had limitations, as the model was evaluated using only 1,890 eyes from 958 participants, which could be considered a relatively small dataset.

DOI: 10.36838/v8i2.13

Additionally, the human bias and subjectivity introduced through the use of handcrafted features demonstrated the inconsistency in conducting this experiment. Recent studies have also emphasized both the opportunities and challenges of early ASD screening. For example, Okoye *et al*. reviewed the clinical benefits and risks of early ASD diagnosis, while Kim *et al*. demonstrated the feasibility of applying deep learning to retinal images.[3,4] Furthermore, disparities in ASD diagnosis related to race and socioeconomic status have been documented, underscoring the need for objective and widely applicable diagnostic methods.[5-18] The present study provides a broader effort to improve accuracy, equity, and efficiency in ASD detection. In particular, the study developed a machine learning and mathematics-based adversarial technology that effectively used medical information. It focused on applying fundus photographs, particularly those of the optic discs, in systematic and mathematical approaches for ASD diagnosis.

ASD:

Individuals with ASD often experience difficult expressing themselves verbally and may rely on nonverbal body language.[1] Due to its detrimental effect on humans, ASD screening at a young age and early diagnosis are crucial in preventing severe impairment. Despite this need, the average age of ASD diagnosis in the United States is five years, even though ASD can be reliably diagnosed by specialists at age two.[5]

Traditional ASD diagnosis involves a thorough evaluation process, which includes collecting a developmental history from parents or caregivers, observing the child's behavior, and using standardized screening tools such as the Modified Checklist for Autism in Toddlers (M -CHAT). Professionals apply the DSM-5 criteria and administer assessments, such as the Autism Diagnostic Observation Schedule and the Autism Diagnostic Interview (ADI) . A team of professionals conducts an evaluation comprising parental interviews, developmental testing, and, if necessary, hearing tests, vision screening, and genetic testing. Throughout the process, continuous monitoring is provided to refine and adjust as needed.[6] Hence, diagnosing ASD is a longlong-term process that can be exhaust ing for both caregivers and children. Additionally, the assessment is subjective and does not guarantee complete accuracy.

(a) (b) (c)

Figure 2: Three scenes portrayportraying ASD diagnoses. (a) Hearing test for ASD screening[7]; (b) ASD screening with a therapist[8]; (c) A toddler undergoing ASD screening.[9]

A previous study demonstrated the potential of using fundus photographs for accurate and objective diagnosis of ASD severity. The calculated AUC-ROC values were 1.00 with a 95% CI for ASD screening and 0.74 with a 95% CI for symptom severity, indica ting that the model was highly reliable.[4] These results demonstrate the importance of accessible, time-efficient, and objective ASD screening and diagnosis.

Fundus Photographs:

Fundus photographs, also known as retinal photographs, show the fundus located at the back of the human eye. It comprises the retina, macula, fovea, optic nerve, and optic disc (Figure 3a). Fundus photography is easily performed in ophthalmology institutes using a fundus camera, which is a non -invasive, painless device. Colored fundus images are obtained and examined to determine the presence of diseases and disorders.[10] A recent device, depicted in Figure 3c, demonstrates a method for taking fundus photographs at home using a cell phone. These new devices, which have made fundus photography more accessible, and machine learning technology, offer a non-invasive approach for observation and diagnosis at a low cost in a flexible environment.

(a) (b) (c)

Figure 3: Fundus photograph. (a) Example of a colored fundus photograph[11]; (b) Traditional method of observing the eye in an ophthalmology institute[12]; (c) Mobile phone-based fundus imaging device.[13]

Adversarial Perturbation:

Adversarial perturbation is a crucial concept in machine learning, originally developed to test the robustness and vulnerability of neural networks by introducing small, imperceptible noise to the input data. Such perturbations, though invisible to the human eye, can lead to significant misclassifications, thereby exposing the limitations of neural network models.[14] Rather than serving as a traditional augmentation technique, adversarial perturbation is designed to challenge models under controlled distortions, enabling the evaluation of model stability and generalization. Among various adversarial methods, this study focused on the fast gradient sign method (FGSM), applying it as a robustness-oriented experiment to assess how convolutional neural networks respond to perturbations in retinal images.[15]

