Auxiliary Classifier Generative Antagonist Network for the Detection of Pneumonia

Robert Langenderfer

Department of Computer Science and Engineering, University of Toledo, Toledo, Ohio, USA
Email: robert.langenderfer@utoledo.edu (R.A.L.)
Manuscript received November 13, 2023; revised December 27, 2023; accepted February 18, 2024; published August 21, 2024

***Abstract*—Pneumonia is an inflammation of the lungs which is caused by bacteria, viruses, mold, and less commonly by environmental toxins. Pneumonia is extremely prevalent worldwide and is the number one cause of death among children under the age of five and is the most common reason for hospitalization for adults. Chest X-rays are a common medical tool for diagnosing this illness, but must be analyzed by trained radiologists, which is often time consuming and expensive. Therefore, it would be beneficial to have an accurate automated system for diagnosing pneumonia from radiological diagnostic imaging. A variety of machine learning techniques have been applied to the problem of medical image diagnostics and have exceeded the accuracy of the average radiologist. Medical datasets often suffer from a sparsity of training examples, so data augmentation is often necessary. Here we implement an auxiliary classifier generative adversarial network method which generates synthetic X-ray images which augment the training of a discriminator network. The described method has an accuracy of 97.7% when trained on the Pneumonia MNIST dataset, which is composed of low-resolution pediatric chest X-rays. Given the common difficulty of acquiring significantly sized medical datasets, the network was trained on a range of datasets sizes to determine the impact on performance given a smaller population of pneumonia examples. Even when trained on a subsample containing only 20 examples, the network achieves an impressive 84.54% accuracy. This system could be used in areas lacking proper medical personnel or act as a verification tool for diagnosticians.**

***Keywords*—auxiliary classifier generative adversarial network, Auxiliary Classifier Generative Adversarial Network (AC-GAN), neural network, pneumonia, X-ray, pediatric, machine learning, medical diagnostics**

I. Introduction

Pneumonia is a serious and potentially life-threatening respiratory infection that is caused by a variety of pathogens, including bacteria, viruses, mold, and is less commonly caused by environmental toxins. Pneumonia is characterized by an infection in the lungs, which leads to inflammation and the accumulation of fluid in the air spaces. It is a leading cause of death among children under the age of five worldwide [1], but it can impact people of all ages. 1.5 million people in the U.S. visit the emergency department each year due to this affliction [2] and in 2019, 2.5 million people died from pneumonia worldwide [3].

Diagnosing pneumonia can be challenging and often requires the analysis of chest X-rays by trained radiologists. However, this process can be time consuming, expensive, and subject to significant occurrence of misdiagnoses [4–6]. To address this challenge, researchers have developed a variety of machine learning techniques that are capable of diagnosing pneumonia from radiological diagnostic imaging with high accuracy. These techniques often outperform the accuracy of the average radiologist and have the potential to improve the efficiency and effectiveness of pneumonia diagnosis.

One approach that has been shown to be highly accurate is the use of an Auxiliary Classifier Generative Adversarial Network (AC-GAN). This method involves the training of two neural networks: a discriminator network that is trained to distinguish between real and synthetic X-ray images, and a generator network that is trained to generate synthetic X-ray images that are indistinguishable from real ones. These synthetic X-ray images are used to augment the training of the discriminator network, improving its ability to accurately distinguish between real and synthetic X-ray images [7].

In this paper, we describe an auxiliary classifier GAN (AC-GAN) method for generating synthetic X-ray images to augment the training of a discriminator network for the diagnosis of pneumonia. The AC-GAN method was trained on the Pneumonia MNIST dataset [8], which is composed of low-resolution pediatric chest X-rays. The described method achieved an accuracy of 97.7%, demonstrating its potential for use as an automated system for diagnosing pneumonia from radiological diagnostic imaging.

The use of machine learning techniques such as the described AC-GAN method has the potential to revolutionize how pneumonia is diagnosed and improve the efficiency and effectiveness of healthcare diagnostics. A potential application of this method is for machine learning-based decision support systems for radiologists. These systems could provide an additional layer of verification for radiologists when interpreting chest X-ray images, potentially leading to more accurate diagnoses [9]. This system could be particularly useful in areas where there is a shortage of trained medical personnel [10].

The remainder of the paper is organized as follows: Section II, Literature Review; Section III, Materials and Methods; Section IV, Results and Discussion; and Section V, Conclusion.

