Enhancing Diagnostic Accuracy in Medical Imaging: Integrating GAN-Based Data Augmentation for Balanced Dataset Creation

Hemanth Karnati
School of Computer Science and Engineering
Vellore Institute of Technology
Vellore, India
karnati.hemanth2020@vitstudent.ac.in

SujayKumar Reddy M
School of Computer Science and Engineering
Vellore Institute of Technology
Vellore, India
sujaykumarreddy.m2020@vitstudent.ac.in

***Abstract*—Advancements in medical imaging have been substantially driven by deep learning technologies, particularly Convolutional Neural Networks (CNNs). A critical hurdle in this domain is the imbalance of datasets, where certain medical conditions are underrepresented, leading to potential biases in diagnostic models. This research addresses the imbalance in medical imaging datasets, specifically in chest radiography, by leveraging Generative Adversarial Networks (GANs) for data augmentation. The study utilizes the ChestXray2017 dataset, which is skewed towards pneumonia cases, resulting in a dearth of normal chest X-ray images. To counter this, Deep Convolution Generative Adversarial Networks (DCGAN) were employed to generate synthetic images of normal chest X-rays, thus aiming to balance the dataset.**

In this study, we conducted a comparative analysis of a Convolutional Neural Network's (CNN) performance on a chest radiography dataset, before and after augmenting it with Deep Convolution Generative Adversarial Network (DCGAN)-generated images. Initially, the CNN trained on the un-augmented dataset achieved 93% training accuracy and 87% validation accuracy. After integrating 400 synthetic normal chest X-ray images, the training accuracy slightly increased to 95%, while the validation accuracy notably improved to 89%. This enhancement in validation accuracy demonstrates the model's improved generalization capabilities due to a more balanced training dataset. Our results indicate that GAN-based data augmentation effectively addresses class imbalances in medical imaging datasets, potentially leading to more accurate and reliable diagnostic models. However, the study also underscores the need for further research into the quality and ethical implications of using synthetic images in medical diagnostics. Overall, the integration of GAN-generated images into CNN training presents a promising method for improving classification performance in medical imaging, offering a practical approach to overcome challenges associated with data scarcity and imbalance.

***Index Terms*—Generative AI, Generative Adversarial Networks (GANs), Convolution Neural Networks (CNNs), Deep Learning.**

I. Introduction

In recent years, the landscape of medical diagnostics has been revolutionized by advancements in artificial intelligence, particularly through the application of deep learning techniques in medical imaging [1]. Among these techniques, Convolutional Neural Networks (CNNs) have emerged as a cornerstone, offering substantial improvements in the accuracy of disease detection and classification from medical images. Despite these advancements, the field faces a significant challenge: the imbalance in medical imaging datasets. This issue arises when certain medical conditions are underrepresented in available datasets, often due to their rarity or the logistical challenges in collecting sufficient data [2]. Such imbalances can lead to diagnostic models that are biased or less accurate, potentially compromising their utility in clinical decision-making.

This research paper delves into this prevalent issue within the context of chest radiography. The imbalance in medical imaging datasets, particularly in the case of chest X-rays, can be a substantial barrier to developing robust and reliable diagnostic tools. To address this, our study focuses on the application of Generative Adversarial Networks (GANs), specifically Deep Convolution Generative Adversarial Networks (DCGAN), to augment existing datasets. We utilize the ChestXray2017 dataset [3], an extensive collection of chest radiograph images, which, like many medical datasets, suffers from an imbalance—predominantly featuring pneumonia cases at the expense of 'normal' cases. Our approach involves generating synthetic but realistic images of normal chest X-rays using DCGAN, which are then integrated into the original dataset. This process aims to create a more balanced distribution of cases, potentially enhancing the accuracy and reliability of CNN-based diagnostic models.

Our research conducts an in-depth comparative analysis, examining the impact of this data augmentation on the performance of a CNN model. Initially, the model trained on the original, un-augmented dataset demonstrated solid performance metrics. However, upon integrating the DCGAN-generated synthetic images, we observed a marked improvement in both training and validation accuracies. This suggests that the augmented dataset not only offers a more balanced representation of conditions but also enhances the model's ability to generalize and accurately classify unseen data.

