School of Information Technology and Engineering at the ADA University

School of Engineering and Applied Science at the George Washington University

A Hybrid MLP-CNN Generator for Medical Image Synthesis with WGAN-GP: Application to Breast Cancer Imaging

A Thesis
Presented to the Graduate Program of Computer Science and Data Analytics
of the School of Information Technology and Engineering
ADA University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in Computer Science and Data Analytics
ADA University

By
Rumiyya Alili

April, 2025

THESIS ACCEPTANCE

This Thesis by: **Rumiyya Alili**
Entitled: *A Hybrid MLP-CNN Generator for Medical Image Synthesis with WGAN-GP: Application to Breast Cancer Imaging*

has been approved **as meeting the requirement** for the Degree of Master of Science in Computer Science and Data Analytics of the School of Information Technology and Engineering, ADA University.

Approved:

Dr. Jamaladdin Hasanov

(Adviser) (Date) 28.08.2025

Dr. Abzatdin Adamov

(Program Director) (Date) 28.08.2025

Dr. Abzatdin Adamov

(Dean) (Date) 28.08.2025

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

Abbreviation	Explanation
GAN	Generative Adversarial Network
WGAN-GP	Wasserstein GAN with Gradient Penalty
cGAN	Conditional GAN
Leaky ReLU	Leaky Rectified Linear Unit
MSE	Mean Squared Error
SSIM	Structural Similarity Index Measure
FID	Frechet Inception Distance
MPV	Mean Pixel Value
MPVD	Mean Pixel Value Distance
GLCM	Gray-Level Co-occurrence Matrix

CONTENTS

I. ABSTRACT

Medical imaging serves as a powerful instrument for diagnosing and treating various diseases; however, the development of robust deep-learning models has been stifled by the problems of scarcity as well as imbalance in the availability of good-quality annotated datasets. Classic data augmentation techniques, such as transformation techniques and intensity manipulations, are also rather limited and do not reflect the complex structural variations inherent in medical images. In the proposed study, an attempt was made to fill this gap by exploring the use of image augmentation via Generative Adversarial Networks (GANs), particularly focusing on the Wasserstein GAN with Gradient Penalty (WGAN-GP).

Herein, the authors propose a framework based on WGAN-GP, which supports training stability while avoiding analog GAN disadvantages like mode collapse or gradient vanishing.

A model will even be able to access medical images as input data, as well as qualitative and quantitative evaluations about the likeness and clinical relevance of images generated from it. The validation process involves Fre´chet Inception Distance (FID), Structural Similarity Index (SSIM), Mean Squared Error (MSE), and texture analysis to compare synthetic with real ones. There is also an expert-based visual assessment of generated images to evaluate the applicability of image generation for end-use within the medical context.

Results indicate that GAN-based augmentation is effective in increasing data diversity and creates images almost indistinguishable from real medical scans yet retaining anatomical components. These implications emphasize the potentiality of GANs addressing data scarcity challenges for deep learning applications in health. It also addresses possible ethical and regulatory issues raised in relation to synthetic medical data, such as data privacy, fairness, and clinical validation in the broader context of regulatory and ethical challenges. This research contributes to medical image synthesis by realizing that GAN-based augmentation is likely to be an effective means for improving deep learning performance in medical diagnostics. The methodology opens the door for future studies interested in the integration of generative models into clinical workflows to improve diagnostic accuracy and model generalizability in real-world applications.

A Hybrid MLP-CNN Generator for Medical Image Synthesis with WGAN-GP: Application to Breast Cancer Imaging

II. Introduction

Medical image synthesis is vital in contemporary healthcare and is the most important process for training diagnostic model systems and enhancing data repositories [1]. Datasets in medical imaging are usually small, one-dimensional, and not suitable for developing and testing deep learning models in detecting, classifying, or segmenting diseases [2]. This diminishment is particularly observed in breast cancer imaging since it has few annotated datasets while requiring accurate and diverse image samples for developing well-functioning machine learning models.

The limited variety of medical image datasets severely affects the training of diagnostic models. In any case, the lack of annotated images makes it harder to teach medical professionals and oncologists – who actually use these images to build their practice. The scarcity of examples makes it difficult to train medical students effectively, as they may not have access to a wide range of cases. In [3], it pointed out that deep learning models have a major impact on the diagnostic procedure but are restricted concerning their size and diversity of training datasets. For breast cancer as an evocative example, the sheer difficulty of gathering clinical-grade annotated data presents itself as a big challenge toward the general accessibility of AI-assisted diagnostics.

Annotated datasets of medical images like mammograms are time-consuming and costly to obtain [4]. For instance, only around 2000 annotated mammogram images are contained in the Breast Cancer Digital Repository (DDSM), making it an insufficient amount for training deep learning models effectively.

Medical images consist of very fine patterns and exhibit very complex structures, such as tissue masses, calcifications, and tumors. Capturing such fine details remains a significant challenge for many contemporary image generation techniques. That is, clinically meaningful patterns should be present in images without any artifacts detecting the model or practitioner perceptually misleading. Furthermore, medical image outputs need to have very high resolution and clinically variable information in outputs that standard models are not capable of reproducing, and the models trained on datasets not related to medical applications hardly generalize the special features of medical images. [5]. In addition, stability in training is often compromised while using deep learning models for the medical generation of images especially when they are Generative Adversarial Networks (GANs). The instability, mode collapse, and vanishing gradients in training make synthesizing realistic medical images, preserving the anatomy as well as pathological details, more complicated [6].

