

A GAN-based Federated Learning Architecture for Data Augmentation of Medical Images

by

Abdullah Al Rakin

21341032

MD. Akib Iqbal Majumder

18201142

Mohammad Farhan Kabir

19101530

Rudmila Arafin

18301105

A thesis submitted to the Department of Computer Science and Engineering
in partial fulfillment of the requirements for the degree of
B.Sc. in Computer Science

Department of Computer Science and Engineering
Brac University
September 2022

© 2022. Brac University
All rights reserved.

Declaration

It is hereby declared that

1. The thesis submitted is my/our own original work while completing degree at Brac University.
2. The thesis does not contain material previously published or written by a third party, except where this is appropriately cited through full and accurate referencing.
3. The thesis does not contain material which has been accepted, or submitted, for any other degree or diploma at a university or other institution.
4. We have acknowledged all main sources of help.

Student's Full Name & Signature:



Abdullah Al Rakin
21341032



Akib Iqbal Majumder
18201142



Mohammad Farhan Kabir
19101530



Rudmila Arafin
18301105

Approval

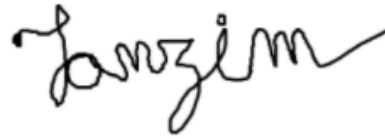
The thesis/project titled “Privacy Preserved Implementation of GAN Leveraging Federated Learning for Medical Images” submitted by

1. Abdullah Al Rakin (21341032)
2. MD. Akib Iqbal Majumder (18201142)
3. Mohammad Farhan Kabir (19101530)
4. Rudmila Arafin (18301105)

Of Summer, 2022 has been accepted as satisfactory in partial fulfillment of the requirement for the degree of B.Sc. in Computer Science on September 22, 2022.

Examining Committee:

Supervisor:
(Member)



Tanzim Reza

Lecturer
Department of Computer Science and Engineering
BRAC University

Co-Supervisor:
(Member)



Dr. Muhammad Iqbal Hossain

Associate Professor
Department of Computer Science and Engineering
BRAC University

Program Coordinator:
(Member)

Dr. Md. Golam Rabiul Alam

Professor
Department of Computer Science and Engineering
BRAC University

Head of Department:
(Chair)

Sadia Hamid Kazi

Associate Professor and Chairperson
Department of Computer Science and Engineering
BRAC University

Ethics Statement

We, hereby declare that this thesis paper is based solely on the results of our extensive research. In this research paper, we guarantee the utmost level of originality. We are assuring the complete confidentiality of all informants and sources.

Abstract

In the medical industry, the availability of precise data limits the scope of deep learning applications. Institutional norms restrict hospitals and research facilities owing to privacy concerns. Therefore, data collection from such sources is unfeasible. Federated Learning (FL) is promising in this scenario, but it does not guarantee data privacy. In this paper, we will use Deep Convolutional Generative Adversarial Network (DCGAN) and Wasserstein Generative Adversarial Network (WGAN) on an OCT dataset to demonstrate that the Federated GAN (FedGAN) architecture fails in these networks due to its innate structure. Additionally, introduce a Distributed Generative Adversarial Network (Distributed GAN) that collects and distributes the weights of each temporary GANs on the client side to the main server to tackle the mode collapse risk of non-iid data. This conserves the optimal distribution of data to all private discriminators while protecting sensitive individual data.

Keywords: GAN; Generator; Discriminator; Federated Learning; OCT; Deep Convolutional Generative Adversarial Network (DCGAN); Wasserstein GAN (WGAN); Distributed GAN; Mode Collapse, Non-iid Data.

Dedication

Our paper is first and foremost dedicated to our cherished parents. Without their unfathomable efforts and inspiration, we would not have been able to reach this point. The contribution of our supervisor and co-supervisor, sir, to the improvement of the research's quality is also noteworthy.

Acknowledgement

Firstly, all praise to the Great Allah for whom our thesis have been completed without any major interruption.

Secondly, to our supervisor Tanzim Reza sir and co-supervisor Dr. Muhammad Iqbal Hossain sir for their kind support and advice in our work. He helped us whenever we needed help.

And finally, to our parents and friends without their support, it may not be possible. With their kind support and prayer, we are now on the verge of our graduation.

Table of Contents

Declaration	i
Approval	ii
Ethics Statement	iv
Abstract	v
Dedication	vi
Acknowledgment	vii
Table of Contents	viii
List of Figures	x
Nomenclature	xi
1 Introduction	1
1.1 BackGround	1
1.2 Motivation	3
1.3 Problem Statement	3
1.4 Research Objectives	4
1.5 Contributions	4
1.6 Thesis Structure	4
2 Literature Review	6
2.1 Related Works	6
3 Model Background and Work Plan	10
3.1 Model Background	10
3.2 GAN	10
3.3 Federated Learning	11
3.4 DCGAN	12
3.5 WGAN	12
3.6 Distributed GAN	13
3.7 Workplan	13

4	Methodology	15
4.1	Methodology	15
4.2	Input Dataset	15
4.3	Data Preprocessing	16
4.4	DCGAN Implementation	16
4.5	Generator	17
4.6	Discriminator	18
4.7	Loss Function	19
4.8	WGAN-GP	19
4.9	Data Preprocessing	19
4.10	WGAN-GP Model Generator	19
4.11	WGAN-GP Model Discriminator	21
4.12	Architecture of the Proposed Model	22
4.13	Generator	22
4.14	Discriminator	22
5	Results and Analysis	25
5.1	Evaluating GAN	25
5.2	DCGAN	25
5.3	WGAN-GP	27
5.4	Distributed GAN	27
6	Conclusion	29
6.1	Challenges	29
6.2	Future Works	29
	Bibliography	32

List of Figures

3.1	Basic GAN Model	11
3.2	Workplan of the Research	14
4.1	Images of the Four Directories: CNV, DME, DRUSEN, and NORMAL	15
4.2	Samples of Reshaped 28x28 Dimension Data Images	16
4.3	Architecture of the DCGAN Generator	17
4.4	Architecture of the DCGAN Discriminator	18
4.5	Architecture of the WGAN-GP Generator	20
4.6	Architecture of the WGAN-GP Discriminator	21
4.7	Proposed Model Architecture	22
4.8	Architecture of the Distributed GAN Generator	23
4.9	Architecture of the Distributed GAN Discriminator	24
5.1	Output Images of the DCGAN	26
5.2	Comparison between Generator and Discriminator Loss	26
5.3	Output Images of the WGAN	27
5.4	Output Images of the Distributed GAN	28

Nomenclature

The next list describes several symbols & abbreviation that will be later used within the body of the document