The results of our study highlight the significant potential of employing GAN-based synthetic data augmentation to mitigate class imbalances in medical imaging datasets. Such an approach is particularly relevant in the realm of medical diagnostics, where the acquisition of large, diverse, and balanced datasets can be inherently challenging. While our findings are promising, they also pave the way for further investigations into the quality, realism, and ethical implications of utilizing synthetic images in medical diagnostics. In summary, this paper presents a comprehensive analysis of integrating GAN-generated images into CNN training regimes, exploring its feasibility and effectiveness in enhancing the performance of medical image classification, especially in scenarios characterized by data scarcity and class imbalance.

II. Related Work

Data augmentation is a crucial aspect of training discriminative Convolutional Neural Networks (CNNs). A study by Hussain et al. (2017) compared various augmentation strategies like horizontal flips, random crops, and principal component analysis (PCA) in the context of medical imaging [4]. Their findings revealed that the effectiveness of augmentation strategies, such as flips and gaussian filters, was significantly higher compared to less effective strategies like adding noise. This work emphasized the importance of choosing the right augmentation strategy to retain the properties of the original medical images, thereby affecting both discriminative and generative performance of the models.

Shin et al. (2018) proposed a generative adversarial network-based method to create synthetic abnormal brain MRI images, addressing the challenge of imbalanced datasets in medical imaging [5]. Their approach not only improved tumor segmentation performance but also served as an anonymization tool, allowing for data sharing while maintaining privacy. This study highlights the dual benefits of synthetic images in enhancing model performance and ensuring data confidentiality.

Mikołajczyk and Grochowski (2018) explored various data augmentation methods, including classical image transformations and advanced techniques like Style Transfer and Generative Adversarial Networks [6]. Their research focused on the impact of these methods on improving deep learning algorithms, particularly in medical imaging contexts where data scarcity is a common issue. The study provided insights into the potential of combining traditional and novel augmentation methods to enhance the efficiency of deep neural networks.

Generative Adversarial Networks (GANs) have become a cornerstone in the field of medical image analysis. Their ability to synthesize images with high levels of realism offers a solution to the chronic scarcity of labeled data in medical imaging. In their review, Kazeminia et al. (2020) provide a comprehensive overview of GAN applications in medical imaging, discussing their potential in tasks such as de-noising, reconstruction, segmentation, and classification [7].

Super-resolution using GANs presents a novel approach to enhancing the quality of medical images. Gupta et al. (2020) demonstrated the effectiveness of GANs in improving the resolution of MRI scans. Their method outperformed standard deep learning models, producing higher resolution images that closely resemble the target scans [8].

Additionally, the application of multi-scale GANs for generating high-resolution medical images has shown promising results. Uzunova et al. (2019) introduced a novel approach for generating large 2D and 3D medical images. Their method addresses the computational demands typically associated with high-resolution image generation, offering a scalable and efficient solution [9].

The evolution of Convolutional Neural Networks (CNNs) has revolutionized the field of medical image analysis. In their survey, Sarvamangala and Kulkarni (2022) provide an in-depth look at the applications of CNNs in medical image understanding [10]. This survey emphasizes the effectiveness of CNNs in various tasks such as image classification, segmentation, localization, and detection, particularly in the context of medical imaging for ailments of the brain, breast, lung, and other organs.

The integration of Generative Adversarial Networks (GANs) with Convolutional Neural Networks (CNNs) has shown significant promise in medical image analysis. Frid-Adar et al. (2018) demonstrated the effectiveness of GAN-based synthetic medical image augmentation in improving CNN performance for liver lesion classification [11]. Their approach involved synthesizing high-quality liver lesion regions of interest (ROIs) using GAN architectures and subsequently enhancing the classification accuracy of a CNN model through synthetic data augmentation.

In another study, Talukdar et al. (2022) employed DCGAN and CNN transfer learning techniques for classifying medical X-ray images [12]. They observed an improvement in the accuracy of various CNN models, including custom CNN and popular architectures like InceptionV3, ResNet50, and VGG16, after incorporating GAN-generated training data.

Furthermore, Bali and Mahara (2023) compared affine and DCGAN-based data augmentation techniques for chest X-ray classification [13]. Their findings indicated that DCGAN not only outperformed traditional models in accuracy and recall but was also capable of identifying fake images with high precision, highlighting its potential in medical diagnostics.

III. Proposed Methodology

In this paper, our methodology adheres to the framework illustrated in Fig. 1, We initiate our investigation by utilizing the original dataset, where our initial focus is on assessing the presence of a class imbalance problem. To address this imbalance, we employ a training strategy involving a DCGAN. Specifically, our approach involves training the minority class with DCGAN to generate synthetic images within the same class.