II. Literature Review

Neural network-based methods are growing in popularity and are increasingly finding a wide range of implementations over every imaginable field. Neural networks applied to medical imaging diagnostics have shown promising results in a wide variety of medical applications. Data from medical diagnostic tools such as Magnetic Resonance Imaging (MRI), Computer Tomography (CT), X-ray images, EKG, and ultrasound. X-ray, MRI, CT, and EKG have trained deep learning algorithms, and have been shown to effectively detect and diagnose a wide range of diseases [11]. Researchers have implemented neural network techniques to detect COVID-19, breast cancer, and Parkinson's [12]. Machine learning has been shown to exceed the diagnostic accuracy of trained medical personnel over many medical domains.

During the COVID-19 pandemic's global spread, Marques *et al*. (2020) leveraged the EfficientNet architecture within a convolutional neural network (CNN) framework to develop an automated diagnostic tool for detecting COVID-19 from chest X-ray images. By conducting binary and multi-class classification experiments, the research achieved remarkable accuracy rates of 99.62% and 96.70%, respectively. These results underscore the potential of advanced deep learning techniques in enhancing medical diagnostics, offering a significant tool for healthcare professionals in managing the pandemic [13].

Alamir *et al*. performed an in- depth meta-analysis of various GAN networks, including cGAN, IAGAN, and AC-GAN, in medical applications using data from MRI, CT, X-ray, and ultrasound images. Their paper covered a variety of diagnostic analyses related to diseases of the brain, heart, liver, lung, and kidney, highlighting their applications in medical imaging for tasks such as cross-modality translation, data augmentation, anomaly detection, classification, and image reconstruction [14].

Akpinar [15] offered a comprehensive comparison of different deep convolutional neural networks (DCNNs). Specifically, AlexNet, ResNet-50, and GoogLeNet were applied to chest X-ray datasets for pneumonia classification. This study not only compared the performance of these models but also examined various pre-processing techniques, identifying the most effective model with a top accuracy of 91.45%.

Ayan *et al.* [16] explored the efficacy of two convolutional neural network (CNN) models, Xception and Vgg16, for diagnosing pneumonia, which utilized transfer learning and fine-tuning methodologies during the training phase. Comparative analysis revealed that the Vgg16 model slightly outperformed Xception in overall accuracy with scores of 87% and 82%.

Röglin *et al*. [17] introduced innovative methods to generate synthetic data from a modest dataset comprised of 2D MRI scans of the spine. By evaluating the synthetic data with a classification network, the study demonstrates the beneficial impact of synthetic data augmentation on improving the original data's classification results. One method proved efficient in generating synthetic imagery from fewer than 50 images, presenting a viable strategy for augmenting datasets related to rare diseases.

Cala *et al*. [18] leveraged the strengths of ACGANs by generating images conditional on labels, to augment data for pneumonia detection from chest X-ray images. The integration of ACGAN-generated images with pretrained CNN models, especially ResNet-18 and variants, markedly improving model performance. This enhancement underscores the effectiveness of ACGAN-generated images in introducing variability to the training set.

III. Materials and Methods

A. Overview of AC-GAN Methodology

The Auxiliary Classifier Generative Adversarial Network (AC-GAN) is a machine learning technique that is used to classify data and generate synthetic data that is indistinguishable from real data. The network consists of two primary components: the discriminator and the generator networks. The discriminator network is trained to distinguish between real and synthetic data, as well as classes in the data. The generator network is trained to generate synthetic data that is indistinguishable from real data. The block diagram of the AC-GAN network architecture is depicted in Fig. 1.

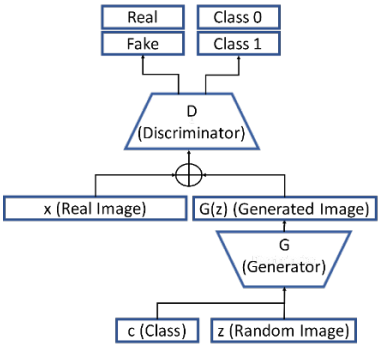

Fig. 1. AC-GAN network.

