Indian Journal of Computer Science and Technology
Special Issue, Volume 5, Issue 1 (January-April 2026), PP: 146-151.
https://www.doi.org/10.59256/indjcst.20260501C024
International Conference on Cyber and AI Security (ICCAIS-2026)
www.indjcst.com
ISSN No: 2583-5300

Enhancing Early Alzheimer's Disease Detection Using GAN-Based MRI Data Augmentation

Lakshmi Priya[1], J. Jerin Jose[2], S. Dinesh[3], G. Ramya Sri[4], A. Bhagya Ratna Sri[5]
[1, 2,3,4,5] *Department of Artificial Intelligence and Machine Learning, Sasi Institute of Technology and Engineering, Tadepalligudem, Andhra Pradesh, India.*

To Cite this Article: *Lakshmi Priya[1], J. Jerin Jose[2], S. Dinesh[3], G. Ramya Sri[4], A. Bhagya Ratna Sri[5], "Enhancing Early Alzheimer's Disease Detection Using GAN-Based MRI Data Augmentation, Indian Journal of Computer Science and Technology Volume 05, Issue 01 (January-April 2026), PP: 146-151.*

Abstract: *Diagnosing Alzheimer's disease early on is still not easy due to two main issues: the lack of medical imaging data and the serious class imbalance between patients with and without Alzheimer's disease at each stage of the disease. Typically, deep learning models do not work well with small or skewed datasets, therefore we believe there is a solution to both of these problems using Generative Adversarial Networks (GANs) to create synthetic MRI images as solutions. We trained a GAN using the actual MRI images of brains from patients diagnosed with Alzheimer's, which allowed us to generate realistic synthetic MRI images (especially for hard-to-find patients in the early stages of Alzheimer's). By mixing the AI-generated images into the real-world data set of MRI images, we significantly balanced the training dataset for our classifier. After training with the newly balanced dataset, there was significant improvement in the classifier's performance in terms of accuracy, recall, and F1-scores. In clinical terms, our biggest improvement was early stage Alzheimer's detection. Our work demonstrates that GAN- generated synthetic medical imaging has the potential to provide a viable method for diagnosing rare diseases that lack sufficient imaging data to support conventional diagnostic methods.*

Key Words: *Alzheimer's disease, GAN, synthetic data, MRI augmentation, deep learning, early detection, class imbalance.*

I. INTRODUCTION

Alzheimer's disease (AD) is a debilitating disease that ultimately deprives individuals of their ability to function independently by gradually stripping them of their cogni- tive function and memory. AD is among the most common causes of dementia globally, creating an enormous burden on healthcare agencies and families affected by the disease. In fact, there appears to be consensus among professionals in the medical community as to the critical importance ofdetecting Alzheimer's disease in the early stages; an early diagnosis may provide opportunities for better treatment, in- tervention, and time spent with a loved one who has been diagnosed. However, it is extremely difficult to detect early- stage Alzheimer's disease since signs of the disease at this point are subtle, and we have little data to provide benchmarks for assessment.

An MRI is an instrumental diagnostic tool for detecting Alzheimer's disease as demonstrated by the significant body of knowledge generated by MRI scans. MRI scans give researchers and clinicians access to information about brain structure, including peeling of the gray matter cortex and atrophy of the hippocampus among others, which are indica- tors of the presence of Alzheimer's disease. Most researchers will need to utilize an automated means of analyzing and interpreting MRI scan data in order to maximize the infor- mation derived from MRI scan results and to speed up the identification of people with Alzheimer's disease. Fortunately, a number of machine learning/deep learning algorithms have been developed by researchers to analyze and interpret MRI scans. However, two primary challenges facing researchers using these algorithms are that they: (1) require substantial amounts of high quality (balanced) training data; and (2) the amount of training data that exist for early onset (or mild) Alzheimer's disease patients is limited.

Enter the world of Generative Artificial Intelligence (AI). Generative adversarial networks (GANs) provide an innovative and efficient means of creating synthetic (illusionary) images based on standard medical imaging (specifically MRI) data so that researchers may both increase their datasets' diversityand alleviate the class imbalance (i.e., disparity between the quantity of images in each MRI class) issue without needing to acquire thousands (or possibly millions) of new images from additional patients. In this study, we will be focused specifically on using GANs to create synthetic (illusions of) images of the brain using standard MRI techniques to assist in the classification of patients with Alzheimer's disease (AD) who may have been excluded from other studies based on access to the standard image-based classification process because they have characteristics/features that fall into the underrepresented categories of AD.