The Generator trained through the DCGAN is employed to generate additional images within the same class, thereby augmenting the dataset. This augmented dataset encompasses a larger number of samples compared to the original dataset. The expanded dataset, enriched through the generative capabilities of the DCGAN, serves as a valuable resource for training

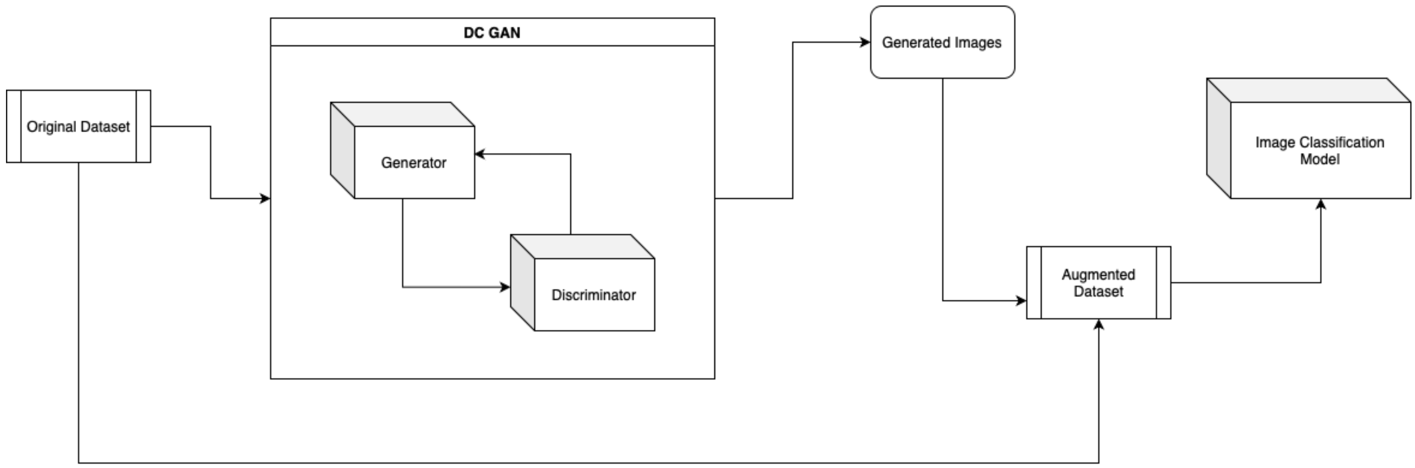

Fig. 1: Proposed Methodology

image classification models such as Convolutional Neural Networks (CNN). By incorporating these augmented samples, we aim to enhance the model's ability to learn diverse features and patterns, ultimately contributing to improved performance and generalization in the image classification task.

A. Datasets

TABLE I: Data Distribution of ChestX-ray2017 dataset

Dataset	Number of Samples
Train-Normal	1349
Train-Pneumonia	3883
Test-Normal	234
Test-Pneumonia	390

This research leverages the comprehensive ChestXray-2017 dataset, as referenced in [3]. The dataset is meticulously curated and encompasses two vital classes: Normal and Pneumonia. Within this dataset, each sample is meticulously captured through the lens of computed tomography (CT) scans, providing a detailed exploration of lung images. These CT scans serve as the foundational basis for our in-depth analysis, offering a rich and diverse set of medical imagery for investigation.

In Fig. 2 and 3, we present a visual representation of the CT images corresponding to the two distinctive labels, Normal and Pneumonia. These images offer a firsthand glimpse into the intricate details of pulmonary conditions. Notably, Fig. 2 provides a clear depiction of the CT images without colormaps, allowing for an unaltered representation of the raw data. In contrast, Fig. 3 introduces colormaps to the CT images, enhancing visual interpretation and emphasizing specific features within the lung scans. This comparison facilitates a comprehensive understanding of the dataset and contributes to the nuanced exploration of medical imaging in the context of chest radiography.

Table 1 illustrates the distribution of the dataset, revealing a notable disparity in counts between the Pneumonia and Normal classes. Here, we designate the Normal CT samples as the minority class, while considering the Pneumonia CT samples as the majority class.

B. DCGAN

Generative Adversarial Networks (GANs) [14] find extensive applications in various domains, including text, audio, and data generation, enabling the creation of novel datasets. However, when confronted with high-dimensional data such as images, GANs encounter challenges in training and struggle to extract meaningful features essential for generating realistic fake images. To address this limitation, Deep Convolutional Generative Adversarial Networks (DCGANs) [15] are employed to enhance the training process and facilitate the extraction of relevant features, particularly in the context of image generation. DCGAN (Deep Convolutional Generative Adversarial Network) is a type of generative adversarial network (GAN) that uses convolutional neural networks (CNNs) in both the generator and discriminator networks. It is a powerful generative model capable of generating high-quality images from random noise.