Generative models such as the Generative Adversarial Networks (GANs), which were introduced in [7], have built some hope of addressing some of the challenges associated with medical image synthesis. GANs have been applied for image synthesis in medical imaging initially for augmentation and segmentation. In [7], the pioneers of GANs demonstrated the basic framework for the practice of adversarial training. Most of the early applications of the usage of GANs in medical imaging largely concerned generic tasks, with very little use on breast cancer applications.

A core motivation for using GANs in breast cancer imaging is to augment the existing datasets. In [8], the authors generated synthetic MRI images of liver lesions utilizing GANs and demonstrated that these images increased the lesion classification performance of deep learning models. While not expressively targeting breast cancer, this study established the proof of principle using GANs to augment medical imaging datasets and led to further work specifically focused on breast cancer imaging.

III. Problem Statement

Deep learning has made an indelible mark in the world of medical image analysis and has improved diagnostic accuracy to a great extent in radiology, pathology, and other related domains. Deep learning models derive their performance from large, diverse, and high-quality datasets. The real challenge with acquiring such datasets has always been in medical imaging. Limited availability of data, class imbalance, and privacy issues are some common challenges that come the way. However, in most cases, data augmentation methods fail to recover from these limitations. Conventional augmentation techniques, such as flipping, rotation, scaling, contrast adjustment, and adding Gaussian noise, are commonly applied in artificially inflating the size of datasets. Despite the enhancement in generalizability offered by conventional augmentation techniques, they suffer from many disadvantages [9]:

- Small variability: Intensity transformations do not impose new modifications on the original image that will replace the data's underlying distribution. Examples include rotating a pathological image or flipping an X-ray; no pathological information is imparted beyond what is present in the original image.
- These features remain constant: original anatomical structures or variations will not be competently reproduced, which may assist a learning algorithm in generalizing to novel unseen cases.
- Some augmentations and transformations, like cropping or changing intensity, could eliminate clinically significant information that hurt the model.
- Class imbalance persists: Even if the sample size is raised, it will not create a true new sample, after which the rare disease case stays underrepresented.

Because of these limitations, there is an increasing demand for more advanced augmentation that is specifically geared toward generating high-quality medical images that are diverse and realistic. Generative adversarial networks have emerged as an ultimate solution in the augmentation of medical imaging datasets by synthesizing new, realistic-looking samples that readily correlate with real patient data. In contrast to traditional methods of augmentation, GANs are capable of learning complex data distributions and synthesizing entirely new images that have retained some medically meaningful characteristics [10]:

- Realism and Diversity: GANs capture complex image patterns with pathological features so that they can produce synthetic samples comparably to real data.
- Class-specific augmentation: Training GANs to generate images of rare pathologies begins solving class imbalance issues.
- Better Preservation of Pathological Features: Pathological features can differ from simple transformations, allowing the generative model to create entirely new pathological variations and increase model robustness.
- Privacy-preserving sharing of data: GANs can generate synthetic images dissociated from any real patients, permitting the sharing of datasets without ethical or legal implications. This research aims to explore the augmentation of medical images using GANs, specifically evaluating the efficacy of Wasserstein GANs with Gradient Penalty (WGAN-GP) in generating high-quality synthetic mammograms. It will address the following questions:

1. How effective is GAN-based augmentation compared to traditional augmentation methods in terms of image quality and dataset diversity? 2. Will the performance of deep learning models used for medical image classification tasks improve using GAN-generated synthetic images? 3. What are the limitations and potential improvements for the application of GANs in medical image augmentation? By investigating these questions, this research contributes to the growing field of AI-based augmentation of medical images where the findings give insight into effectively using GANs to improve deep learning-powered diagnostic tools.

IV. Ethical considerations and data privacy

The rapid development of generative adversarial networks (GANs) in medical imaging leaves ethical concerns and data privacy issues wide open for debate. GANs generate synthetic images of the medical domain which can help alleviate drawbacks concerning data scarcity, privacy, and the need for data augmentation in medical image analysis, yet, such an advancement should be cautiously considered lest benefits come at the expense of patients' rights, safety, and well-being.

One pressing ethical issue of applying GANs to medical imaging is the privacy of patient data. Medical data, including images, diagnosis, and personal health information, is considered sensitive and is protected by law under statutes such as the Health Insurance Portability and Accountability Act (HIPAA) in the U.S. and General Data Protection Regulation (GDPR) in the EU. GANs generating synthetic medical images pose a lethal risk of inadvertently disseminating identifiable patient information through the model in question. This especially threatens the sanctity of privacy when GANs are trained on real medical data with poor or no anonymization, allowing for synthetic images that shall openly present private data.

The reverse is also true; the model could memorize unique characteristics of individuals in the training dataset such that sensitive information can be reverse-engineered from otherwise generated images. A pertinent defense against this will be to ensure that adequate anonymization and de-identification techniques are applied to training data and synthetic images. Additionally, for federated learning or differential privacy methods, individuals whose data is used would remain private and not compromise the privacy of others.

Another significant aspect of the ethical challenges of medical image generation is bias. Certainly, GANs that were trained on data from unrepresentative datasets may yield images that consider and even perpetuate the same biases carried in the original data. A medical image dataset, for example, may not capture the different demographics of certain racial, ethnic and gender groups; thus, a GAN model generates synthetic images that may fail to reflect these populations appropriately. Consequently, the medical diagnosis and treatment may differ because of biased AI training on inaccurate data, producing less accurate or even harmful results for a certain underrepresented group.