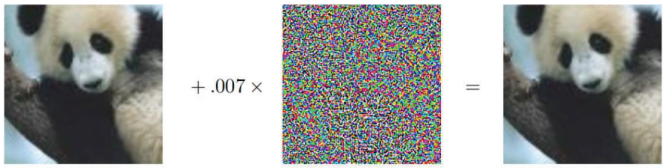

Figure 4: Adversarial perturbation.[1] Minimal adversarial noise (0.007 magnitudes) significantly shifts the model prediction from panda (57.7%) to gibbon (99.3%) illustrating the vulnerability of convolutional neural networks (CNN).

The FGSM utilizes the gradients of the loss function with respect to the input data to determine the perturbations applied to the input. For instance, the neural network identifies

the image as that of a panda with a confidence level of 57.7% (Figure 4) . However, when a small amount of noise (denoted as 0.007 times a specific color pattern) is added by calculating the gradients of the loss function of the input image, it can effectively mislead the model into classifying the image ina ccurately. The model identifies the perturbation pattern as a nematode with 8.2% confidence, which is irrelevant, but illustrates the additional randomness of the adversarial pattern. Although the resulting image appears to be a panda to human eyes, it is classified as a gibbon with 99.3% confidence, indicating a misclassification by the neural network and a significant shift caused by a slight alteration in the input image.[16]

Methods

This study applied FGSM, a commonly used adversarial perturbation method, to add noise that disturbs the learning process. Three novel methods were explored: additivity of the FGSM attack on the fundus photograph, additivity of the FGSM attack solely on the optic nerve head of the photograph, and complete removal of the optic nerve head from the photograph. These methods were examined to estimate changes in the accuracy of the method when adversarial perturbation was added , as well as the role of the optic nerve head in ASD diagnosis.

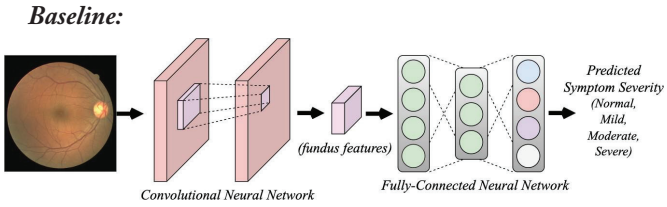

Figure 5: Baseline convolutional neural network (CNN) architecture for ASD severity classification.

Figure 5 illustrates the basic process for predicting symptom severity using the baseline model. This is the basic architecture of the classification network used in this study. The network uses a fundus photograph as an input $I \in RHW$ and generates feature maps. H and W denote the height and width of the fundus photograph, respectively. Fundus features, the output of a convolutional neural network, are represented as a three-dimensional matrix denoted by z € ⟦R⟧ ^l. This leads to the FCNN, which outputs different probability values for each of the four possibilities: normal, mild, moderate, and severe. As illustrated by the different colors in Figure 5, each element has a score value, which can be altered into a probability using the Softmax function. This probability represents the model's prediction of the possibility of this symptom range. This process can be defined as: $FCNN: Z \rightarrow P$.

Equation 1. Softmax function.

$$P_k = \frac{e^{S_k}}{\sum_j e^{S_j}}$$

Equation 1 illustrates the Softmax function, which converts a set of raw scores into probabilities that are easier to interpret and work with when utilizing machine learning. Pk is the output or the probability assigned to class k, Sk is the score for k, and $\sum_j e^{s^j}$ It is the sum of the exponentials of all the raw scores. By exponentiating each score, the equation checks all outputs. Normalizing these values by dividing by the sum of all the exponentials of the scores ensures that output probabilities, when added up, equal one.

Equation 2. Cross-entropy loss function.

$$L_{ce} = -log_e P$$

Equation 2 presents the cross-entropy loss function, which evaluates a model's performance by comparing its predicted probability distribution with the actual distribution. $\boldsymbol{L_{ce}}$ is the output of the cross-entropy loss or the probability value between 0 and 1, where $\boldsymbol{log_e}$ represents the natural logarithm, and P is the predicted probability of the correct class. Specifically, the loss value is quantified by taking the negative logarithm of the predicted value, thereby minimizing this loss value and improving the model's ability to make accurate predictions. A loss value closer to one indicates a lower loss, whereas a loss value closer to zero indicates a higher loss.

