Improved Pneumothorax Classification using DCGAN-Based Data Augmentation

Subhan Chaudry
schaudry@utexas.edu
University of Texas at Austin
Austin, Texas, USA

Abstract

The detection of pneumothorax based on medical imaging presents unique challenges to clinicians, as they can be frequently misinterpreted and have variable presentations. Achieving high accuracy with AI models is difficult due to limited labeled datasets and restricted public access to clinical data. This study investigates how Deep Convolutional Generative Adversarial Networks (DCGANs) can be used to augment datasets with synthetic examples of chest x-rays, providing a simple solution to data scarcity concerns.

A DCGAN was trained to generated synthetic pneumothoras images, which were then utilized alongside the training dataset to balance class distributions for positive and negative samples. A ResNet-18 classifier was then trained on the augmented dataset and achieved training accuracy of 98.60% and ROC-AUC of 1.000. Grad-CAM and t-SNE visualizations were used to validate that the model focused on key clinical regions and learned features properly. The results in this paper demonstrate that DCGAN-based data augmentation can improve pneumothorax classification and offer a simple approach for scaling AI-driven medical imaging analysis, especially with limited data availability.

1 Introduction

Automated diagnosis of pneumothorax through machine learning has the potential to improve patient outcomes and result in a faster path to diagnosis. Pneumothorax is the presence of air in the pleural cavity, leading to lung collapse and severe respiratory distress, requiring immediate intervention to treat properly. Chest X-rays are the most common imaging modality for diagnosis by radiologists but challenging to read even for experienced professionals. Machine learning models have been developed to assist physicians but they usually rely on very large, labeled datasets.

One major difficulty with pneumothorax datasets is that they are typically not balanced properly with only a small number of pneumothorax examples to train on. It is not easy to collect additional labeled medical images since it requires expertise from radiologists and healthcare organizations have their own barriers to entry for data sharing. Traditional data augmentation techniques like rotations and scaling images are not complex enough to provide clinical variation to X-rays and build too strongly from training examples.

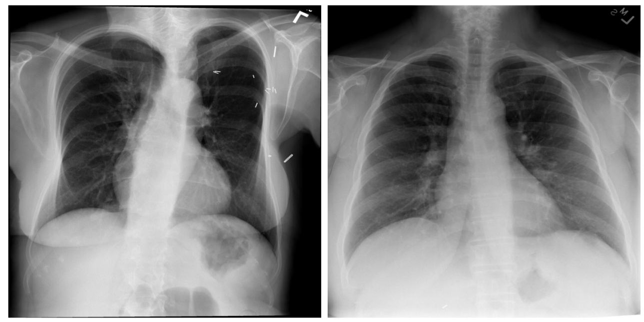

Figure 1. Pneumothorax Chest X-Rays

Deep Convolutional Generative Adversarial Networks (DCGANs) introduce a better alternative to these traditional augmentation techniques by creating synthetic data that more closely mimics the training data leveraged. With this, it is easier to generate realistic examples of pneumothorax chest X-rays to increase the diversity and size of training sets. DCGANs work by understanding the underlying patterns for the original data and producing examples that closely resemble key features in positive or negative samples.

In this project, DCGAN-based data augmentation is leveraged to improve the performance of classification models. The hypothesis is that synthetic images will assist us in identifying positive samples of pneumothorax and achieving better accuracy for medical machine learning models. This aims to illustrate how realistic data generation can be used to address diagnosis of rare conditions where sufficient labeled, training data may not be easily available.

2 Related Work

Previous studies on the efficacy of data enhancement with generative models have already shown significant improvements in model performance and classification models in healthcare.

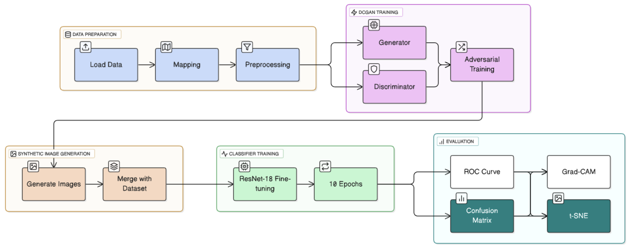

Figure 2. Pneumothorax Classification with DCGAN Augmentation Architecture Diagram

2.1 Liver Lesion Classification using GANs

Frid-Adar et al. [1] demonstrated that usage of GANS for synthetic liver lesion images significantly improved CNN performance with small datasets. Their approach focused on how increasing the diversity of training images helped the models generalize better and reduced overfitting during training pipelines. This set the stage for subsequent studies on how synthetic data could reproduce meaningful patterns found in medical imagining data.