The AC-GAN technique is an extension of the standard Generative Adversarial Network (GAN) architecture. The GAN architecture is trained using a min-max function, which is a type of optimization that involves finding the balance between the two conflicting objectives. This function causes the generator network to minimize the performance of the discriminator network, while causing the discriminator network to maximize its performance.

$$\min_D \max_G \; V(D,G) = E_{x\sim P_{data}(x)}[\log D(x)] + E_{z\sim P_z(Z)}[log(1 - D(G(z))] \quad (1)$$

The "min-max" objective (Equation 1) is composed of two parts: the "min" component, which represents the optimization problem for the discriminator network (D), and the "max" component, which represents the optimization problem for the generator network (G). The min-max objective can be considered a two-player zero-sum game, where the discriminator attempts to maximize its performance while the generator is trying to minimize its performance. The goal of the equation is to find a balance between these two objectives, such that the generator generates realistic synthetic data that is indistinguishable from real data, and the discriminator accurately distinguishes between real and synthetic data. This formula expresses the adversarial nature of the technique [19].

The min-max objective is defined as the sum of two values: 1) The value of the log probability that the discriminator assigns to real data, $E_{x\sim P_{data}(x)}[\log D(x)]$. This term represents the performance of the discriminator on real data. 2) The value of the log probability that the discriminator assigns to synthetic data, $E_{z\sim P_z(Z)}\left[\log\left(1 - D\left(G(z)\right)\right)\right]$. This term represents the performance of the discriminator on synthetic data generated by the generator. By optimizing the min-max objective, the generator can continuously improve its ability to generate synthetic data that is indistinguishable from real data, and the discriminator can improve its ability to distinguish between real and synthetic data. This process repeats until the generator and discriminator reach a Nash equilibrium, when they generate and distinguish synthetic and real data with equal ability.

The difference between the two architectures is that the AC-GAN includes an additional classifier component that is trained to predict the class label of the input data. This

classifier network is known as the auxiliary classifier. The overall goal of the AC-GAN is to find a balance between these two objectives, such that the generator network generates synthetic data that is indistinguishable from real data, and the discriminator network accurately distinguishes between real and generated data, as well as distinguish between the data class labels [20, 21]. The advantage of this process is that the generator acts to augment the actual data with generated information, and essentially learns the features in the dataset from this generative perspective. In data domains where there is a sparsity of labeled examples, commonly the data is modified by mathematical transformations, or by adding noise, which serves to augment the supply of examples which helps the network algorithm generalize and improve accuracy [22]. The use of this technique may eliminate the need for traditional augmentation methods [23]. In this paper no additional augmentation was required to produce high levels of accuracy.

The objective function for the AC-GAN is composed of two parts:

$$L_s = E[\log P(S = \text{real} \mid X_{\text{real}})])] + E[\log P(S = \text{fake} \mid X_{\text{fake}})] \quad (2)$$

where L_s in Eq. (2) is the likelihood of predicting the source, and S is the source and X the input image, and P is the probability.

$$L_c = E[\log P(C = c \mid X_{\text{real}})] + E[\log P(C = c \mid X_{\text{fake}})] \quad (3)$$

where L_c in Eq. (2) is the likelihood of predicting the class, and c is the class label.

The latent space z in an ACGAN is a representation of data in a lower-dimensional space that is learned by the model. In this case, random latent space values were presented to the input of the generator network, which was used to generate synthetic image examples. This random latent space value can be thought of a "seed" for generating images. These synthetic images were sent to the input of the classifier network which attempts to both classify the image type and determines if it is real or generated by the generator. Note that the generator begins with a smaller dimension latent space value and expands on this information layer by layer in the network into a higher dimensional image. Conversely, the classifier takes the higher dimension image data and compresses it layer by layer in the network into the low dimensional classes and real/synthetic outputs [24].

B. Network Architecture

Discriminator: The purpose of the discriminator is to classify the input images so therefore the outputs of the discriminator consist of the real/not real output and the class label outputs. The real/not real output uses the sigmoid activation function and indicates the probability of whether the input image is real or generated by the generation network. The class label out output is a probability of the image belonging to each class via the SoftMax activation function, like any given multi-class classification neural network model, and is optimized using categorical cross entropy.

Generator: The output of the generator network is the image intended to fool the discriminator. The inputs are the latent dimension and real/not real. The image output is 28× 28 which matches the size of the training data.