Proposed Noise Model (Fundus):

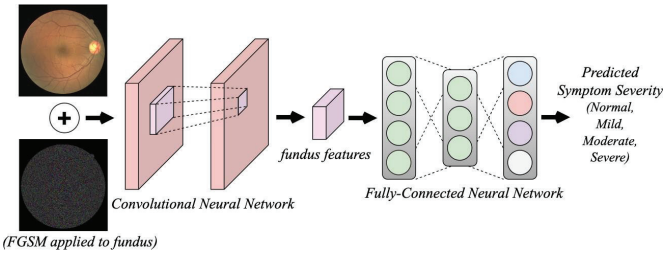

Figure 6: Noise model with fast gradient sign method (FGSM) applied to entire fundus images.

Figure 6 illustrates the architecture of the first additive proposed in this study for classifying ASD symptom severity. All processes were identical to the baseline architecture, except for the input of the image, which included an FGSM-applied fundus photograph. The FGSM is a picture comprising small dots of color, which makes no difference in how a human views the photo; however, it renders machine learning more challenging for computers. The noise value was denoted as $\boldsymbol{N_{total}}$. This FGSM attack is mathematically constructed by reverse-engineering a typical gradient-descent algorithm. A typical gradient descent algorithm iteratively uses input and gradient values to produce a better optimized result through extensive calculations. Instead of using the gradient descent algorithm to increase our output value positively, the gradient values are included in $\boldsymbol{N_{total}}$ to make training more difficult. To improve the results, the noise value increases in every sample.

DOI: 10.36838/v8i2.13

Proposed Noise Model (Optic Nerve Head):

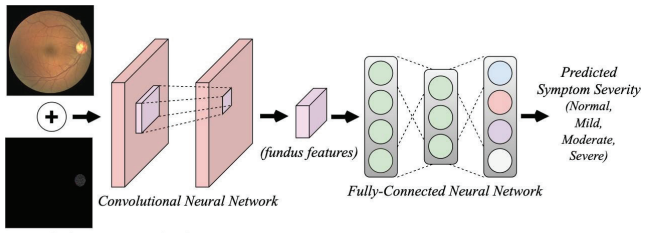

Figure 7: Noise model with FGSM applied only to the optic disc region.

Figure 7 shows the architecture of the second additive proposed in this study, in which noise was added to the optic nerve head area. Similar to the first additivity, all processes are identical except for the input, which is a picture with noise or FGSM applied solely to the optic disc of the fundus photograph. The input is denoted by N_{disc}. As the optic disc is a crucial part of ASD diagnosis, it can be hypothesized that the accuracy does not increase by adding noise to this specific system. However, this experiment further investigated whether the accuracy would be maintained by adding noise to the optic disc specifically, rather than the entire fundus photograph, and, if not, the rate of decrease in comparison to the proposed noise model in Equation 2. This experiment also used a cross-entropy loss function.

Proposed Noise Model (Optic Nerve Head Removed):

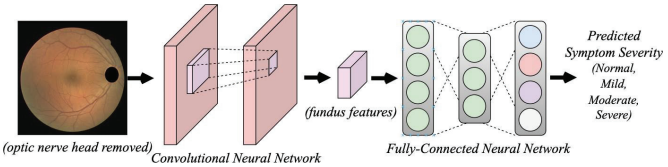

Figure 8: Noise model with the optic disc region removed from fundus images.

Figure 8 illustrates the architecture of the last model, in which the optic nerve head has been removed from the photograph. The process is identical; however, the input differs because the optic nerve head is completely removed. This input is referred to as I_{nodisc}. It is created through simple coding by changing the existing pixel values of the optic disc to black, given the disc area.

Fundus Dataset:

This study used a dataset from the AI Hub, a government-funded database in Korea, to conduct experiments.[17] The most recently updated version of the data dated January 19, 2024 was utilized. From 1,038,674 samples representing various diagnoses of disorders in children and adolescents, 57,195 samples consisting solely of fundus photographs from children and adolescents with ASD and those without any diagnosed disorders were collected.