To accomplish this, we will combine GAN-generated syn- thetic (illusionary) images with real (actual) images of the brain to create a fully balanced (equal number of patients with AD at each stage of the disease) training set that can therefore be used to train our classifier to detect (diagnose) each stage of AD. Through this process, we aim to deter- mine whether this novel use of Generative AI will provide us with an improved ability to diagnose patients with AD, especially those patients that present with findings indicative of the earliest possible stages of AD because current methods are focused on determining whether patients already have a diagnosis of dementia (regardless of what type). We believe that this innovative application of Generative AI will provide an opportunity for the AI-based diagnosis of/for patients presenting with undetermined or inconclusive results or in cases where the number of traditional MRI image scans is limited.

A. Related Work: What Others Have Tried

Several research teams have explored generative models and deep learning for medical data augmentation. Let's see what they've accomplished and where gaps remain.

The foundational work came from Goodfellow and col- leagues back in 2014 [1]. They introduced GANs as a groundbreaking framework for generating synthetic data. Their work established the theoretical foundation for everything that followed, though they focused on generic image datasets rather than medical applications.

Shin et al. [2] demonstrated that GANs could synthesize medical images across different modalities—MRI, CT, you name it. They showed that GAN-generated images could effectively augment datasets, but their work wasn't disease- specific.

Bowles and his team [3] applied GAN-based augmen- tation specifically to brain MRI scans. Their results were promising—synthetic images did improve model performance. However, they were primarily interested in segmentation tasks rather than disease classification.

Suk et al. [4] took a different approach, using CNNs to classify Alzheimer's disease from MRI data. They achieved decent results, but they were clearly limited by insufficient training data and class imbalance issues.

Payan and Montana [5] built a 3D CNN model specifically for Alzheimer's diagnosis using structural MRI. This was pioneering work in applying deep learning to this problem, but again, the small dataset size restricted how well their model could generalize.

Frid-Adar's group [6] combined GAN-generated synthetic data with CNN classifiers for medical image classification. They saw performance improvements, but their approach wasn't tailored for Alzheimer's specifically.

More recently, Gupta et al. [7] investigated combining GAN-based augmentation with CNNs for Alzheimer's classification. Their results showed that synthetic MRI data improved accuracy, but they didn't thoroughly validate the approach for early-stage cases—which is where the real clinical need lies. What makes our work different? We're specifically targeting early-stage Alzheimer's detection in MRI scans. We're directly addressing the class imbalance problem and aiming to improve diagnostic accuracy where it matters most—in those initial, subtle stages of the disease. Our emphasis is on creating balanced datasets and achieving better model performance
when real medical data is scarce.

B. Workflow of the Proposed System

The workflow defines how the GAN-based method for data augmentation improves classification accuracy for Alzheimer's Disease in accordance with a project schedule. The first step is to collect all related MR for patients with Alzheimer's Disease and healthy patients from accessible public databases.

The next step involves preprocessing the collected MR im- ages by normalizing pixel intensities, removing noise, resizing each to the same size, labelling images according to their associated pathologies (i.e., stage and type of Alzheimer's Disease), and determining the presence of class imbalance in the overall database. In the preprocessing stage, images that are inappropriate for analysis may be removed manually, but a large proportion of these images will be retained in the final database for analysis as part of the convolution depth of the CNN.

The third step involves training a GAN (or DCGAN) based on real patient MR images. Once trained, the GAN will be able to generate synthetic patient MR images of underrepresented classes of patients with Alzheimer's Disease. The newly created synthetic MR images may also be verified to ensure they follow a similar structure and appear consistent with the real patient images.

Post-verification of the synthetic patient MR images, the synthetic images are added to the original data set to produce a more complete, representative training data set. Subsequently, the provided training data set is utilized to train a deep learning based classification model (e.g. CNN), which will categorize each patient and determine their respective stage of Alzheimer's Disease.