AI Artificial Intelligence

CNN Convolutional Neural Networks

DCGAN Deep Convolutional Generative Adversarial Network

DCNNs Deep Convolutional Neural Networks

FE Federated environment

FedGAN Federated Generative Adversarial Network

FID Frechet Inception Distance

FL Federated Learning

GAN Generative Adversarial Network

GPU Graphics processing unit

LeakyReLU Leaky Rectified Linear Unit

OCT Optical Coherence Tomography

RAM Random Access Memory

SNNs Simulated Neural Networks

WGAN Wasserstein Generative Adversarial Network

WGAN – GP Wasserstein Generative Adversarial Networks Gradient Penalty

Chapter 1

Introduction

1.1 BackGround

Generative Adversarial Networks (GANs) are a breakthrough contribution to deep learning where generative models are trained in an adversarial manner to replicate data from a given dataset [33]. GANs require a large quantity of data with a wide variety of properties in order to generate a robust model. Unfortunately, privacy concerns and institution restrictions have led to a lack of varied data concealed by sources. Such scenarios are exemplified by hospitals and research institutions. For instance, the United States of America, the European Union, and a large number of other nations prohibit the export of patient data [31]. As a direct result, numerous medical facilities and research establishments are leery of cloud platforms and choose to manage their own servers on-premises. Yet within the same country, working on medical data is fraught with significant difficulties. Applications of deep learning have been quite successful in several technical domains [4].

GAN applications can play a significant part in addressing the shortage of medical data. The eventual goal of GAN applications in medical imaging prospects is to enhance the performance of models, such as classification and segmentation models [27]. It can be used to enrich data during centralized classification training after learning in a decentralized unsupervised manner [12], [21]. The inaccuracy of deep learning medical image models is diminished by small dataset sizes and poor image quality [38]. To avoid these challenges, some academics use GAN technology for data augmentation, such as super-resolution, image denoising, reconstruction, registration, and dataset augmentation.

Federated Learning (FL) is an example of distributed learning where a data-private method of collaborative learning in which many heterogenous contributors simultaneously train a machine learning model and then communicate their model changes to a central server so that they can be aggregated into a consensus model [37]. After that, the aggregation server will provide the consensus model to all of the participating colleges so that it can be utilized and/or extra training can be provided. A federated round is an iteration of this technique that consists of parallel training, update aggregation, and the distribution of new parameters. Each iteration of this procedure is referred to as a federated round. FL could be used to circumvent the data-sharing barrier for GAN training. Each server maintains (usually private) lo-

cal training datasets for the training of global models. This Federated Environment (FE) enables a Federated GAN (FedGAN) in which the generator is located on a global server and the client is hosted privately by discriminators. Each local discriminator is educated on distributed heterogeneous data on local servers before feeding back to the generator. FedGAN is stable with diverse data without jeopardizing confidentiality.

It is pretty apparent that in order to train an effective machine learning algorithm for medical image processing, considerable data volume is required [7]. Things are significantly more challenging with medical data. Patient data regulations protect patient health data for a good cause. Unfortunately, norms vary widely between countries, further complicating the matter. Recently, numerous large hospitals, organizations, and health authorities made anonymized data public to advance deep learning research [9]. Those datasets range in size from tens to thousands, and the annotations vary widely, as they are often done only once per dataset. Some countries are currently processing medical data in government-run databases accessible to researchers and industry. For example, Australia, Denmark, and Estonia are already on this track. Future developments in these nations can be promising [22].

Although predictive deep learning models have the potential to enhance medical diagnosis and treatment, in order for these models to be generally applicable, they require huge quantities of data from a variety of sources [16]. A recent study found that deep learning models tend to overfit to subtle institutional data biases and perform poorly when tested on data from institutions whose data were not observed during training. Deep learning medical imaging algorithms may choose to rely on the confounding elements connected with institutional biases in order to make their predictions as opposed to basing their conclusions on the evaluated apparent pathology. When compared to data that has been held out from the same institution, such models may provide correct results; nevertheless, they may not generalize well to data from other institutions or even between departments that are located within the same institution. The utilization of data from many universities to train a single model is an example of collaborative learning, which is a reasonable way to increase the volume and diversity of data. GANs are successively trained utilizing boosting algorithms in the work of Tolstikhin et al. to incrementally raise the performance of the final model [23].

An Optical Coherence Tomography (OCT) scan combined with a standard eye exam will provide the most detailed information about our overall eye health. It is a non-invasive imaging technique used in ophthalmology to perform quantitative analyses of eye tissues. It also generates real-time images by the application of the interferometry principle [2]. OCT makes it possible for clinicians to diagnose multiple eye illnesses, including glaucoma, Age-related Macular Degeneration (AMD) [29], and Diabetic Macular Edema (DME) [32]. Because more and more images are being obtained by OCT and saved in increasingly large electronic databases, medical professionals are devoting more time and energy to the process of interpreting those images. Although OCT was initially implemented in the field of ophthalmology, where it is currently utilized in clinical as well as research settings, it has since been shown to be effective in a wide range of other fields. To assess if OCT alone or in

conjunction with other imaging modalities has clinical use for identifying susceptible plaques, further research is essential to be carried out in this area [17].

1.2 Motivation

Data augmentation is required for deep learning training because centralized training requires a massive amount of data. However, medical image classification is not commonly preferred in the industry due to privacy concerns. Due to the fact that training data is occasionally non-IID, there are often many images devoid of the disease but few with it [35]. Therefore, it is essential to provide additional disease images to balance the data. Therefore, it may be beneficial to use GAN to generate synthetic data for augmentation. However, because the sharing of medical data creates security concerns, GAN training demands a large amount of data and poses a barrier in terms of training data availability. Here is where FL could alleviate the security issues while preserving the accessibility of GAN data. We chose to evaluate multiple versions of GAN in an FL framework, in addition to augmentation and data privacy.

1.3 Problem Statement

As the number of diseases increases at an alarming rate, novel categorization strategies based on various algorithms are being brought to us on a regular basis. It is challenging to accelerate the upgrading of algorithms in this province due to a lack of publicly available medical images and data. Furthermore, systematizing medical data repositories is challenging since it requires a high level of precision and accuracy in the field. In the realm of medical imaging, where aberrant results are by definition rare, this might be a concern. As a consequence, there is a scarcity of labeled data, which makes training an accurate algorithm difficult. Traditional data augmentation approaches (e.g., cropping, translation, and rotation) may help with some of these difficulties, but they still result in highly correlated image training data. In Goldbaum, M. et al. (2018), artificial intelligence (AI) withholds the ability to upgrade illness diagnosis and treatment by completing categorization tasks that are challenging for humans and rapidly evaluating enormous volumes of data [24]. Despite its capabilities, AI is difficult to evaluate in clinical settings and to prepare due to a lack of data.