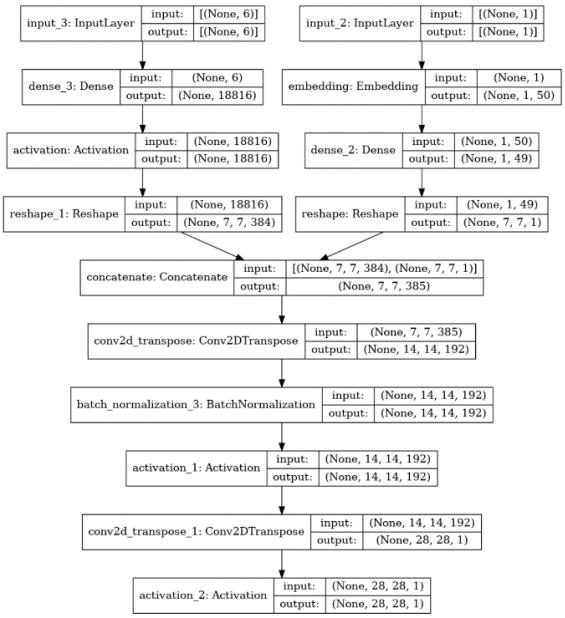

Fig. 2. Generator network.

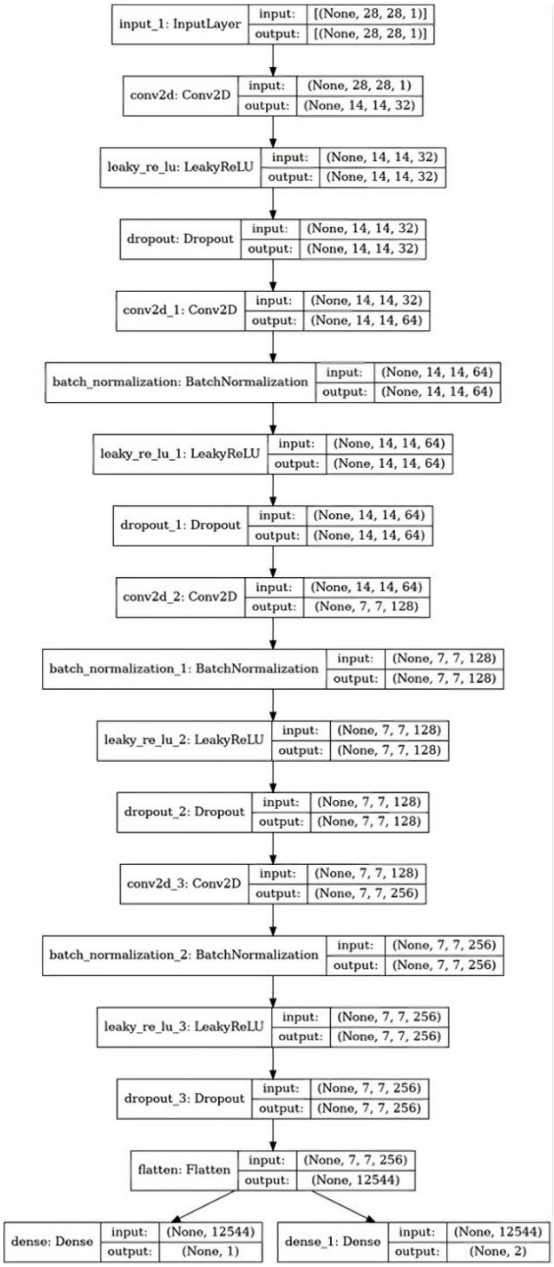

Fig. 3. Discriminator network.

The network parameters and layer types for the generator and discriminator are detailed in Figs. 2–3.

C. Dataset

The machine learning techniques in this paper have been trained and tested using the Pneumonia MNIST dataset. The Pneumonia MNIST dataset has been widely used for the development and evaluation of machine learning algorithms for automated pneumonia diagnosis and are composed of low-resolution pediatric chest X-ray images [25]. The dataset consists of 5,232 images divided into two classes: normal and pneumonia. The images were obtained from children aged 1 to 5 years and were annotated by trained radiologists. This dataset was originally composed of 5,856 anterior-posterior (front view) chest X-Ray images which have a range of sizes: (384×127)–(2,916×2,713). The final images in the dataset are one byte grayscale values from 0 to 255 and have been cropped and resized to a resolution of 28×28 pixels. The dataset consists of 5,332 training images and 524 test images. The training data contains 1,835 normal and 3,497 pneumonia examples, while the test data consists of 135 normal and 389 pneumonia samples. Fig. 4 displays two pneumonia and two normal images from the data set.

A significant challenge in the use of automated pneumonia diagnosis is the variability in the appearance of pneumonia on chest X-ray images. Pneumonia can manifest as a range of patterns on chest X-ray images, including nodules, infiltration, and consolidation, which can be difficult for machine learning algorithms to classify accurately. Also, the extremely low resolution of the Pneumonia MNIST images can make it difficult to discern the subtle features that may be predictive of pneumonia.