II.OUR METHODOLOGY

A. Dataset Collection and Preprocessing

Publically available magnetic resonance imaging (MRI) scans and images of the brains of Alzheimer's Disease (AD) patients who have undergone therapy in the past few years will form part of this research. Each MRI containing AD brain images was acquired from previously created global databases of medical images previously approved according to international standards.

For this research, the MRI scans acquired were prepared for further analysis through sequential steps including;

- The images were prepared to standardize them (normal- ization).
- Each image was assessed regarding the quality of the indi- vidual scan, with poor-quality and corrupted images excluded from future analyses.
- The images remained high quality and therefore were included in the training set once the data had been processed,

determining whether the participants had ’AD’.

After completing the image processing, a trial and evalua- tion process was utilized to create three distinct groups of data i.e. (1) training subset; (2) validation subset; (3) test subset.

B. Analyzing the Class Imbalance Problem

From our investigation into the statistical data, we have found that there are a great number of individuals who are currently in a very early stage of cognitive impairment com- pared to what we consider to be ”normal” models. Because of this imbalance, it will not be possible to create deep learning classifiers designed to accurately detect early-stage Alzheimer’s Disease. Additionally, models will receive very little information about early-stage individuals as they are so sparsely distributed compared to larger classes due to the fact that most of them will experience drift towards the major classes. Fig. 1 illustrates the workflow of the proposed system.

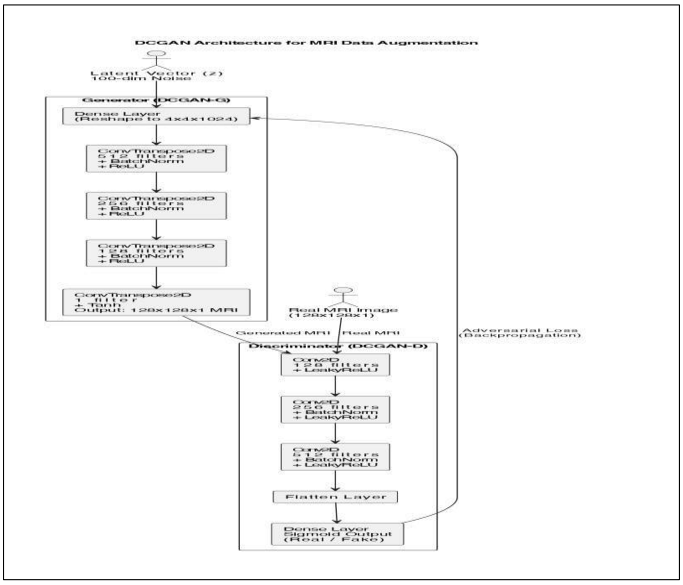

Fig. 1. Workflow of the proposed GAN-based Alzheimer’s disease diagnosis Framework.

C. GAN Architecture Design

We utilized a Deep Convolutional GAN (DCGAN) archi- tecture for our generative model, where two neural networks are used in an adversarial gaming environment (a generator and a discriminator). The generator receives random noise and outputs a synthetic MRI scan that looks as much like a real one as possible. The discriminator evaluates images and attempts to determine if they are real or synthetic by looking at a database of real MRI scans. The mechanism is similar to a forger and an expert art appraiser — in other words, each neural network pushes the other to become more proficient in their respective domains.

Through this adversarial training process, both the generator and the discriminator experience a continuous progressive improvement in performing their respective roles. As the generator creates more and more realistic representations of real-world data, the discriminator gets better and better at identifying counterfeit data sources. Eventually, the generator creates synthetic MRI images that are so much like actual MRI scans that the discriminator can no longer consistently distinguish between them.

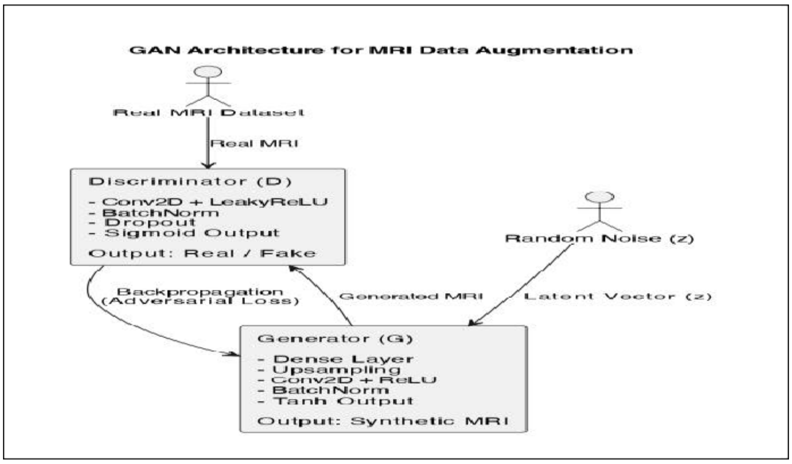

Fig. 2. Architecture of the Generative Adversarial Network (GAN)