2.2 Cardiac Image Segmentation with GANs

Sandfort et al. [2] leveraged CycleGANS to make CT scans sourced from different machines look more similar without matching sets of images. The results showed that synthetic data generation helps to segment features on unseen targets and understand patterns within the data more easily. The key takeaway was that GANs could be utilized to make medical images more consistent across different samples and to better unify data sets for machine learning development.

2.3 Chest X-Ray Augmentation for Cardiovascular Abnormalities

Madani et al. [3] also used GANs to generate synthetic chest x-rays to help with identifying heart abnormalities more easily, especially for rare conditions. They found that these images helped models to better train to recognize underrepresented conditions and improve their sensitivity to positive samples. Their work illustrated how GANs could balance class distributions and strengthen model performance for medical applications.

3 Methodology

This project includes a standard workflow for data preparation, DCGAN training and generation, model classification, and performance visualizations. The overall workflow for this project is mapped out in Figure 2.

3.1 Data Loading and Preparation

The pneumothorax dataset for this project was provided by the Society for Imaging Informatics in Medicine (SIIM) as part of their Pneumothorax Segmentation challenge. The dataset consisted of 11,582 chest x-rays stored in a DICOM format along with annotations in the form of run-length encoded (RLE) binary masks. As the file organization for the dataset was given in a nested form, a mapping from the unique identifier (UID) to the corresponding image path was created to provide access to the raw files.

The RLE masks were processed through decoding in order to create pixel-wise binary masks that showed regions presenting signs of pneumothorax. Images without pneumothorax were simply given a mask value of -1. For simplicity in processing images, all DICOM chest x-rays were converted to PNGs. Data preparation involved resizing images, rotations/flips for variation, and additional augmentation techniques. Some example images can be seen in Figure 1.

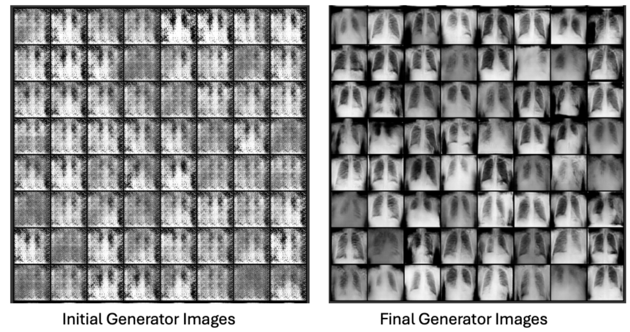

Figure 3. DCGAN Image Generation Iterations

3.2 DCGAN Architecture and Training

For generating synthetic images, a Deep Convolutional Generative Adversarial Network (DCGAN) was developed with a discriminator and generator. The generator model began by transforming random noise vectors selected through normal distribution. It then passes them through multiple layers of convolutions, normalization, and ReLU activations to upsample and transform appropriately. The end result is a structured 128x128 grayscale chest x-ray that closely resembles our training dataset. This is critical to ensure that synthetic images replicate patterns found in real pneumothorax images taken from clinical settings.

The discriminator model functionally acts as a binary classifier to determine whether an image produced by the generator is real or synthetic. It passes images through a series of convolutional layers with ReLU activations and downsampling. This is processed into a sigmoid output that provides the probability of how likely an image matches real pneumothorax chest x-rays.

Together, these two models act in an adversarial setting. The generator tries to avoid having the discriminator identify images as synthetic, while the discriminator attempts to continue to have high classification accuracy. This relationship allows for generated images to continue to improve over multiple training epochs. The DCGAN in this project was trained with binary cross-entropy loss and the Adam optimizer with tuned learning rates for training stability. Figure 3 demonstrates how the generator initially performed in its first iteration vs the final iteration and how it drastically improves over its training duration.

3.3 Synthetic Image Generation

Once the DCGAN completed training, the generator produced a series of synthetic pneumothorax images. These generated images closely matched realistic features of the training dataset like line visibility and lung collapse patterns. The synthetic images were added to the training dataset for greater diversity and addressing class imbalance issues.

3.4 Classification Model

For the final classification, a pretrained ResNet-18 Convolutional Neural Network (CNN) was used for pneumothorax detection. The final connected layer was mapped to a two-class output and fine-tuned on the training dataset consisting of original and synthetic images.