D. Implementation

The implementation of the Auxiliary Classifier Generative Adversarial Network (AC-GAN) was carried out using Python, a versatile programming language known for its efficacy in data analysis and machine learning tasks. The open-source libraries TensorFlow, Keras, and Scikit-learn each played valuable roles in the development process. TensorFlow, renowned for its powerful computational abilities, was employed for building and training neural network models, thanks to its robust handling of large datasets and its capability to perform complex numerical computations. Keras, an open-source neural network library, was utilized as an interface for TensorFlow, providing a more user-friendly platform to rapidly design the neural network architecture. Scikit-learn was implemented for pre-processing the data and for evaluating the model's performance using various metrics.

The training and testing were conducted on the Kaggle platform, which provided a conducive environment for machine learning experiments. The use of Kaggle's P100 GPU acceleration was a strategic choice, significantly reducing the computation time and enhancing the efficiency of the training process. This setup was particularly advantageous for handling the intensive computational demands of training the AC-GAN model.

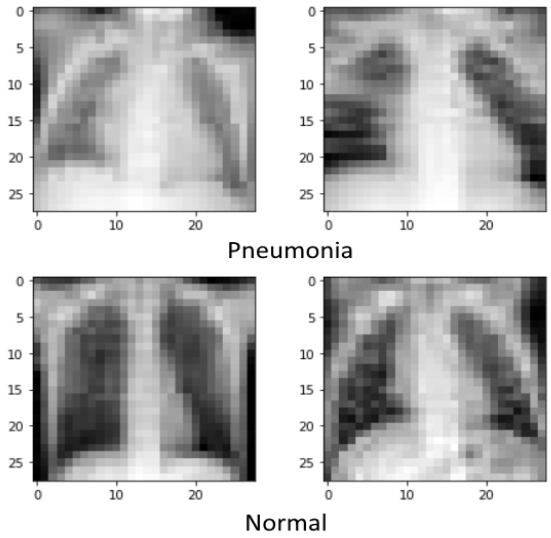

Fig. 4. Pneumonia MNIST dataset: pneumonia and normal images.

Table 1: Performance metrics for full dataset

Training Examples Normal/Pneumonia	Accuracy	Precision	Sensitivity	Specificity	F1	AUC
1835/3497	0.9779	0.9701	0.9846	0.9556	0.9701	0.9968

For data augmentation, we moved away from standard techniques and instead capitalized on the unique strengths of the AC-GAN's generator network. This method allowed us to synthetically expand our dataset, producing additional training images. This approach not only supplemented our dataset but also introduced a level of diversity in the training examples, which is often crucial for the robustness and generalizability of machine learning models. By generating new, synthetic images, the model was exposed to a broader range of data scenarios, aiding in its ability to learn and adapt more effectively.

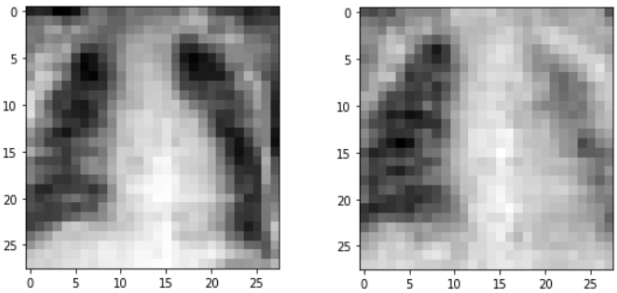

Fig. 5. Normal (left) and Pneumonia (right) synthesized images from the generator network.

IV. Results and Discussion

After training, the generator network synthesizes realistic X-ray images that closely mimic those from the dataset. Fig. 5 depicts two synthesized images from the trained generator network. These images are difficult to distinguish from the originals in Fig. 4.

The F1, precision, sensitivity, specificity, and accuracy metrics were employed to evaluate the performance of the network. These metrics are calculated using the following formulas:

$$precision = \frac{TP}{TP+FP} \tag{4}$$

$$sensitivity = \frac{TP}{TP+FN} \quad (5)$$

$$specificity = \frac{TN}{TN+FP} \quad (6)$$

$$F1\ Score = 2\frac{sensitivity*precision}{sensitivity+precision} \quad (7)$$

$$accuracy = \frac{TP+TN}{TP+FP+TN+FN} \quad (8)$$

Each reflects a different aspect of the diagnostic ability of the network. True Positives (TP) refer to the instances where the network correctly identified pneumonia cases, whereas True Negatives (TN) indicate the correctly identified normal cases. Conversely, False Negatives (FN) represent the pneumonia cases that were incorrectly classified as normal, and False Positives (FP) denote normal cases mistakenly identified as pneumonia. These metrics are crucial for understanding the model's diagnostic accuracy in different scenarios. These values are shown in Table 1.