Although the samples were exclusively from South Korea and comprised data on South Koreans, this should not affect the accuracy of the experiment, as ASD is not correlated with a specific ethnicity.[18] Among the 57,195 samples, 37,145 (64.9%) were normal and consisted of fundus photographs of children and adolescents without the disorder, and 20,050 (35.1%) were samples of children and adolescents with ASD. In terms of age distribution, 27.70% of the samples were from children under seven years of age, 43.57% were from children aged 7–12 years, and 28.73% were from adolescents aged 13–20 years. Additionally, the ratio of the samples used for training and testing was 8:2.

Result and Discussion

Evaluation of the FGSM Applied Model:

Table 1: Evaluation result of the FGSM applied model (fundus): a comparison of performance metrics (accuracy, recall, precision, and F1-score) across four CNN architectures under adversarial perturbation applied to entire fundus images.

FGSM (Fundus)	Accuracy	Recall	Precision	F1-Score
ConvNeXt[19]	86.70	86.12	85.97	86.04
DenseNet-201[20]	87.55	87.12	86.60	86.86
ResNet-101[21]	87.93	87.34	86.55	86.94
ResNet-152[21]	89.11	88.58	87.82	88.20

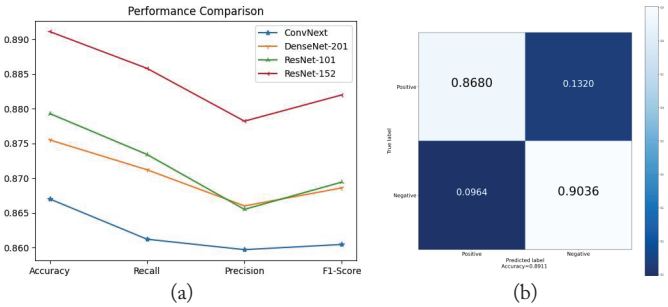

Figure 9: Evaluation results. (a) Comparison of accuracy, recall, precision, and F1-score across four CNN architectures with FGSM applied to the entire fundus image, showing that ResNet-152 outperforms the others in all metrics. (b) Confusion matrix for ResNet-152 under FGSM perturbation, achieving 89.11% accuracy. High true-positive and true-negative rates indicate strong robustness to adversarial noise.

Figures 9a and 9b, along with Table 1, present the performance comparison graph, confusion matrix, and summary table, respectively, offering insights into the experimental results of evaluating the performance of various CNNs, particularly focusing on how they handle adversarial perturbations (FGSM) in the symptom severity assessment of ASD through fundus photographs. Figure 9a shows the performance metrics (accuracy, recall, precision, and F1-score) for the four CNN architectures: ConvNeXt, DenseNet-201, ResNet-101, and ResNet-152. ResNet-152, which is a deep network with 152 layers, outperformed the other models across all metrics, indicating that it was the most effective model for this task.

ConvNeXt, with 50 layers, showed the lowest performance, particularly in terms of recall (a measure of true positives from all positive samples) and precision (a measure of positive predictions), suggesting that it may not be as reliable for correctly identifying both positive and negative cases. The table in Figure 9a provides a detailed breakdown of the performance comparison graph for each model, reinforcing the observation that ConvNeXt lags behind the other models, presumably because of its shallow layers, which hinder complex studies with FGSM.

Thus, the findings suggest that machine learning models, particularly deep-learning CNNs, can serve as powerful tools for ASD screening and symptom severity assessment by accurately analyzing retinal images. Spesifically, the confusion matrix shown in Figure 9b indicates an overall accuracy of 89.11%. The gradation scale measures accuracy, with dark blue indicating the least accurate and light blue indicating the most accurate. Visually, the true-positive (correctly identified ASD patients using fundus photographs) and true-negative (correctly identified non-ASD patients using fundus photographs) rates were low. In contrast, there were false-positive results (identified as non-ASD patients with ASD). The false-negative rates (identifying ASD patients as non-ASD) are low, indicating that the model correctly identifies the majority of the given dataset.