The generator network in DCGAN is responsible for creating fake images from random noise. It consists of a series of transposed convolutional layers, also known as deconvolutional layers, that gradually upsample the input noise to the desired image size. Each transposed convolutional layer is followed by a batch normalization layer and a ReLU activation function. The final layer is a convolutional layer with a tanh activation function, which outputs the generated image. The discriminator network in DCGAN is responsible for distinguishing between real images and fake images generated by the generator network. It consists of a series of convolutional layers, each followed by a batch normalization layer and a Leaky ReLU activation function. The final layer is a fully connected layer with a sigmoid activation function, which outputs a probability score indicating the likelihood that the input image is real.

The generator and discriminator networks in DCGAN are trained simultaneously in an adversarial manner. The generator

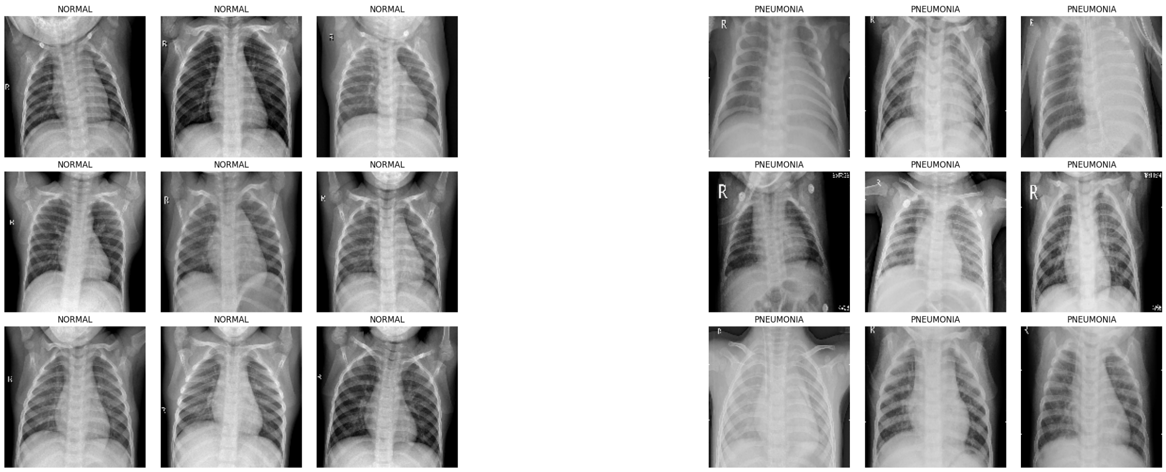

(a) Normal Lung

(b) Pneumonia Lung

Fig. 2: Data Samples for Pneumonia Classification Dataset

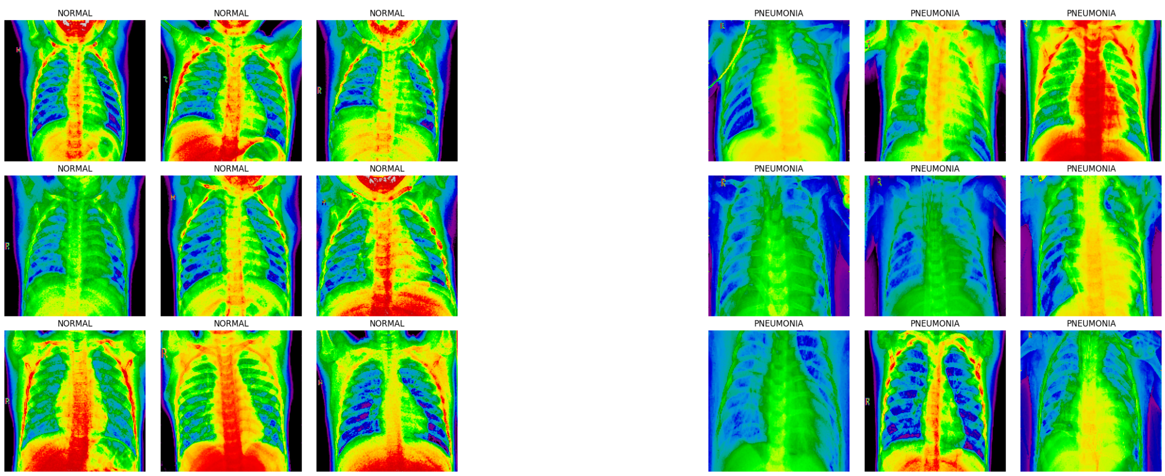

(a) Normal Lungs

(b) Pneumonia Lungs

Fig. 3: Data Samples using Colormaps (nipy_spectral) for Pneumonia Classification Dataset

network tries to generate fake images that are indistinguishable from real images, while the discriminator network tries to distinguish between real and fake images. This adversarial training process results in the generator network learning to produce high-quality images that are realistic and visually appealing.