Fig. 6 depicts a collation of these metrics in a confusion matrix. This matrix is a valuable tool for visualizing the model's performance across all categories, offering an at-a-glance understanding of its accuracy in diagnosing pneumonia using the full Pneumonia MNIST dataset. The confusion matrix not only quantifies the true and false diagnoses but also illustrates the balance between sensitivity and specificity, two key indicators of a diagnostic tool's effectiveness.

Comparative results including area under the curve and accuracy values are shown in Table 2. The ACGAN method is compared to several of the popular CNN models, such as ResNet, AutoKeras, and Google AutoML Vision. The proposed ACGAN method outperformed all other methods.

While the Pneumonia MNIST dataset contains a generous number of example images, large medical datasets are often difficult to acquire due to the rarity of certain conditions. To evaluate the adaptability and effectiveness of the ACGAN technique on smaller samples, the full dataset was divided into a range of smaller symmetrical datasets. Detailed in Table 3, the network was also trained on 2000, 200, 20, and 10 image datasets composed of half normal and half pneumonia images. There was only a very small difference of 0.27% in accuracy between the full 1835/2497 test dataset and the reduced, 1000/1000 version. Stepping down the training set by an order of magnitude to 100/100 reduced the accuracy to 95.04%, which is still an excellent result. The miniscule 10/10 training set saw a significant performance decrease to 84.54%, but this is still a good level of performance, especially considering the size of the dataset.

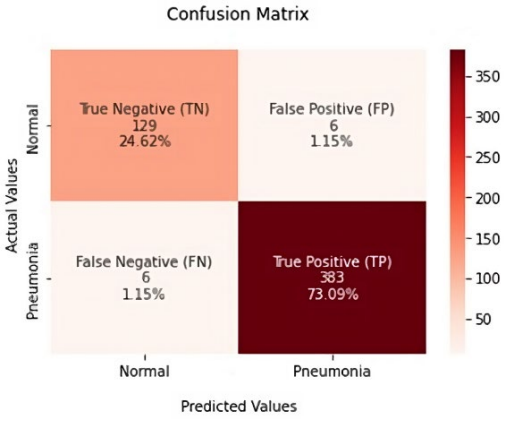

Fig. 6. Confusion matrix for ACGAN trained on complete dataset.

Table 2. Comparative performance metrics

Performance Comparison		
Network	**AUC**	**Accuracy**
ResNet-18	0.956	86.4%
ResNet-50	0.962	88.4%
Auto-sklearn	0.942	85.5%
AutoKeras	0.947	87.8%
Google AutoML Vision	0.991	94.6%
ACGAN (Proposed)	**0.997**	**98.7%**

Table 3. Performance metrics for network trained on reduced training sets

Training Examples Norm/Pneumonia	Accuracy	Precision	Sensitivity	Specificity	F1	AUC
1835/3497	0.9779	0.9701	0.9846	0.9556	0.9701	0.9968
1000/1000	0.9752	0.9645	0.9794	0.9629	0.9678	0.9929
100/100	0.9504	0.9318	0.9614	0.9185	0.9657	0.9828
10/10	0.8454	0.7958	0.8714	0.7704	0.8065	0.8765
5/5	0.7881	0.7250	0.8432	0.6296	0.7301	0.8466

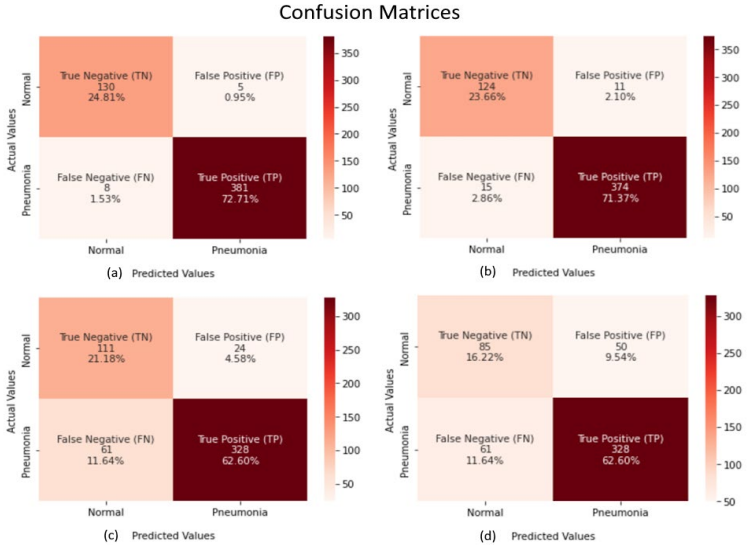

Fig. 7. Confusion matrices for (a) 1000/1000, (b) 100/100, (c) 10/10, and (d) 5/5 training sets.