Therefore, the application of FGSM should primarily be interpreted as a robustness test, assessing the stability of CNN models under perturbation. The observed improvements in accuracy suggest that, beyond withstanding adversarial noise, the models demonstrated enhanced generalization. This reframing highlights FGSM's role in testing model robustness rather than serving as a conventional data augmentation method. These results confirm that deep-learning CNNs, particularly ResNet-152 among the ones tested, are effective tools for ASD screening and symptom assessment.

Optic Disc Area Applied:

Table 2: Comparison of model accuracy when the FGSM perturbation is applied only to the optic disc area versus baseline without perturbation. All four CNN architectures showed accurate improvements, with ResNet-152 achieving the highest increase (3.06%). These results suggest that localized perturbation to the optic disc can enhance model performance.

FGSM (Optic disc)	FGSM (Optic Disc; Accuracy)	Baseline (Accuracy)
ConvNeXt	86.5	83.54
DenseNet-201	87.54	84.18
ResNet-101	87.82	84.34
ResNet-152	88.89	85.83

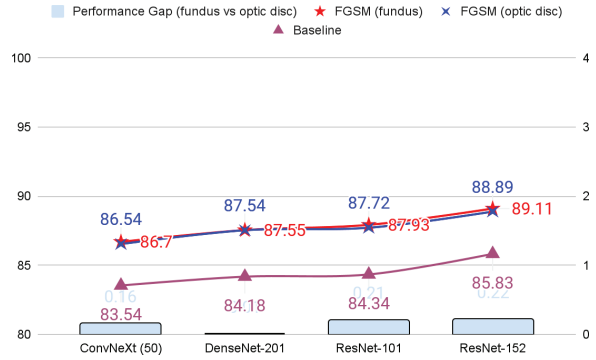

Figure 10: Ablation study results. Accuracy comparison across four CNN architectures for baseline, FGSM applied to the entire fundus, and FGSM applied only to the optic disc. Both FGSM conditions improved accuracy over baseline, with full-fundus FGSM showing slightly higher gains. ResNet-152 achieved the highest accuracy in all settings, indicating strong robustness to perturbation.

This ablation study compared the effects of applying FGSM solely to the optic disc area with a baseline (without FGSM) model, using a different approach to define the impact of the optic disc on diagnosing ASD symptom severity. It also compares the accuracies of four different neural networks in this architecture. Applying the FGSM to the optic disc resulted in performance improvements for all tested models. Specifically, ConvNeXt, which was the least accurate when tested in baseline architecture, showed a performance increase of 2.96, whereas ResNet-152 maintained its high accuracy with an increase of 3.06. These improvements highlight the efficacy of the FGSM in enhancing model performance by effectively preprocessing the input image. The graph further emphasizes the performance gap when the FGSM was added to the baseline model, highlighting that applying the FGSM to the fundus photograph was more accurate than applying it solely to the optic disc. However, for DenseNet-201, the performance gap was 0.01, indicating that applying the FGSM to optic discs or full fundus photography did not make a noticeable difference in diagnosing the severity of ASD symptoms.

Evaluation of Optic Nerve Head Removal Model:

Table 3: Accuracy comparison across four CNN architectures for baseline and when the optic nerve head is entirely removed from fundus images. The removal of the optic disc resulted in a substantial performance drop for all models, with ConvNeXt showing the most significant decrease (-9.07%) and ResNet-101 showing the smallest (-6.67%). These results highlight the critical role of optic disc information in ASD severity classification from retinal images.

	Optic Nerve Head Removed (Accuracy)	Baseline (Accuracy)
ConvNeXt	74.47	83.54
DenseNet-201	76.35	84.18
ResNet-101	77.67	84.34
ResNet-152	77.98	85.83

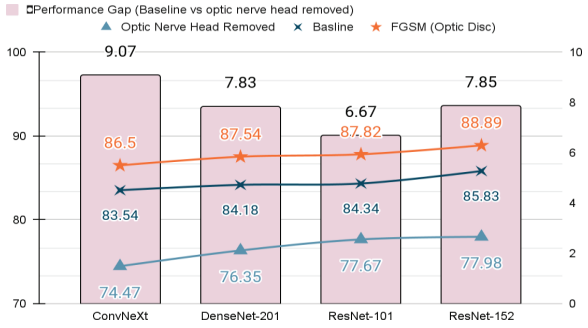

Figure 11: Evaluation result of optic nerve head removal experiment: accuracy comparison across CNN architectures, highlighting the performance drop when the optic nerve head is removed versus the baseline and FGSM (optic disc) conditions.