One feasible solution to this challenge by utilizing a technique to generate synthetic pictures that offers a different type of data augmentation. When it comes to training deep learning models, data variety is crucial. As medical imaging data sets are often unbalanced, posing substantial hurdles for training deep learning models conducted by Rogers, J. K. et al. (2018) [25]. Further, which acts as an effective technique of data anonymization. However, implementing any machine learning model in the medical sector is demanding to train since it requires a large quantity of data. Medical picture data collection may be complex given the sensitive nature of health care data and the necessity to ensure privacy at all times. Some hospitals also do not want their data to be pushed onto the cloud; instead, they want it to be available in a safe manner inside their facility.

In summary, we find that the problem of medical data management is far from resolved. We found three viable strategies for resolving the issue: industry investment, state control, or non-governmental organizations. While each of them is conceivable, we must decide which one we like. In any case, the issue is critical and must be resolved to further deep learning research in medicine.

1.4 Research Objectives

The goal of this research is to ensure data augmentation of medical images while ensuring the privacy of the local data. We will be using OCT images as our dataset and implement different GANs (DCGAN, WGAN & Distributed GAN) while preserving the anonymity of the decentralized data in the federated environment. Additionally, by tweaking the parameters of our models of the different GANs we will try to find the optimum data privacy protection and high-quality generated images from GAN.

- Ensuring optimum data augmentation of the training data.
- Preserving the privacy of the client data.
- Experimenting with various GAN architectures on federated environment.
- Scrutinizing the drawbacks of FL
- Achieving target distribution of local data to avoid mode collapse

1.5 Contributions

Our key aim was is to assure data enhancement in a protected environment provided by federated learning paradigm. We will conduct experiments with Deep Convolutional Generative Adversarial Network (DCGAN), Wasserstein Generative Adversarial Network (WGAN), and Distributed GAN to determine the output variations. In addition, we will adjust the parameters of the models to improve the stability of the GANs and the accuracy of data augmentation. In conclusion, we will analyze the various GAN outputs to assess the quality of the generated data and also determine whether GAN architecture is compatible with a federated environment while protecting local data.

1.6 Thesis Structure

Chapter 1: Introduction where background, motivation, problem statement, research objectives, and contributions are discussed.

Chapter 2: Literature review where relevant works on our topic are elaborated.

Chapter 3: Model Background and Work Plan explicitly clarifies all the used models in depth. Also, the work plan shows the transition of our work

Chapter 4: Methodology carefully demonstrates the whole process of the study.

Chapter 5: Results and Analysis visualizes the desired output of our model.

Chapter 6: Conclusion restates the statement while jotting down the challenges and future works

Chapter 2

Literature Review

2.1 Related Works

This research aims to use different Generative Adversarial Networks (GANs) to augment data on Optical Coherence Tomography (OCT) image data while protecting the privacy of the local data by implementing a federated environment. In cases requiring medical image analysis, it's crucial to uncover and understand data patterns. Relevant or task-related features were typically produced by human specialists using domain expertise, making it difficult for non-experts to use machine learning techniques. Widespread use of deep learning in medical imaging improves image interpretation, representation, and categorization [19], [24]. AI-based models demand lots of data. A huge database of anatomy, pathology, and input data is needed to train an AI-based tumor detection. Due to the sensitivity of health information and the constraints on its use, obtaining this data may be difficult. Simulation and synthesis of medical images are gaining popularity in medical imaging [8].

In order to properly comprehend the potential drawbacks of conducting more studies, it is essential to emphasize previous studies conducted in contexts comparable to the present one. While some of the materials will directly address our theme, others will be relevant only in specific settings. We set out to evaluate a small selection of substantial research on the following topics.

OCT is a novel biomedical imaging method that permits noninvasive micron-scale, cross-sectional, and three-dimensional imaging of biological tissues [6]. Also, OCT is quite similar to ultrasound, however instead of using sound waves, it makes use of light waves. These light waves are able to reconstruct a profile of the eye by reflecting off of different depths within the eye, while a light beam that is laterally scanned is able to offer an image of the eye in three dimensions. Different layers of the retina become visible and the thickness of them is measured [3]. The images are helpful in the diagnosis, planning of treatment, and monitoring of retinal repair for many different eye conditions [6]. OCT image denoising uses low-pass, median, and mean filters [9]. OCT images can be digitally speckle-reduced in two ways (when $N > 1$). After angle compounding, the digital filter is added to the final image; N -filtered photos are combined to get the final image. OCT was the most frequently utilized imaging modality in ophthalmology in 2012, 2013, and 2014, with 5.35 million, 4.93 million, and 4.50 million OCTs conducted, respectively [18]. In cardiology,

OCT permits in-vivo visualization of coronary arteries, which can be used to diagnose stenosis, learn more about the processes of stenosis and thrombosis, and as an adjunct to percutaneous coronary operations [10]. In addition, neurology can use OCT pictures to evaluate degenerative diseases and other visual neuropathies. Human professionals such as radiologists and physicians have traditionally performed medical image interpretation in clinics [28]. Researchers and doctors have begun to benefit from computer-assisted therapy due to the vast variety of diseases and the likely depletion of human expertise. Medical imaging is extremely important in today's medicine. Modern imaging techniques such as X-rays, ultrasonography, CT scans, and MRIs may disclose fine details about the architecture inside our bodies. With the acceptance of deep learning in computer vision in 2012, the use of deep learning algorithms in medical imaging has grown significantly. X-rays, ultrasounds, magnetic resonance imaging (MRI), and optical coherence tomography are only a few examples of medical picture treatments (OCT) [28]. OCT images are typically very noisy, however, they can be rebuilt with the use of deep learning models such as GAN [26]

Due to artificial intelligence and increased computer performance, DL has made enormous progress in image processing. SNNs and CNNs recognize patterns and classify images [10], [19], [24]. Deep convolutional neural networks (DCNNs) are capable of recognizing, segmenting, and differentiating image objects and regions [15]. At the moment, it is utilized for purposes including classification, detection, segmentation, and the addition of data to medical images. Zhang et al. presented a Synergic Deep Learning (SDL) model [13]. This model includes many DCNNs that can learn from each other in order to address intra-class and inter-class variations that are the result of clinical variation and medical imaging [30]. In their study, Webb and colleagues suggested a classification system for WBC fluorescence imaging characteristics that was based on neural networks. The classification approach that was taken was successful. The discovery of out-of-distribution inputs causes an increase in the estimation of the degree of uncertainty that exists when a DC-GAN is included in a model. The parameters of the model are modified via transfer learning, which speeds up the convergence process and makes the model more accurate.

DCGAN was used to create reliable training data, a modified loss function was used to increase intra-class classification variability, transfer learning based on ImageNet was utilized to give finely tuned pre-trained network parameters, and used a modified version of ImageNet [30]. It was planned to incorporate T1-weighted MRI scans from the ADNI, AIBL, and NACC groups in the study published in [36]. The research consisted only of scans acquired from the 151 individuals included in the ADNI dataset. This is because the methods employed to get MRIs for each of the other pictures were unique. Teslas are the units of measurement for the magnet strength of an MRI scanner. There are numerous possible strengths ranging from 0.5T to 3T. (Tesla). Higher values need greater financial investment but provide superior photographs. During the scanning procedure, a 1.5T and a 3T scanner were utilized on the same individual. Simultaneously, a GAN and a fully convolutional model were trained. In order to train the GAN, MRIs of individuals with moderate cognitive impairment were employed. The FCN utilized various scans since it could only differentiate between normal cognition and Alzheimer's disease.