4 Results

The performance of this project was evaluated using a variety of quantitative metrics and visualizations. Metrics like training accuracy, ROC-AUC, and confusion matrix analysis assessed how well the classifier performed. Visualizations with Grad-CAM and t-SNE embeddings provided interpretability of model results and how it learned key features.

4.1 Quantitative Metrics

During the training pipeline, it was clear that over the 10 epochs that loss decreased along with stronger accuracy. By the final epoch, the classification model achieved an accuracy of 98.60% demonstrating strong convergence.

- Epoch 1/10: Loss = 0.0009, Accuracy = 77.59%
- Epoch 2/10: Loss = 0.0007, Accuracy = 84.86%
- ...
- Epoch 9/10: Loss = 0.0000, Accuracy = 98.20%
- Epoch 10/10: Loss = 0.0001, Accuracy = 98.60%

When evaluating our final classification model, it showed strong performance with an ROC-AUC score of 1.000 which meant that it had almost perfect separation between pneumothorax and normal cases within patient groups. This can be seen in Figure 4.

The confusion matrix in Figure 5 also emphasized the strong observed performance with proper classification of 3128 pneumothorax and 8290 normal cases. There was a small number of misclassifications, but true positive and true negative rates were very promising for this project.

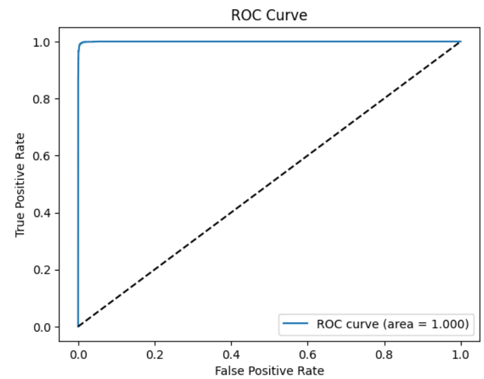

Figure 4. ROC Curve for model performance with AUC

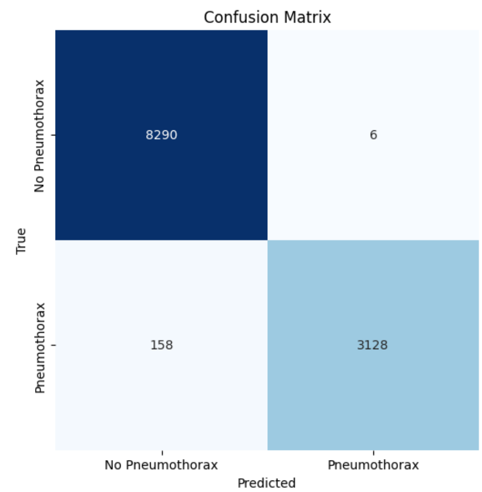

Figure 5. Confusion Matrix of predictions

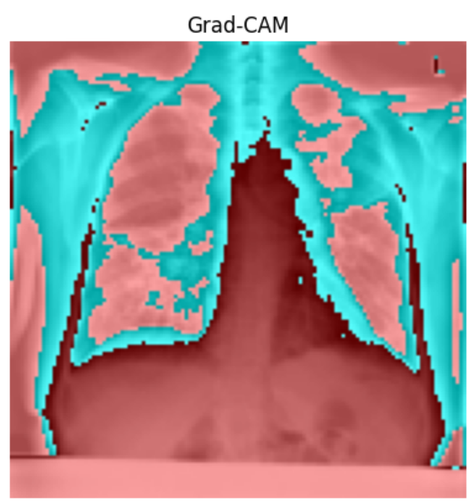

Figure 6. Grad-CAM heatmap for key areas of attention

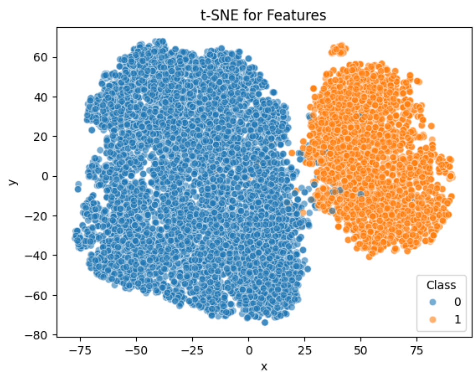

Figure 7. t-SNE projection with clear clustering

4.2 Visualizations of Performance

In order to properly visualize how the model interprets images, Gradient-Weighted Class Activation Mapping (Grad-CAM) was utilized to generate heatmaps of key focus areas. Grad-CAM takes the gradients for the classification score in the final convolutional feature map to identify which parts of the image most impact the final prediction. In Figure 6, one of the chest x-rays is referenced and the model correctly highlights relevant clinical regions like the lung apex and pleural space when determining a diagnosis.