The second ablation study focused on the performance of various neural network architectures when the optic nerve head was removed entirely from the fundus photographs. As shown in Table 3, the accuracy of the different models in diagnosing the symptom severity of ASD decreased drastically to 70% when the optic nerve head was removed from the input image, compared to the baseline (the original study without FGSM applied). ConvNeXt, with 50 layers, showed the most significant performance drop of 9.07, indicating a high dependency on optic nerve head information. Simultaneously, ResNet-101

DOI: 10.36838/v8i2.13

experienced the smallest performance drop of 6.67, suggesting greater robustness in removing this feature. Overall, all neural network architectures experienced a significant decline in performance, underscoring the importance of the optic nerve head in medical imaging tasks, particularly in the diagnosis of ASD.

Figure 11 summarizes the performance gap between the baseline, optic nerve head removed, and FGSM applied to the input images. The FGSM results consistently showed a higher performance than both the baseline and the removed optic nerve head, suggesting that FGSM may be a more effective preprocessing method for enhancing model performance. Among the various models, ResNet-152 demonstrated the highest accuracy in evaluating fundus photographs for both the baseline and FGSM, underscoring the significance of the depth of the neural network in its performance. These insights provide a step-ahead solution for accurately diagnosing symptom severity in ASD, which is valuable for future model design and selection in medical imaging applications.

Conclusion

This study proposed and evaluated a novel system to diagnose ASD symptom severity using fundus photographs, focusing on the optic disc and the FGSM. This study applied FGSM as a robust-oriented perturbation technique to the entire fundus photograph and specifically to the optic disc, demonstrating that adversarial perturbation enhanced smodel performance. Furthermore, this study systematically the significance of the optic disc by comparing the accuracy of ASD diagnosis following its complete removal. The findings revealed that applying the FGSM to the optic disc significantly improve diagnostic accuracy across multiple neural network architectures, surpassing the baseline performance. Performance noticeably declined when the optic disc was removed entirely, underscoring the critical role of the optic disc in medical imaging tasks. Moreover, an analysis of various models showed that the deeper layers of feature maps were correlated with performance accuracy. Overall, findings could help develop robust, effective, and non-invasive diagnostic tools for ASD, thereby improving early detection and intervention strategies.

Despite these promising results, this study has several limitations. First, the dataset was limited to pediatric and adolescent fundus images from South Korea. Further validation on more diverse, multi-ethnic cohorts is necessary. Second, while FGSM perturbations were useful as a robustness test, they represent only one type of adversarial approach; future research should explore additional techniques such as PGD (Projected Gradient Descent) or DeepFool. Finally, this study focused exclusively on retinal imaging. Future studies integrating multimodal data (e.g., genetic, behavioral, or linguistic features) could enhance diagnostic performance. Addressing these limitations will be essential to ensure the clinical applicability of this approach.