The GAN model was subdivided for training, testing, and validation purposes. This work demonstrates how image-to-image translation may be applied to 1.5T MRIs to enhance the visibility of Alzheimer’s disease. This would make it simpler to estimate a person’s emotions.

WGAN outperforms DCGAN in terms of training, which is significant given DCGAN’s concerns with mode collapse and the vanishing gradient problem. This indicates that the underlying training issue with GAN has been resolved and is no longer a concern. WGAN reduces the amount of time required to compute an accurate and effective EM distance. After each iteration of the gradient update of the evaluation function, the weights must be bound to a certain, limited range. The fact that WGAN training advances at a slower rate than DCGAN training is a prevalent issue. WGANs are an alternate training method for GANs that attempts to address the fundamental issues posed by the conventional method. Regarding the training of WGANs in particular, it is not necessary to strike a balance between training discriminators and training generators. In addition, training WGANs does not necessitate the development of specific network architecture. Additionally, mode dropping in GANs is not nearly as prevalent as it formerly was. One of the characteristics that make WGANs one of the most effective types of neural networks is their capacity to repeatedly estimate the EM distance by refining the training of the discriminator. Visualizing these learning curves not only aids in debugging and selecting the optimal hyperparameter values, but also has a substantial bearing on the quality of the samples.

The most prominent problem of using DCGAN and WGAN in a Federated Learning paradigm is the mode collapse problem, vanishing gradient problem, and the risks of data leakage via the weight transformation in the local client data. Durugkar et al. provide a centralized multidiscriminators technique to improve discriminator judgment on produced data to evade mode collapse. In a similar vein, Hoang et al. propose a centralized multi-generator solution that aims to increase generator capacity while minimizing the issue of mode collapse. Wang et al. [39] built an aggregated model by deploying a group of GANs that were separately trained and then assembled in a cascade.

Recent research has sought to improve GAN convergence by taking into account a large number of generators and discriminators. Despite this, the focus of these efforts is not on working with distributed datasets. This is where distributed GANs on non-iid data flourish. We feel that a large number of discriminators is preferable to a single generator. Recent research has demonstrated that certain fundamental techniques based on a single generator and several discriminators, or a mixture of generators and a single discriminator can outperform a single GAN. Hosting multiple private discriminators on the local entities and aggregating the feedback to the global generator helps in terms of protecting the sensitive data. When designing an architecture for machine learning, it is vital to include techniques for protecting privacy. This is due to the increasing number of privacy issues surrounding the sharing of data, as well as the rules implemented to preserve privacy, such as HIPAA [37].

In 2016, around 79 million CT scans and 38 million MRI scans were performed in the United States [31]. Despite this, the datasets available for machine learning research are still only modestly enlarged. The most extensive public collection of medical imaging data is consisting of 32,000 CT scans; nevertheless, this only accounts for 0.02% of all images acquired in the United States each year. The ImageNet [39] project, on the other hand, is a massive visual dataset created for study on visual object recognition. It has over 14 million photos organized into over 20,000 distinct categories. Protecting sensitive data requires a combination of privacy controls, synthetic data transfer, and design adaptability. Only information that presents a misleading image is transmitted. Privacy is protected despite the fact that the central generator does not have access to the raw data. The primary generator includes a few details about raw hospital photographs. The generator transmits the phony image to hospital discriminators. This strategy keeps raw data private from the central generator. Because synthetic data are what they are, the creator of synthetic images can freely distribute them. With this approach to gathering and disseminating data, a trustworthy, publicly accessible medical database can be created. Researchers, physicians, and the expansion of medical intelligence can benefit from the database's infinity. Using numerous GANs and their architecture habitat in federated learning, we experiment with the data privacy issue and highlight the limitations of the federated learning paradigm.

Chapter 3

Model Background and Work Plan

3.1 Model Background

Machine learning is used in the vast majority of AI technology [7]. Supervised machine learning is the most common method. Algorithms of supervised learning are fed examples of inputs and outputs. They figure out how to create a mapping by associating each input with its corresponding output. Images, natural language sentences, and audio waveforms are examples of input, but output examples are basic. The most common type of supervised learning, classification, yields a simple number code (a photo might be recognized as coming from category 0 containing apples, or category 1 containing oranges, etc.) [1]. SL is typically more accurate than individuals after training, which is why it is utilized in a wide range of products and services. The goal of this article is to discuss generative adversarial networks, an approach for unsupervised learning that employs generative modeling [12], [33].

3.2 GAN

GAN is a promising deep neural network that has been used to a number of machine learning problems, such as semantic segmentation, text-to-image transformation [1], image classification, generation and editing, and super-resolution image reconstruction [11], [20]. In medical imaging, GAN is utilized in two ways. The first is the generative element, which aids in the research and the structure of training data and learning the process to generate new synthetic images. GANs are particularly promising for addressing data scarcity and patient privacy. The second examines the selective feature of the discriminator, which can be viewed as a learned prior for normal images [34]. This uses as a detector when irregular images are dispensed. The equation described for GAN in the paper of Goodfellow et al is [33]:

$$\min_G \max_D V(D, G) = \min_G \max_D (E_{x \sim P_{data}(x)}[\log D(x)] + E_{z \sim p_z(z)}[\log(1 - D(G(z)))])$$