For better understanding the distribution of feature representations, t-distributed Stochastic Neighbor Embedding (t-SNE) is leveraged to map the high-dimensional feature embeddings into two dimensions that are easier to analyze. Figure 7 shows that there is a clear distinction made between pneumothorax and normal cases within the feature space. This validates the results so far and that the models learned key features from the synthetic and real images provided.

Overall, the results from this project support that DCGAN-generated synthetic images can be valuable in improving the detection of pneumothorax from clinical images. The classification model generalized very well and achieved strong results across all quantitative metrics with minimal errors. Visualizations with Grad-CAM and t-SNE confirm that the model identified key features within chest x-rays and the data augmentation approach was effective in pneumothorax classification.

5 Conclusion

This project confirmed the initial hypothesis that synthetic images generated by Deep Convolutional Generative Adversarial Networks (DCGANs) could assist in identifying positive samples of pneumothorax and improve classification models for clinical machine learning. By augmenting our training dataset with realistic and diverse synthetic chest x-rays, the final classification model achieved near-perfect ROC-AUC score along with other performance metrics. The Grad-CAM heatmaps and t-SNE projections supported this even further. These results validate that data augmentation with generative technologies can be incredibly effective for enhancing classification performance, especially when labeled, extensive datasets are not readily available.

5.1 Implications of Results

The results of this project have significant implications in leveraging artificial intelligence and machine learning techniques for better diagnostic models. The success of generating synthetic data for training pipelines can guide future research and development around classification models that will be more sensitive to rare conditions like pneumothorax and reduce diagnosis error rates. This project emphasizes the growing role of GAN-based data augmentation to avoid requiring large, extensive, manually labeled datasets for developing effective models. This paves a path forward for organizations of all sizes scale their machine learning development for training iterations. One critically important finding in this project was that synthetic augmentation not only improved the performance of classification but also retained interpretability of clinical images.

5.2 Improvements to Project

Although this project demonstrated strong results, there could be additional steps taken to improve the quality and diversity of generated synthetic images. The DCGAN produced effective data augmentations, but there are now more complex architecture like StyleGAN2 and Diffusion Models that may show even greater performance with more anatomical details from medical imaging. Introducing other datasets from a variety of clinical settings could improve the broader generalization of the model and reduce biases that may exist based on the source of the dataset or subset of clinicians responsible for labeling.

5.3 Future Directions

Future development on the work in this project could build upon the data augmentation strategies for multi-label disease detection. In this scenario, models must correctly classify multiple abnormalities in chest x-rays like pneumothorax, effusions, consolidations, etc. A conditional GAN would enable the ability to target generating multiple examples of these abnormalities and balance class distributions appropriately. Models in healthcare require high levels of sensitivity to safely deploy them in practice, especially with the inclusion of synthetic data. Models should be calibrated much more carefully with more advanced techniques.

5.4 Final Reflection

Based on the results of this project, it is clear that synthetic data generation will have an impactful role in overcoming issues with data availability and current medical imaging analysis. With continuous advancements in data processing and complex model training approaches, building accurate

AI solutions in healthcare should continue to improve and bring institutions closer to tackling some of the most difficult clinical challenges in this space.

Cite This Work

Subhan Chaudry. 2025. *Improved Pneumothorax Classification using DCGAN-Based Data Augmentation.* Technical Report, University of Texas at Austin.

References

[1] Frid-Adar, M., Diamant, I., Klang, E., Amitai, M., Goldberger, J., and Greenspan, H. (2018). GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification. *Neurocomputing*, 321, 321-331.

[2] Sandfort, V., Yan, K., Pickhardt, P. J., and Summers, R. M. (2019). Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks. *Scientific Reports*, 9(1), 16884.

[3] Madani, A., Moradi, M., Karargyris, A., and Syeda-Mahmood, T. (2018). Chest X-ray generation and data augmentation for cardiovascular abnormality classification. *Medical Imaging 2018: Image Processing*, 10574, 105741M.