First we update the discriminator with the real batch once the Gradients are calculated we calculate the loss D_x now, we make the generator to generate a fake batch of Images which is then passed over discriminator

Let z be a latent space vector sampled from a standard normal distribution. The generator function $G(z)$ maps z to data-space. The goal of G is to estimate the distribution that the training data comes from (p_{data}), generating samples from the estimated distribution (p_g).

The discriminator $D(G(z))$ computes the probability (scalar) that the output of the generator G is a real image. In the minimax game described by Goodfellow, the discriminator D seeks to maximize the probability of correctly classifying real and fake samples ($\log D(x)$), while the generator G aims to minimize the probability that D predicts its outputs as fake ($\log(1 - D(G(z)))$).

IV. Experimental Results and Discussion

Our study utilized the ChestXray2017 dataset [3], which consists of two classes: pneumonia and normal. We observed a significant class imbalance, with a lower number of normal chest X-ray images. To address this, we synthesized 400 normal chest X-ray images using DCGAN and integrated them into the dataset to balance the class distribution and improve the CNN's performance.

The primary focus of our study was to train the DCGAN on the minority class, i.e., Normal Images. The dynamics of this adversarial training process are illustrated in Fig. 6, where we observe the fluctuating loss values for both the generator and

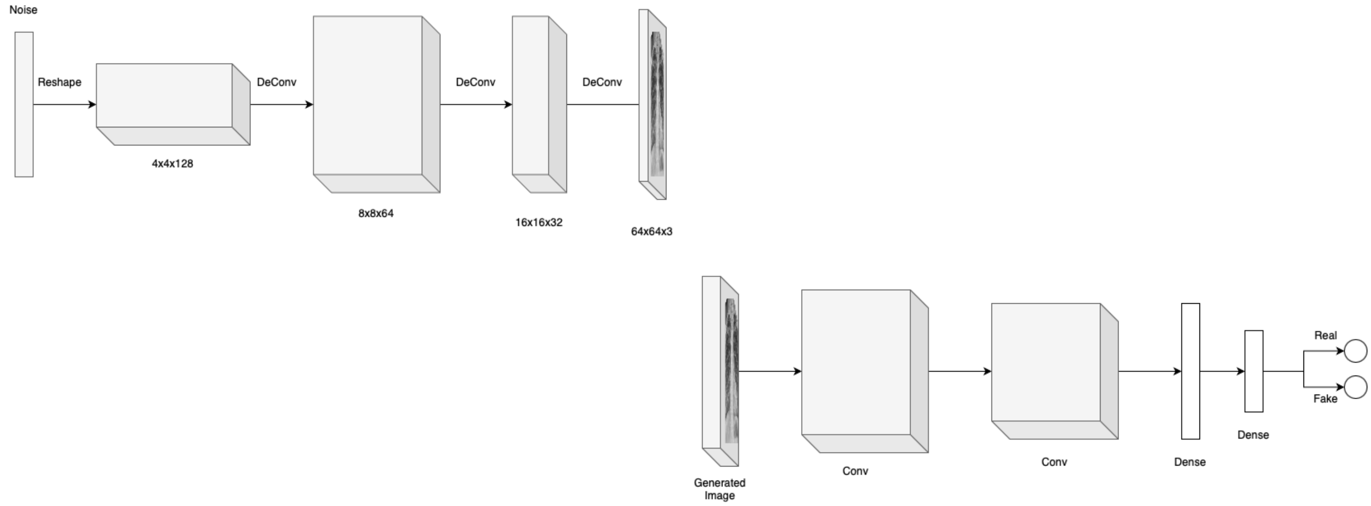

Fig. 4: Deep Convolution Generative Adversarial Network (DCGAN)