Where, D = Discriminator

G = Generator

$P_z(z)$ = Input noise distribution

$P_{data}(x)$ = Original data distribution

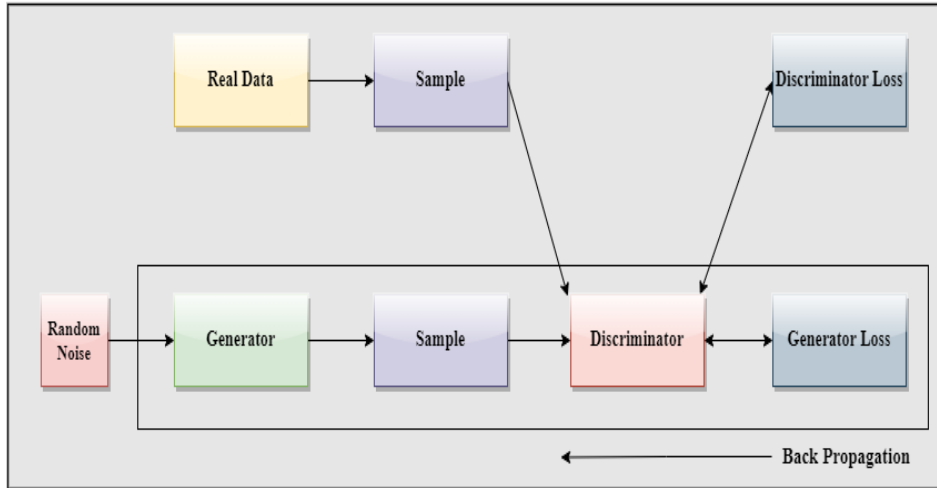


Figure 3.1: Basic GAN Model

Adversity implies that the model was trained in a real-world setting using neural networks as AI techniques. The generative model assesses the data distribution after training to improve the discriminator’s error rate. It is based on a model that estimates the real-data origin of the sample. Training data was used to generate the real sample. The generator network generates noise at random. The discriminator network then determines whether or not the sample is genuine. We train the discriminator on real data for n ”epochs” before freezing the generator and discarding its training set. The network will only send forward data and will not accept backward data. The discriminator is trained on real data to forecast them and on fictitious data to identify them as fictitious. The generator is then trained to misrepresent samples. The discriminator is then frozen, allowing us to build upon and mislead it. The discriminator is then trained using fictitious data to predict generator network samples. The discriminator correctly predicts values. We educate the generator to fool it after collecting discriminator predictions. Adversarial networks function effectively if the generator and discriminator are stable. The generator will not operate if the discriminator is too powerful, because fraudulent samples will always be discovered as such. Creating network images is worthless if the discriminator network is extremely permissive. GANs are incapable of determining how frequently an object should appear in a particular area. We must ensure that the model investigates the GAN challenge. We’ll also exclude 3D photographs from our medical dataset because GAN can’t recognize them and renders them flat. In addition, the general structure requires attention.

3.3 Federated Learning

In 2017, federated averaging was introduced as FL for the first time. Many cases are required for AI algorithms to produce models as good as medical specialists while protecting personal data through federated learning. By eliminating data pooling, federated learning decentralizes deep learning. The model is instead trained in multiple locations. Assume three hospitals worked together to develop a model to automate diabetic retinopathy image analysis. Each hospital that employs a client-server federated model would be given a copy of the global deep neural net-

work to train on its own dataset. Participants submitted their modified version to a centralized server after training the model locally for a few rounds and kept their dataset private. The central server would compile participant feedback. Participating universities would receive modified settings for local instruction. Because the model is not dependent on any specific data, if one of the hospitals were to leave the training team, the model’s training would continue. Meanwhile, a new hospital could join the program at any time. This is just one of the numerous federated learning strategies. Every participant gains global knowledge through the analysis of local data, which is a common thread that runs through all methodologies. As it is obvious that medical data must be secure and private, we will train the data in a federated environment to protect the data silo. There is a significant paucity of medical photographs for research purposes. OCT pictures are a good example of this. As a result, we want to use distributed GAN to protect data privacy while generating correct OCT pictures, ergo data augmentation.

3.4 DCGAN

DCGAN is a new GAN design extension that combines a CNN in addition to the symmetric generator and discrimination [10]. Alec Radford introduced it in 2015 [5] to compensate between CNNs for supervised and unsupervised learning. Its convolutional structure allows it to balance GAN training. DCGAN omits the pooling layer seen in CNNs, allowing the framework to spatially sample up and down on its own. To boost the training’s stability, the complete network uses a fractional-strided convolutional layer in lieu of an up-sampling layer and pooling layer. The DCGAN significantly enhanced the consistency of GAN training and the caliber of the results provided [30]. DCGAN originated as an innovative technique for creating graphics, audio, and films. Due to the layering, CNN feature extraction is also utilized [14]. DCGAN has demonstrated impressive performance on large-scale datasets including CelebA, LSUN, and Google Image Net [36].

3.5 WGAN

An adversarial network is a network that attempts to reduce an approximation of the Earth-distance Mover’s (EM) divergence [12]. One sort of adversarial network is the Wasserstein GAN, also known as WGAN. This variation of the GAN formulation, in contrast to the original GAN formulation, which focused on minimizing the Jensen-Shannon divergence, concentrates on increasing the Jensen-Shannon convergence. It generates training that is more stable than the original GANs, with less evidence of mode collapse, and it also produces meaningful curves that may be exploited for troubleshooting and detecting hyperparameters. WGAN is a reflection of the absence of convergence since GANs are unstable. For the purpose of correctly differentiating the data, the discriminator in WGAN, who is also referred to as the critic, strives to maximize the distance that exists between the genuine data and the data that was created. The goal of the generator network is to produce data that is as realistic as is practically possible by narrowing the gap between the created data and the actual data as much as is possible.

3.6 Distributed GAN

Distributed GAN functions with multiple local temporary discriminators and utilizes the feedback by aggregating to a central generator. In this manner the nature of the individual sensitive data from the local servers are intact. Such readiness of this architecture is a lucrative insight for researchers with private data holders. Distributed GAN inherently does the task of the federated environment by protecting data privacy, whereas traditional federated learning has the risk of data leakage by reverse engineering through the weights of the local data and sharing gradient information. Local discriminators in a distributed GAN architecture operate as a shield protecting sensitive data from the querying entity, ensuring privacy. Furthermore, the nature of artificial data enables the generator to freely spread it, which is crucial for privacy-sensitive applications.

3.7 Workplan

The primary goal of this study is to improve data for medical images. Consequently, we must utilize GAN (Generative adversarial learning). GAN can be used to produce data, but it requires a substantial quantity of training data. The accessibility of training data is an obstacle for GAN because the sharing of medical data generates security concerns. Thus, we offer a novel decentralized strategy for producing more images for our GAN model. The source of our dataset is the 2018 "Large Dataset of Labeled Optical Coherence Tomography (OCT)". The data was split into training and testing sets. Due to the varied dimensions of the dataset photos, the image sizes will also vary. For data preprocessing, we were required to resize the image to 28×28 pixels and convert it to a NumPy array. Then, we selected the DCGAN model for our initial GAN model and trained it with the data, in addition to training WGAN-GP with this dataset. However, neither of these GANs could be applied to the Federated learning environment due to the inherent architecture of FL. As a result, we designed a GAN that can be distributed across clients and sends weighted values to the server in order to train the generator.