D. DCGAN Architecture

DCGAN works as a successful enhancement for creating synthetic MRIs from scratch using DCGAN to improve the quality of generated MRI images as well as provide stability in their generation. DCGAN is an advanced method utilizing the basic generative adversarial network (GAN) architecture for use with images. Instead of fully-connected layers, which are necessary to build the basic GAN, DCGAN uses convolutional layers and transposed convolutional layers to identify both the spatial and structural characteristics of the images.

To create a synthetic MRI from a noise vector, the DCGAN generates a synthetic MRI by using a series of transposed convolutional layers to progressively convert a random noise vector into an MRI image. Each of these transposed convo- lutional layers will "up-sample" the spatial attributes of the generated MRI image, while increasing the number of associ- ated data points associated with the generated MRI image, thus learning what attributes characterize the spatial relationships of the generated MRI image. Batch normalization is applied after the activation of each of the layers in the DCGAN, which assists in stabilizing the training of the DCGAN and helps speed up the time it takes to train the DCGAN. The activation functions commonly used in the generator are ReLU; however, the activation function most often used in the output layer to create the normalized pixel value of the produced MRI image is Tanh.

The discriminator of the DCGAN utilizes a series of con- volutional layers to extract different levels of hierarchical features from MRIs. When provided with real MRIs (i.e., from a magnetic resonance imaging machine) or synthetic.

MRIs generated by the DCGAN pre-trained generator, the discriminator will learn the distinguishing features between the two types of MRIs. The most common activation function used in the discriminator is leaky ReLU, which eliminates the "vanishing gradients" problem associated with the gradient descent method of training the DCGAN. Fig. 3 illustrates the architecture of the proposed DCGAN model.

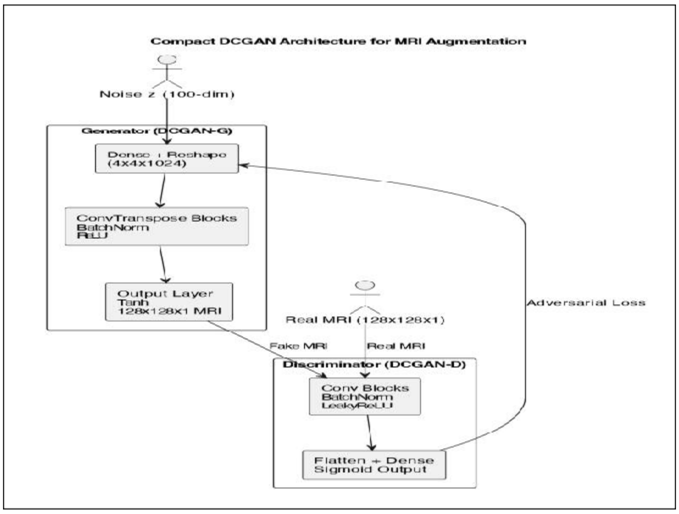

Fig. 3. Architecture of the Deep Convolutional Generative Adversarial Network (DCGAN).

E. Network Configuration Details

To create the proposed method, transposed convolutional layers, batch normalization, and ReLU activations were uti- lized with the generator network. The layers progressively upsampled the random noise input to produce a full-resolution MRI image at the other end of the generator network (i.e., the resulting image).

We used usual convolutional layers (with leaky ReLU activations) combined with dropout for regularization on the discriminator side. We did not want to see any signs of overfitting when training either of the two networks.

Both networks were trained together using the Adam op- timizer. Determining the correct learning rates was necessary to maximize learning and required some experimentation on our part. Using a loss function that integrated the adversarial nature of each network together guided both networks towards converging on producing test images that looked very similar to actual MRI scans.