References

1. National Institute of Mental Health. *Autism Spectrum Disorder*. https://www.nimh.nih.gov/health/topics/autism-spectrum-disorders-asd
2. Centers for Disease Control and Prevention. *Data and Statistics on Autism Spectrum Disorder*. https://www.cdc.gov/autism/data-research/index.html
3. Okoye, C.; Obialo-Ibeawuchi, C. M.; Obajeun, O. A.; Sarwar, S.; Tawfik, C.; Waleed, M. S.; Wasim, A. U.; Mohamoud, I.; Afolayan, A. Y.; Mbaezue, R. N. Early diagnosis of autism spectrum disorder: a review and analysis of the risks and benefits. *Cureus* **2023**, 15 (8), e43226. DOI: 10.7759/cureus 43226
4. Kim, J. H.; Hong, J.; Choi, H.; Kang, H. G.; Yoon, S.; Hwang, J. Y.; Park, Y. R.; Cheon, K.-A. Development of deep ensembles to screen for autism and symptom severity using retinal photographs. *JAMA Network Open* **2023**, *6* (12), e2347692–e2347692. DOI: 10.1001/jamanetworkopen.2023.47692
5. Autism Speaks. *Autism Statistics and Facts*. https://www.autismspeaks.org/autism-statistics-asd
6. NYU Langone Health. *Diagnosing Autism Spectrum Disorder in Children*. https://nyulangone.org/conditions/autism-spectrum-disorder-in-children/diagnosis
7. Special Learning. *Hearing Evaluation for Children with Autism*. https://special-learning.com/hearing-evaluation-for-children-with-autism
8. Waypoint Wellness Center. *Autism Screenings*. https://www.waypointwellnesscenter.com/services/autism-screenings
9. Otto, F. Toddler screening essential for autism detection despite the national task force's reservations. *Drexel News*, February 22, 2016. https://drexel.edu/news/archive/2016/february/early_autism_screening_rec
10. Dhanorkar, A. What is the purpose of fundus photography? *MedicineNet*. www.medicinenet.com/what_is_the_purpose_of_fundus_photography/article.htm
11. Carver College of Medicine, Department of Ophthalmology and Visual Sciences. *Diagnostic Imaging Services*. https://eye.medicine.uiowa.edu/patient-care/imaging-services
12. The University of British Columbia. *Color Fundus Photography*. https://ophthalmology.med.ubc.ca/patient-care/ophthalmic-photography/color-fundus-photography
13. RO Staff. A 'smart' way to get into fundus photography. *Review of Optometry*, January 23, 2021. https://www.reviewofoptometry.com/article/jumpstart-healing-with-new-amniotic-membrane-graft-1
14. Nightfall. *Adversarial Attacks and Perturbations*. https://www.nightfall.ai/ai-security-101/adversarial-attacks-and-perturbations
15. Goodfellow, I. J.; Shlens, J.; Szegedy, C. Explaining and harnessing adversarial examples. *arXiv* **2014**, *1412.6572*. DOI: 10.48550/arXiv.1412.6572
16. Knagg, O. Know your enemy: Why adversarial examples are more important than you realize. *Medium*, January 2, 2019. https://medium.com/data-science/know-your-enemy-the-fascinating-implications-of-adversarial-examples-5936bccb24af
17. AI Hub. *Fundus Image Data for Symptoms of Mental Illness in Children and Adolescents*. https://aihub.or.kr/aihubdata/data/view.do?currMenu=115&topMenu=100&aihubDataSe=data&dataSetSn=71516
18. Durkin, M. S.; Maenner, M. J.; Baio, J.; Christensen, D.; Daniels, J.; Fitzgerald, R.; Imm, P.; Lee, L.-C.; Schieve, L. A.; Van Naarden Braun, K. Autism spectrum disorder among US children (2002–2010): socioeconomic, racial, and ethnic disparities. *American Journal of Public Health*, 2017, 107(11), 1818–1826. DOI: 10.2105/AJPH.2017.304032
19. Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; Xie, S. A Convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, LA, USA, June 18–24, 2022; IEEE: New York, NY, USA, 2022; pp 11976–11986. DOI: 10.1109/CVPR52688.2022.01167

20. Huang, G.; Liu, Z.; Van Der Maaten, L.; Weinberger, K. Q. Densely Connected Convolutional Networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, HI, USA, July 21–26, 2017; IEEE: New York, NY, USA, 2017; pp 4700–4708. DOI: 10.1109/CVPR.2017.243
21. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, USA, June 27–30, 2016; IEEE: New York, NY, USA, 2016; pp 770–778. DOI: 10.1109/CVPR.2016.90

Author

Seyoung Park attends The Webb School in Claremont, CA. She recently won first place at the California Science Fair (Behavior Medicine) for the research presented in this paper. She has previously published research on medical imaging classification using machine learning. She aspires to study science to impact human health positively.

DOI: 10.36838/v8i2.13