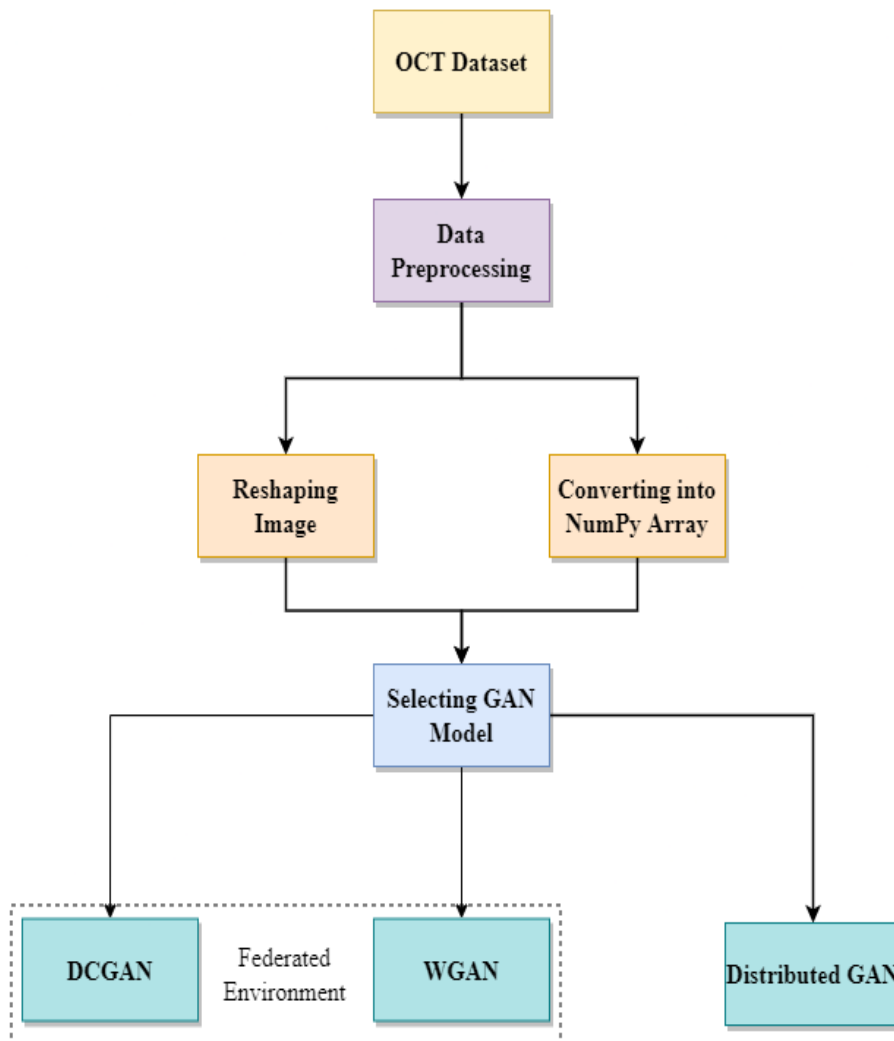


Figure 3.2: Workplan of the Research

Chapter 4

Methodology

4.1 Methodology

In the beginning of our research, we were assigned with obtaining a dataset to feed to GANs in order to generate realistic OCT images. A dataset was initially obtained using an open-source portal. There are four types of OCT images in this collection: CNV, DME, DRUSEN, and NORMAL. The training dataset was run via GAN models (DCGAN, WGAN, and Distributed GAN) to generate OCT images for the data augmentation challenge.

Hence, this section discusses the design and implementation of different types of GANs. To ensure that this technique remains flexible and adaptable, the complete task is divided into smaller subcategories. Several of the most essential processes include the following:

4.2 Input Dataset

The dataset we used focuses on common treatable blinding retinal diseases, mainly on OCT images. Therefore, we use CNN in this paper for image classification. It is from the 2018 "Large Dataset of Labeled Optical Coherence Tomography (OCT)" were used in the process of compiling the dataset for the research.

There are JPG photos of varying resolutions (512×496 , 1536×496 , and 768×496) contained within it. There is only one color channel and 83,484 training images. The dataset included two photo files that were respectively labeled "train" and "test." Each training and assessment dataset consists of a total of eight subcategories, which are denoted as CNV, DME, DRUSEN, and Normal respectively (Shown in figure 4.2).




Figure 4.1: Images of the Four Directories: CNV, DME, DRUSEN, and NORMAL

4.3 Data Preprocessing

The training data was utilized to enhance the raw data. Scaling the photos was necessary because the training data came in a wide variety of resolutions (28,28,1).




Figure 4.2: Samples of Reshaped 28x28 Dimension Data Images

The data was then converted to a NumPy array format so that it could be used in subsequent GAN training. Scale the values from -1 to 1 using DCGAN, WGAN-GP, and Distributed GAN, which were the methods utilized in the process of adding training data to the dataset. We normalized the dataset because, in general, normalization lowers the amount of time required for training (shown in figure 4.2).

4.4 DCGAN Implementation

To begin, we used a Deep Convolutional Generative Adversarial Network, or DCGAN, to simulate the training data for the OCT dataset. Vanilla GAN is another name for DCGAN, which is more often used. DCGAN is widely utilized for a wide variety of image generation tasks. Two neural networks, a discriminator and a generator, composed the DCGAN. The generator attempts to generate realistic visuals that can trick the discriminator in determining which one is real and which one is fake.

Convolutional layers were used to construct the discriminator and transposed convolutional layers were used to construct the generator, which was capable of creating fake images from the dataset. Our discriminator did not apply a maxpool layer for downsampling. For downsampling, stride was utilized. This is the DCGAN architectural standard.

4.5 Generator

The purpose of our generator is to produce $28 \times 28 \times 1$ images that can mislead a discriminator into believing they are real. The generator's input is a 100-dimensional vector of noise. This noise is sent to the sequential model we've established for our generator. We are adding a dense layer with an input size of $7 \times 7 \times 256$ to the sequential model. Then, we apply BatchNormalization to the layer and use LeakyReLU as the activation function. Therefore, $7 \times 7 \times 256$ is a low-resolution version of the output image, and we pass the random noise. Then, in the subsequent layer, we use Conv2Dtranspose to upsample our data, as the size of the first layer was $7 \times 7 \times 256$. For upsampling our data, Conv2Dtranspose is used. Then, we add two more Conv2Dtranspose operations to upsample from $7 \times 7 \times 256$ to $14 \times 14 \times 64$ to $28 \times 28 \times 1$ samples. This function contains a neural network that will perform the generator's functionality (shown in figure 4.4).