F. Generating Synthetic Images

Once we had trained the GAN (generative adversarial network), we used the generator to create synthetic MRI im- ages. We targeted under-represented categories of Alzheimer's disease to correct our class imbalance issue. The generator will take in a random noise vector as input and will output a synthetic brain scan. We produced multiple images for each under-represented category, based on both the severity of the class imbalance and the desired final distribution of classes.

G. Quality Validation of Generated Images

It is an unwise assumption that a synthetic image is "good enough" to use without validation. We had domain experts visually inspect synthetic images to determine whether or not they appropriately depicted the anatomy of the brain, with no visible artifacts. If the generated images do not have brain structures which resemble real brain structures, then those images cannot be used.

We used additional "quantitative" measures in addition to the domain expert visual review to validate synthetic images. Based on Inception Score and Fre´chet Inception Distance measurements, we were able to establish objective image quality/quantity measures for both synthetic and real images. We also performed statistical comparisons of the characteristics of real versus synthetic images to ensure that they shared similar characteristics. If all verification measures were passed, then the synthetic image could be used.

H. Dataset Augmentation Strategy

The combination of synthetic MRIs from a previous trial and our data yielded a larger diversity of disease types and allowed us to present both our data and thesynthetic MRI dataset in an equitable manner. The end result resulted in a better classifica¬tion model as a result of the variances in the number of cases between the synthetic train¬ing dataset and the original training dataset for the various categories or classes represented in the datasets. In addition, the synthetic training dataset also provided an equable num¬ber of cases representing the various categories/classes than it had prior to augment¬ing its training dataset prior to producing it.

I. Training the Classification Model

The classification task for diagnosing Alzheimer's disease involved using a deep convolutional neural network (CNN) architecture with several convolutional blocks that each con- tained images, Batch normalization processing, and Max pool- ing operations, followed by several fully connected output layers that were activated using Softmax functions to generate multi-class predictions.

We performed supervised training of the CNN model using data augmentation (e.g., random rotation, flipping, scaling etc.) to introduce additional regularization into the training set and minimize overfitting. For model training, we employed categorical cross-entropy as our loss function.

J. Performance Evaluation

To assess the performance of our trained models, we com- pared them against a test dataset that was not used to train the models. We have recorded performance results using common evaluation metrics - accuracy, precision, recall and F1 score for each disease category. The performance of early-stage Alzheimer's disease detection was of greatest interest to us as it represents our primary clinical concern.

We performed a model comparison between a model trained with only the original dataset and a model trained using our augmented dataset. This comparison would indicate whether or not the use of a GAN-based method of generating synthetic data will yield meaningful results.

III.RESULTS AND DISCUSSION

In this section, weReport our findings from the experiment using a synthetically-generated augmentation using GANs. The goal of this research was very simple, does adding new synthetic MRI-images generated from GANs to an imbalanced set of small original MRI-images correlate to better perfor- mance in classifying Alzheimer's disease?

A. Evaluation Metrics Explained

We evaluated our classification model using standard perfor- mance metrics. Accuracy gives us the overall correct prediction rate. Precision tells us what percentage of our positive disease predictions were actually correct. Recall measures how many actual disease cases we successfully caught. The F1-score balances precision and recall together. In medical diagnosis, these metrics are crucial because missing a disease case can have serious consequences for patients.

B. Baseline Performance: No Augmentation

First, we established a baseline by training our classi- fier using only the original MRI dataset—no augmentation whatsoever. Given the dataset's small size and heavy class imbalance, the model struggled to learn good features for early-stage Alzheimer's. While it performed reasonably well on majority classes, recall and F1-scores for minority classes were disappointing. This clearly showed the model was having trouble detecting early-stage cases.

C. Enhanced Performance: With GAN Augmentation

The addition of synthetic MRI images produced by GANs into the training database led to a fundamental change in our ability to build a model that would perform significantly better than previous models. The new data provided us with the opportunity to obtain a more balanced representation of all diagnostic categories resulting in a larger number of features to learn from. All metrics used to evaluate the model (i.e., accuracy, precision, recall, and F1-scores) improved. The most unexpected improvement was seen with respect to recall (i.e., detecting previously untreated cases of Alzheimer's disease).

D. What These Results Mean

When examining whether augmented or non-augmented models perform better, there are clear advantages with augmentation. When evaluated on augmented data, the model generalizes with less bias toward the majority class.