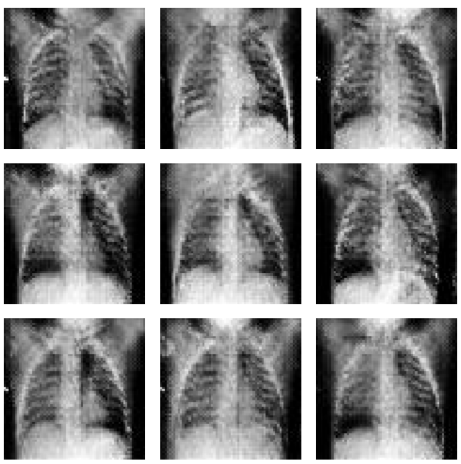

(a) Generated Normal Lungs image using GAN

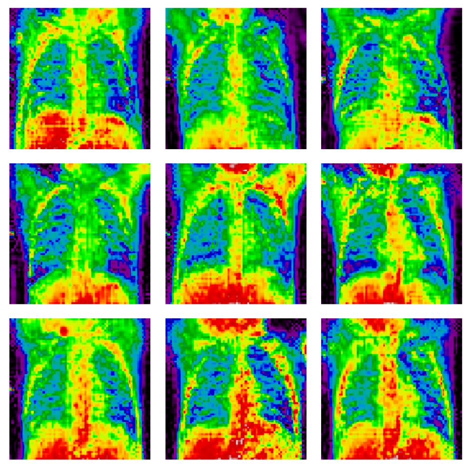

(b) Colormap of Generated Normal Lungs image using GAN

Fig. 5: Generated new Data Samples using DCGAN

discriminator. These oscillations are typical of GAN training, indicating the generator's improvement in creating realistic images and the discriminator's refinement in distinguishing real from synthetic images.

TABLE II: Comparison of Model Performance Before and After Augmentation

Metric	Before Augmentation	After Augmentation
Training Accuracy	93.2%	95%
Validation Accuracy	87.1%	89.3%
Precision	0.84	0.87
Recall	0.97	0.98

Table II illustrates the enhancements in the performance of a CNN model for medical imaging diagnostics after incorporating synthetic images via GAN-based data augmentation.

The augmentation led to notable improvements across all evaluated metrics: training accuracy increased from 93.2% to 95%, and validation accuracy rose from 87.1% to 89.3%. Additionally, precision and recall saw improvements, moving from 0.84 to 0.87 and from 0.97 to 0.98, respectively. These results underscore the effectiveness of utilizing synthetic data to balance datasets, which in turn enhances the model's accuracy in diagnosing medical conditions, reduces false positives, and improves its ability to identify true positive cases.

Fig. 7 displays the CNN model's accuracy and loss over epochs without GAN-based data augmentation, with the left plot for accuracy and the right plot for loss. Finally, Fig. 8 provides a comprehensive view of the CNN model's performance with GAN-based data augmentation.

These results demonstrate the effectiveness of GAN-based synthetic data augmentation in addressing class imbalances in medical imaging datasets. The improved validation accuracy indicates enhanced learning of discriminative features by the CNN model. However, the quality and ethical implications of using synthetic images warrant further investigation.