Halving the dataset to 5/5 dropped the performance to 78.81%, which is still a commendable result, considering the extreme sparsity of examples.

Fig. 7 provides representations in the form of confusion matrices (a) through (d) for each of the four reduced datasets. These matrices illustrate the performance of our model across different training scales, offering a clear interpretation of the results. Each confusion matrix corresponds to a specific subset of the data – 1000/1000, 100/100, 10/10, and 5/5 – and showcases the model's ability to correctly identify cases of pneumonia versus normal cases in each scenario.

This series of experiments, conducted with datasets of varying sizes, not only underscores the resilience and adaptability of the ACGAN technique but also emphasizes its invaluable potential in contexts where medical data is scarce or hard to obtain. The ability of the ACGAN model to generate high-quality synthetic data plays a pivotal role in this context. By supplementing limited datasets with these synthetic images, the network effectively broadens its exposure to diverse data scenarios. This adversarial augmentation is crucial for enhancing the model's capability to generalize from smaller datasets, a key challenge in the field of medical diagnostics. The successful application of the ACGAN technique in such constrained data environments illustrates its promise in advancing diagnostic accuracy where data limitations are a significant obstacle.

V. Conclusion

Pneumonia is the leading cause of death of children 1 to 5 years of age worldwide, a sobering statistic that underscores the urgent need for improved diagnostic methods. Even when radiological imaging equipment is available, there exists a shortage of radiologists who can interpret diagnostic images, which leads to increased mortality from this treatable illness.

This paper presents the Auxiliary Classifier Generative Adversarial Network (AC-GAN), which demonstrates remarkable diagnostic accuracy in identifying pneumonia from chest X-rays. With an accuracy rate of 97.7%, our AC-GAN network not only surpasses other advanced neural network models but also outperforms expert radiologists, especially when analyzing low-resolution images. Such a high level of accuracy is pivotal in settings where high-resolution imaging is not available. Furthermore, the unique image augmentation capability of our model shows significant promise, maintaining a high level of diagnostic performance even with limited data. Impressively, it achieves an accuracy of 84.54% when trained on datasets as small as 20 patient images, demonstrating its potential in resource-constrained environments.

It is hoped that this work could lead to the development of machine learning-based decision support systems for radiologists, or even a fully automated system for the diagnosis of pneumonia. A system based on the techniques described herein could lead to more accurate and consistent diagnoses, which could potentially reduce worldwide mortality. Further research is needed to assess the effectiveness of this method in clinical settings, as well as its applicability to other medical domains.

Conflict of Interest

The author declares no conflict of interest.

Acknowledgements

I extend my sincere appreciation to Ezzatollah Salari and Jared Olouch for their support and for taking the time to review this work.

References

[1] Pneumonia in children. (2022). [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/pneumonia

[2] FastStats - Pneumonia. (Dec. 29, 2022). [Online]. Available: https://www.cdc.gov/nchs/fastats/pneumonia.htm

[3] Pneumonia - Our World in Data. (Dec. 30, 2022). [Online]. Available: https://ourworldindata.org/pneumonia

[4] H. F. H. System. (Apr. 22, 2022). Pneumonia often misdiagnosed on patient readmissions, studies find—ScienceDaily. [Online]. Available: https://www.sciencedaily.com/releases/2010/10/101022123749.htm

[5] G. J. Williams *et al.*, "Variability and accuracy in interpretation of consolidation on chest radiography for diagnosing pneumonia in children under 5 years of age," *Pediatr. Pulmonol.*, vol. 48, no. 12, pp. 1195–1200, Dec. 2013. doi: 10.1002/ppul.22806

[6] M. I. Neuman *et al.*, "Variability in the interpretation of chest radiographs for the diagnosis of pneumonia in children," *J. Hosp. Med.*, vol. 7, no. 4, pp. 294–298, Apr. 2012. doi: 10.1002/jhm.955