Model: "sequential_1"

Layer (type)	Output Shape	Param #
dense_1 (Dense)	(None, 12544)	1254400
batch_normalization_3 (Batch Normalization)	(None, 12544)	50176
leaky_re_lu_3 (LeakyReLU)	(None, 12544)	0
reshape_1 (Reshape)	(None, 7, 7, 256)	0
conv2d_transpose_3 (Conv2DTranspose)	(None, 7, 7, 128)	819200
batch_normalization_4 (Batch Normalization)	(None, 7, 7, 128)	512
leaky_re_lu_4 (LeakyReLU)	(None, 7, 7, 128)	0
conv2d_transpose_4 (Conv2DTranspose)	(None, 14, 14, 64)	204800
batch_normalization_5 (Batch Normalization)	(None, 14, 14, 64)	256
leaky_re_lu_5 (LeakyReLU)	(None, 14, 14, 64)	0
conv2d_transpose_5 (Conv2DTranspose)	(None, 28, 28, 1)	1600
=====		
Total params: 2,330,944		
Trainable params: 2,305,472		
Non-trainable params: 25,472		

Figure 4.3: Architecture of the DCGAN Generator

4.6 Discriminator

We've added a supplementary function to the discriminator. Typically, the discriminator is a classifier that determines if a picture is real or fraudulent. As our actual dataset is $28 \times 28 \times 1$, the discriminator's input will also be $28 \times 28 \times 1$ and the image from the generator. The discriminator's output is a binary value that indicates whether an image is real or fraudulent. We employ two Conv2D layers with a stride of two. The maxpool layer was not used for this discriminator. The 2,2 stride is for downsampling. And for the activation function LeakyReLU, we are flattening the layer, followed by a dense layer with sigmoid activation, which will reveal if the image is real or fraudulent (shown in figure 4.5).

Model: "sequential_3"

Layer (type)	Output Shape	Param #
conv2d (Conv2D)	(None, 14, 14, 64)	1664
leaky_re_lu_9 (LeakyReLU)	(None, 14, 14, 64)	0
dropout (Dropout)	(None, 14, 14, 64)	0
conv2d_1 (Conv2D)	(None, 7, 7, 128)	204928
leaky_re_lu_10 (LeakyReLU)	(None, 7, 7, 128)	0
dropout_1 (Dropout)	(None, 7, 7, 128)	0
flatten (Flatten)	(None, 6272)	0
dense_3 (Dense)	(None, 1)	6273

=====
Total params: 212,865
Trainable params: 212,865
Non-trainable params: 0
=====

Figure 4.4: Architecture of the DCGAN Discriminator

4.7 Loss Function

We require two loss functions. One loss function is responsible for the computation of the generator, while the other is responsible for the calculation of the discriminator. Therefore, for the loss function, we employ binary cross entropy, and we have developed two distinct techniques, one for discriminator loss and the other for generator loss. Method of discriminator loss, this method quantifies the discriminator's ability to discern between actual and fraudulent images. It compares the discriminator's predictions on genuine photos to an array of 1s and on fake (made) images to an array of 0. And with regard to generator loss, it quantifies how effectively it was able to fool the discriminator. The discriminator will intuitively classify the bogus images as real if the generator is operating correctly (or 1). Here, the discriminator judgments on the created images are compared against an array of 1s.

4.8 WGAN-GP

The original Wasserstein GAN uses the Wasserstein distance to create a function f that, in terms of mathematical properties, outperforms the functional form used in the original GAN publication. The discriminator or critic must be in the domain of 1-Lipschitz functions for WGAN to function. The authors advised weight loss as a means of meeting this criterion. While weight clipping is effective, it is not necessarily the optimal tactic for enforcing the 1-Lipschitz constraint and may lead to unexpected results. The WGAN-GP method is a viable alternative to weight clipping for preventing training errors. In lieu of weight clipping, the researchers propose a "gradient penalty," which is obtained by adding a loss term that ensures the L2 norm of the discriminator gradients remains near 1.

4.9 Data Preprocessing

To train our WGAN-GP, we will utilize the existing OCT image dataset. In addition, it has been previously cropped to a grayscale image with the dimensions (28,28,1) and the label 4.

4.10 WGAN-GP Model Generator

For each training batch, we were required to train the generator and calculate generator loss. After training the discriminator, the loss is determined and the gradient penalty is computed. After determining the gradient penalty, it must be multiplied by a predetermined weight factor. After performing the multiplication, the gradient penalty must be included in the gradient loss, and the discriminator and generator losses must be returned as a dictionary. Prior to training our model, we must first train the discriminator. Unlike the generator, the discriminator is trained in three steps. We employ a learning rate of 0.0002 and beta values of 0.5 and 0.9 for our generator and discriminator optimizers, respectively (shown in figure 4.6).

Model: "generator"

Layer (type)	Output Shape	Param #
input_2 (InputLayer)	[(None, 128)]	0
dense_1 (Dense)	(None, 4096)	524288
batch_normalization (Batch Normalization)	(None, 4096)	16384
leaky_re_lu_4 (LeakyReLU)	(None, 4096)	0
reshape (Reshape)	(None, 4, 4, 256)	0
up_sampling2d (UpSampling2D)	(None, 8, 8, 256)	0
conv2d_4 (Conv2D)	(None, 8, 8, 128)	294912
batch_normalization_1 (Batch Normalization)	(None, 8, 8, 128)	512
leaky_re_lu_5 (LeakyReLU)	(None, 8, 8, 128)	0
up_sampling2d_1 (UpSampling2D)	(None, 16, 16, 128)	0
conv2d_5 (Conv2D)	(None, 16, 16, 64)	73728
batch_normalization_2 (Batch Normalization)	(None, 16, 16, 64)	256
leaky_re_lu_6 (LeakyReLU)	(None, 16, 16, 64)	0
up_sampling2d_2 (UpSampling2D)	(None, 32, 32, 64)	0
conv2d_6 (Conv2D)	(None, 32, 32, 1)	576
batch_normalization_3 (Batch Normalization)	(None, 32, 32, 1)	4
activation (Activation)	(None, 32, 32, 1)	0
cropping2d (Cropping2D)	(None, 28, 28, 1)	0

=====
Total params: 910,660
Trainable params: 902,082
Non-trainable params: 8,578
=====

Figure 4.5: Architecture of the WGAN-GP Generator

4.11 WGAN-GP Model Discriminator

The shape of the dataset of samples is as follows: (28, 28, 1). In order to create the discriminator, we are going to apply strided convolution in each of the layers. As a consequence, the form that we end up with will have odd dimensions. Additionally, we are altering the geometry of the input by using a "zero pad" in order to shift it to (32,32,1) for each sample.

```

Model: "discriminator"

```

Layer (type)	Output Shape	Param #
input_1 (InputLayer)	[(None, 28, 28, 1)]	0
zero_padding2d (ZeroPadding 2D)	(None, 32, 32, 1)	0
conv2d (Conv2D)	(None, 16, 16, 64)	1664
leaky_re_lu (LeakyReLU)	(None, 16, 16, 64)	0
conv2d_1 (Conv2D)	(None, 8, 8, 128)	204928
leaky_re_lu_1 (LeakyReLU)	(None, 8, 8, 128)	0
dropout (Dropout)	(None, 8, 8, 128)	0
conv2d_2 (Conv2D)	(None, 4, 4, 256)	819456
leaky_re_lu_2 (LeakyReLU)	(None, 4, 4, 256)	0
dropout_1 (Dropout)	(None, 4, 4, 256)	0
conv2d_3 (Conv2D)	(None, 2, 2, 512)	3277312
leaky_re_lu_3 (LeakyReLU)	(None, 2, 2, 512)	0
flatten (Flatten)	(None, 2048)	0
dropout_2 (Dropout)	(None, 2048)	0
dense (Dense)	(None, 1)	2049

```

=====
Total params: 4,305,409
Trainable params: 4,305,409
Non-trainable params: 0
=====

```

Figure 4.6: Architecture of the WGAN-GP Discriminator

If we are not careful when upsampling in the generator component of the network, we will not acquire the same output shape as the images from the original dataset. We will instead obtain a different shape (shown in figure 4.7).