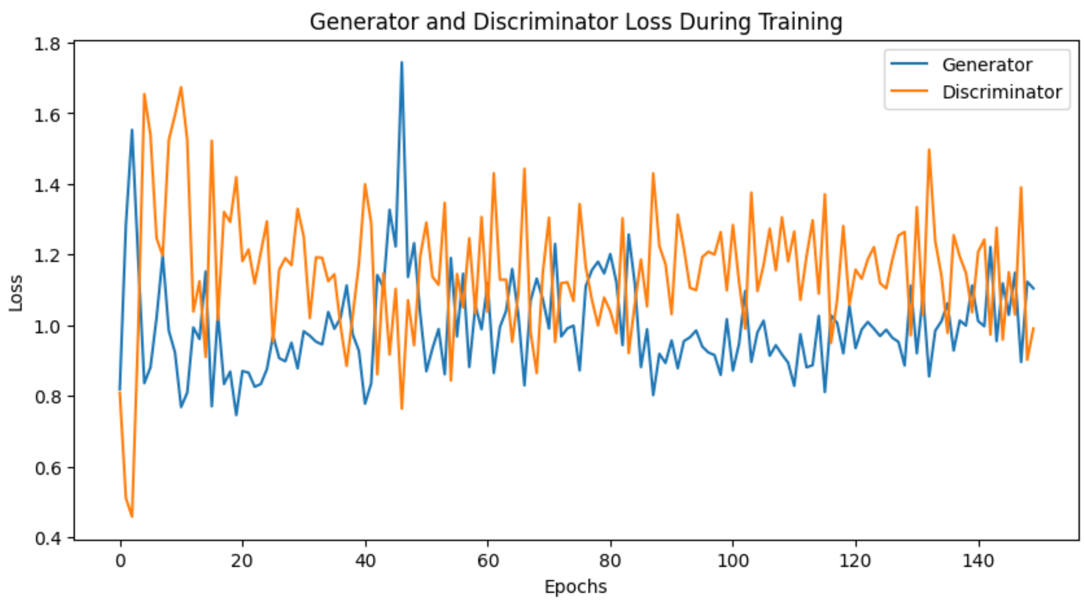

Fig. 6: Generator and Discriminator Loss During Training. This plot tracks the loss of both the generator (in blue) and the discriminator (in orange) across 150 epochs, reflecting the ongoing learning and adaptation within the GAN framework.

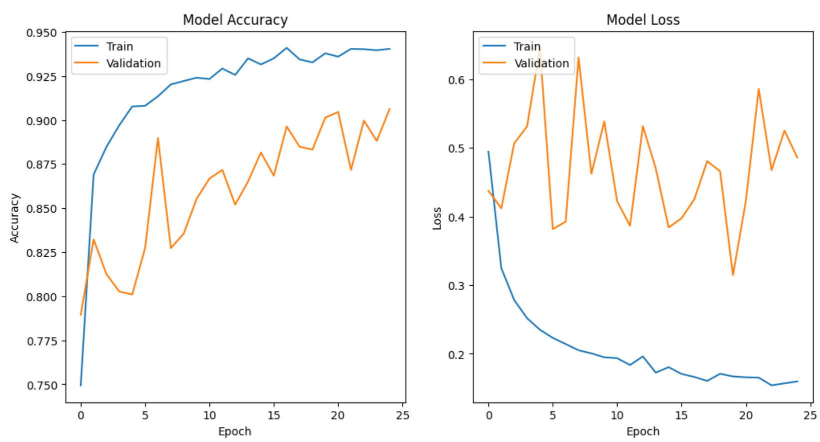

Fig. 7: Model Accuracy and Loss Over Epochs for a CNN Model Without Data Augmentation via GAN.

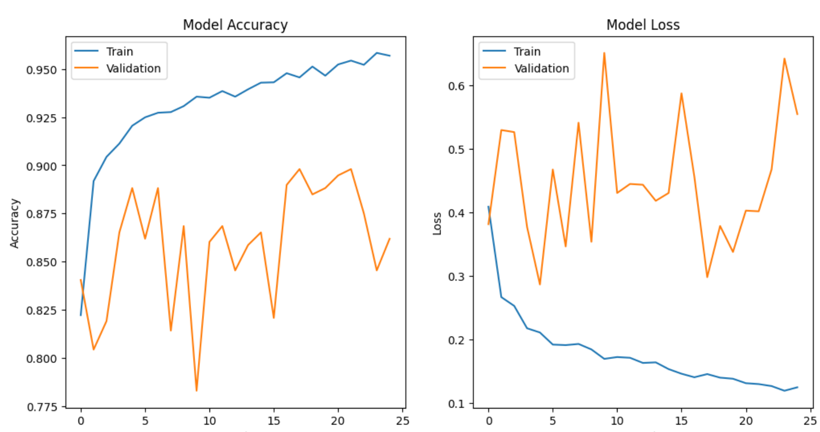

Fig. 8: Model Accuracy and Loss Over Epochs for a CNN Model With Data Augmentation via GAN.

In conclusion, integrating GAN-generated images into CNN training appears to be a viable strategy for enhancing classification performance in medical imaging, particularly in scenarios with data scarcity and imbalance.

V. Conclusion and Future Work

This research has successfully demonstrated the potential of Deep Convolution Generative Adversarial Network (DCGAN)-generated synthetic images in enhancing the performance of Convolutional Neural Networks (CNNs) for medical image classification. By incorporating artificially generated normal chest X-ray images into the ChestXray2017 dataset, we achieved a notable increase in the CNN's validation accuracy, from 87% to 89%. This enhancement highlights the viability of synthetic data augmentation as a solution to the prevalent issue of imbalanced datasets in medical imaging.

The utilization of DCGANs to generate supplementary training data represents a significant stride towards improving the generalizability and accuracy of CNN models in medical diagnostics. Addressing dataset imbalance through this method is a key step in advancing the development of reliable and precise diagnostic tools in healthcare.