The benefits of better detecting early-stage Alzheimer's are particularly evident since they have relevance to clinical practice. This shows that synthetic data from generative AI can be used to improve accuracy in diagnosing diseases by automating the process with reliable information.

Our experiments support the conclusion that using gener- ative AI as a method for augmenting medical data will lead to better classification accuracy in situations where there is limited availability of medical records. The results of the study show that generating synthetic magnetic resonance images (MRI) using GAN technology is a feasible and effective technique for diagnosing rare diseases.

IV.CONCLUSIONS AND FUTURE DIRECTIONS

In this paper, we studied how to use Generative Adversarial Networks (GANs) to reduce the problem created by limited MRI image data for diagnosing Alzheimer's Disease. The synthetic images generated by GANs are realistic enough that they were able to help address both the problems of limited data and classifiers' imbalanced dataset when diag- nosing patients. Based on the experiments we performed, the accuracy, precision, recall, and F1-Score of our generated image results were significantly better than those reported using traditional measurement techniques. Additionally, this approach resulted in improved detection of early-stage disease and increased reliability of models. Thus our evidence suggests that using GANs to augment the data may be a practical way to increase the accuracy of medical imaging based diagnosis of Alzheimer's Disease.

There are also a number of future research directions to explore. For example, other alternative generative models could be used besides GANs for generating synthetic medical images like Variational Auto-Encoders and Diffusion Models. Another potential area of future research would be using multi- modal imaging (such as MRI and PET scans) and multi- modal learning techniques to combine results from these two different image modalities, which could help improve diagnostic accuracy. Most importantly, this technique requires clinical validation; partnering with hospitals and/or medical facilities will provide evidence on how useful the proposed technique will be.

REFERENCES

1. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," in *Proc. Advances in Neural Information Processing Systems (NeurIPS),* 2014, pp. 2672–2680.
2. H.-C. Shin, N. A. Tenenholtz, J. K. Rogers, C. G. Schwarz, M. L. Senjem, J. L. Gunter, K. Andriole, and M. Michalski, "Medical image synthesis for data augmentation and anonymization using generative adversarial networks," *arXiv preprint arXiv:1807.10225*, 2018.
3. C. Bowles, L. Chen, R. Guerrero, P. Bentley, R. Gunn, A. Hammers, D. Dickie, M. Herna´ndez, J. Wardlaw, and D. Rueckert, "GAN augmenta- tion: Augmenting training data using generative adversarial networks," *arXiv preprint arXiv:1810.10863*, 2018.
4. H.-I. Suk, S.-W. Lee, and D. Shen, "Deep learning-based feature representation for Alzheimer's disease and mild cognitive impairment classification," *IEEE Trans. Biomedical Engineering*, vol. 62, no. 12, pp. 2941–2949, 2014.
5. A. Payan and G. Montana, "Predicting Alzheimer's disease: A neu- roimaging study with 3D convolutional neural networks," in *Proc. Int. Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015, pp. 355–362.
6. M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification," *IEEE Trans. Medical Imaging*, vol. 37, no. 7, pp. 1661–1673, 2018.
7. V. Gupta, S. Jain, and S. Mishra, "GAN-based data augmentation for improved Alzheimer's disease classification using MRI images," in *Proc. Int. Conf. Artificial Intelligence and Data Engineering*, 2021, pp. 1–6.
8. M. Stone, *Dialogue Systems in Artificial Intelligence*. Oxford, U.K.: Oxford Univ. Press, 2019.
9. D. Jurafsky and J. H. Martin, *Speech and Language Processing*, 3rd ed. Pearson, 2023.
10. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. Bharath, "Generative adversarial networks: An overview," *IEEE Signal Processing Magazine*, vol. 35, no. 1, pp. 53–65, 2018.
11. X. Yi, E. Walia, and P. Babyn, "Generative adversarial network in medical imaging: A review," *Medical Image Analysis*, vol. 58, pp. 101552, 2019.
12. L. Cui, S. Biswal, L. M. Glass, G. Lever, J. Sun, and C. Xiao, "CONAN: Complementary pattern augmentation for rare disease detection," in *Proc. AAAI Conf. Artificial Intelligence*, vol. 34, no. 4, pp. 614–621, 2020.