[7] J. Lubis, M. Nasir, Z. Zakaria, and M. Zhu, "To cite this article: Mingjing Zhu," *J. Phys*, p. 12165, 2021. doi: 10.1088/1742-6596/1827/1/012165

[8] J. Yang *et al.*, "MedMNIST v2: A large-scale lightweight benchmark for 2D and 3D biomedical image classification," Arxiv, Oct. 2021. doi: 10.48550/arxiv.2110.14795

[9] A. Kesselman *et al.*, "2015 RAD-AID conference on international radiology for developing countries: The evolving global radiology landscape," *J. Am. Coll. Radiol.*, vol. 13, no. 9, pp. 1139–1144, Sep. 2016. doi: 10.1016/J.JACR.2016.03.028

[10] S. G. Armato *et al.*, "The Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI): A completed reference database of lung nodules on CT scans," *Med. Phys.*, vol. 38, no. 2, pp. 915–931, 2011. doi: 10.1118/1.3528204

[11] M. Mirbabaie, *et al*, "Artificial intelligence in disease diagnostics: A critical review and classification on the current state of research guiding future direction," *Heal. Technol. 2021 114*, vol. 11, no. 4, pp. 693–731, May 2021. doi: 10.1007/S12553-021-00555-5

[12] G. Litjens, "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, Dec. 2017. doi: 10.1016/j.media.2017.07.005

[13] G. Marques, D. Agarwal, and I. Díez, "Automated medical diagnosis of COVID-19 through EfficientNet convolutional neural network," *Appl. Soft Comput.*, vol. 96, p. 106691, Nov. 2020. doi: 10.1016/J.ASOC.2020.106691

[14] M. Alamir, *et al.*, "The role of generative adversarial network in medical image analysis: An in-depth survey," *ACM Comput. Surv.*, vol. 55, p. 96, Dec. 2022. doi: 10.1145/3527849

[15] S. Karagol, K. N. Akpinar, and S. Genc, "Comparison of deep learning approaches in classification of the chest X-Ray," *J. Eng. Res.*, Nov. 2022. doi: 10.36909/JER.17121

[16] E. Ayan and H. M. Ünver, "Diagnosis of pneumonia from chest X-ray images using deep learning," *2019 Sci. Meet. Electr. Biomed. Eng. Comput. Sci. EBBT 2019*, Apr. 2019. doi: 10.1109/EBBT.2019.8741582

[17] J. Röglin, *et al.*, "Improving classification results on a small medical dataset using a GAN; An outlook for dealing with rare disease datasets," *Front. Comput. Sci.*, vol. 4, p. 102, Aug. 2022. doi: 10.3389/FCOMP.2022.858874/BIBTEX

[18] C. Florence Cala-Or, *et al.*, "Detection of pneumonia in chest X-Ray images using deep transfer learning and data augmentation with auxiliary classifier generative adversarial network," *Manila J. Sci.*, vol. 14, pp. 35–54, 2021.

[19] I. Goodfellow *et al.*, "Generative adversarial networks," *Commun. ACM*, vol. 63, no. 11, pp. 139–144, Oct. 2020. doi: 10.1145/3422622.

[20] ACGAN Architectural Design – coding. (Nov. 15, 2022). [Online]. Available: https://stephan-osterburg.gitbook.io/coding/coding/ml-dl/tensorfow/chapter-4-conditional-generative-adversarial-network/acgan-architectural-design

[21] A. Odena, C. Olah, and J. Shlens, "Conditional image synthesis with auxiliary classifier GANs," in *Proc. 34th Int. Conf. Mach. Learn. ICML 2017*, Oct. 2016, vol. 6, pp. 4043–4055.
[22] (May 04, 2022). A survey on image data augmentation for deep learning. *Journal of Big Data*. [Online]. Available: https://journalofbigdata.springeropen.com/articles/10.1186/s40537-019-0197-0
[23] S. Motamed, P. Rogalla, and F. Khalvati, "Data augmentation using Generative Adversarial Networks (GANs) for GAN-based detection of Pneumonia and COVID-19 in chest X-ray images," *Informatics Med. Unlocked*, vol. 27, p. 100779, Jan. 2021. doi: 10.1016/J.IMU.2021.100779
[24] M. Mirza and S. Osindero, "Conditional generative adversarial nets," arxiv, Nov. 2014. http://arxiv.org/abs/1411.1784
[25] MedMNIST. (Apr. 01, 2022). [Online]. Available: https://medmnist.com/

Copyright © 2024 by the authors. This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited (CC BY 4.0).