Looking ahead, our research opens avenues for investigating various GAN architectures and their impact on CNN performance, along with a rigorous analysis of the quality and realism of synthetic images. Ethical considerations and the expansion of this approach to other imaging modalities, such as MRI or ultrasound, are also crucial. Conducting clinical trials will be essential for evaluating the practical effectiveness of CNN models trained with synthetic data augmentation in real-world medical diagnostics.

References

[1] S. A. Alowais, S. S. Alghamdi, N. Alsuhebany, T. Alqahtani, A. I. Alshaya, S. N. Almohareb, A. Aldairem, M. Alrashed, K. Bin Saleh, H. A. Badreldin, et al., "Revolutionizing healthcare: the role of artificial intelligence in clinical practice," *BMC Medical Education*, vol. 23, no. 1, pp. 689, 2023, Springer.

[2] L. Gao, L. Zhang, C. Liu, and S. Wu, "Handling imbalanced medical image data: A deep-learning-based one-class classification approach," *Artificial Intelligence in Medicine*, vol. 108, pp. 101935, 2020, Elsevier.

[3] D. Kermany, K. Zhang, and M. Goldbaum, "Labeled Optical Coherence Tomography (OCT) and Chest X-Ray Images for Classification," *Mendeley Data*, vol. 2, 2018. [Online]. Available: https://doi.org/10.17632/rscbjbr9sj.2

[4] Z. Hussain, F. Gimenez, D. Yi, and D. Rubin, "Differential data augmentation techniques for medical imaging classification tasks," in *Proc. AMIA Annu. Symp.*, 2017, pp. 979–979.

[5] H.-C. Shin, N. A. Tenenholtz, J. K. Rogers, C. G. Schwarz, M. L. Senjem, J. L. Gunter, K. P. Andriole, and M. Michalski, "Medical image synthesis for data augmentation and anonymization using generative adversarial networks," in *Proc. 3rd Int. Workshop on Simulation and Synthesis in Medical Imaging (SASHIMI 2018), Held in Conjunction with MICCAI 2018*, Granada, Spain, Sept. 2018, pp. 1–11, Springer.

[6] A. Mikołajczyk and M. Grochowski, "Data augmentation for improving deep learning in image classification problem," in *Proc. 2018 Int. Interdisciplinary PhD Workshop (IIPhDW)*, 2018, pp. 117–122, IEEE.

[7] S. Kazeminia, C. Baur, A. Kuijper, B. van Ginneken, N. Navab, S. Albarqouni, and A. Mukhopadhyay, "GANs for medical image analysis," *Artificial Intelligence in Medicine*, vol. 109, p. 101938, 2020, Elsevier.

[8] R. Gupta, A. Sharma, and A. Kumar, "Super-resolution using GANs for medical imaging," *Procedia Computer Science*, vol. 173, pp. 28–35, 2020, Elsevier.

[9] H. Uzunova, J. Ehrhardt, F. Jacob, A. Frydrychowicz, and H. Handels, "Multi-scale GANs for memory-efficient generation of high resolution medical images," in *Proc. 22nd Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI 2019)*, Shenzhen, China, Oct. 13-17, 2019, Part VI, vol. 22, pp. 112–120, Springer.

[10] D.R. Sarvamangala and R.V. Kulkarni, "Convolutional neural networks in medical image understanding: a survey," *Evolutionary Intelligence*, vol. 15, no. 1, pp. 1–22, Springer, 2022.

[11] 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," *Neurocomputing*, vol. 321, pp. 321–331, 2018, Elsevier.

[12] Md. A. Talukdar, A. Siddika, A. H. Abir, M. Z. Hassan, and M. I. Hossain, "Medical X-ray image classification employing DCGAN and CNN transfer learning techniques," in *Proc. Seventh Int. Congress on Information and Communication Technology (ICICT 2022), London*, Vol. 1, pp. 839–848, 2022, Springer.

[13] M. Bali and T. Mahara, "Comparison of Affine and DCGAN-based Data Augmentation Techniques for Chest X-Ray Classification," *Procedia Comput. Sci.*, vol. 218, pp. 283–290, 2023, Elsevier.

[14] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Networks," *arXiv preprint arXiv:1406.2661 [stat.ML]*, 2014. [Online]. Available: https://arxiv.org/abs/1406.2661

[15] A. Radford, L. Metz, and S. Chintala, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks," *arXiv preprint arXiv:1511.06434 [cs.LG]*, 2016. [Online]. Available: https://arxiv.org/abs/1511.06434