4.12 Architecture of the Proposed Model

We have ensured that our GAN model would run in a distributed environment with privacy protection. We will call it Distributed GAN. For this model, we distributed the standard DCGAN models to the client, and the clients will transmit back the discriminator weights of their side for synchronization, as well as the weights of each client-side generator to the main server generator for synchronization. This will result in an improved main server generator and discriminator. Then the generator can be used for data augmentation while protecting the privacy of the client's data.

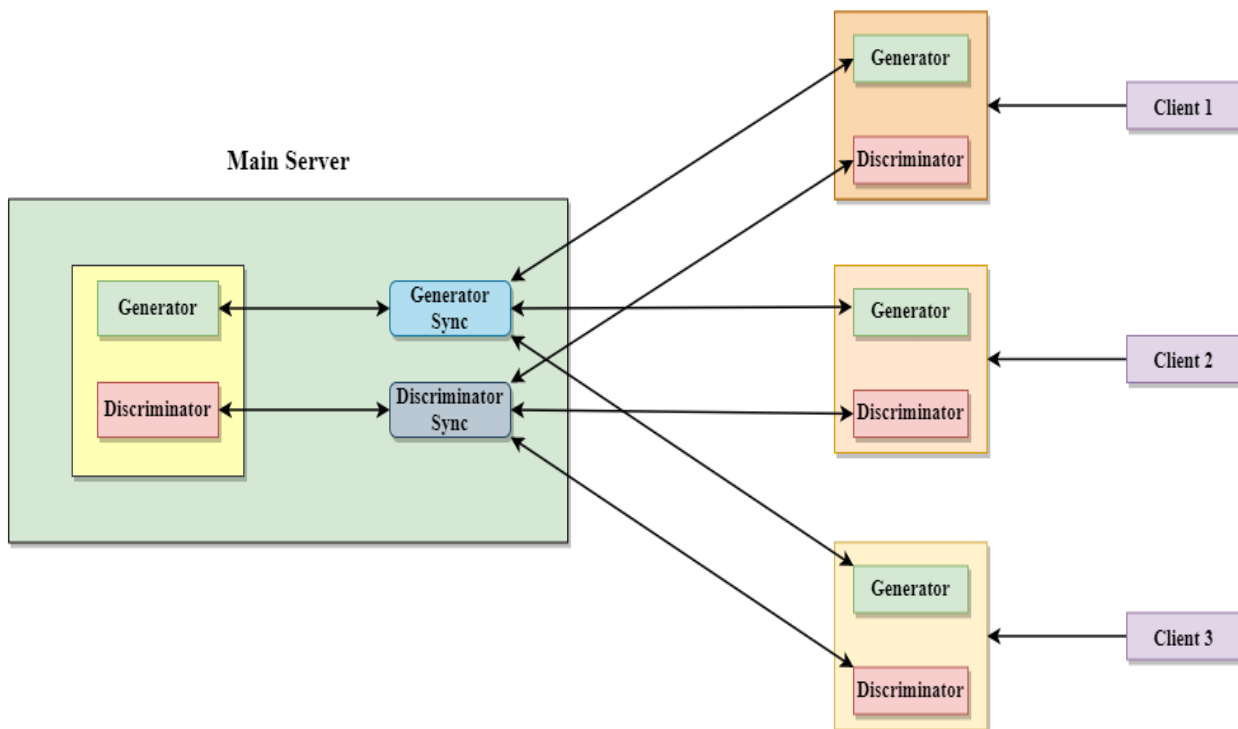


Figure 4.7: Proposed Model Architecture

4.13 Generator

In our generator model, there is one dense layer and three Conv2Dtranspose layers, all of which include the activation function LeakyReLU (Shown in figure 4.8). In a Federated learning environment, our GAN model contains both the discriminator and the generator, and it may be distributed across several networks.

4.14 Discriminator

Our discriminator model will take as input $(28,28,1)$ images from the real dataset and the generator. The discriminator model consists of three conv2D layers, batch normalization, LeakyReLU, and its activation function (shown in figure 4.9).

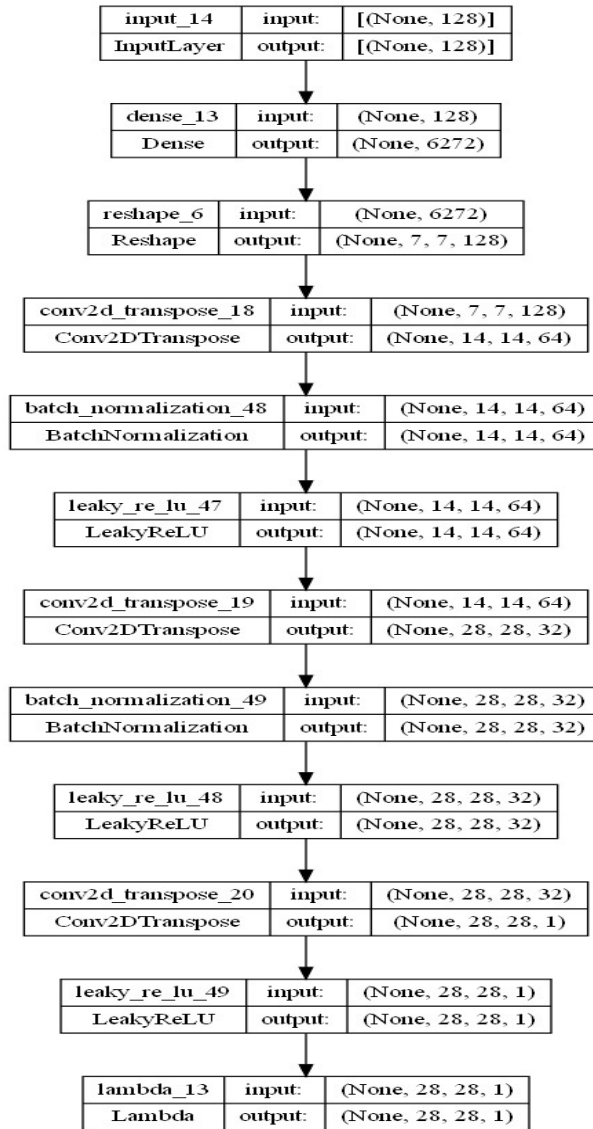


Figure 4.8: Architecture of the Distributed GAN Generator

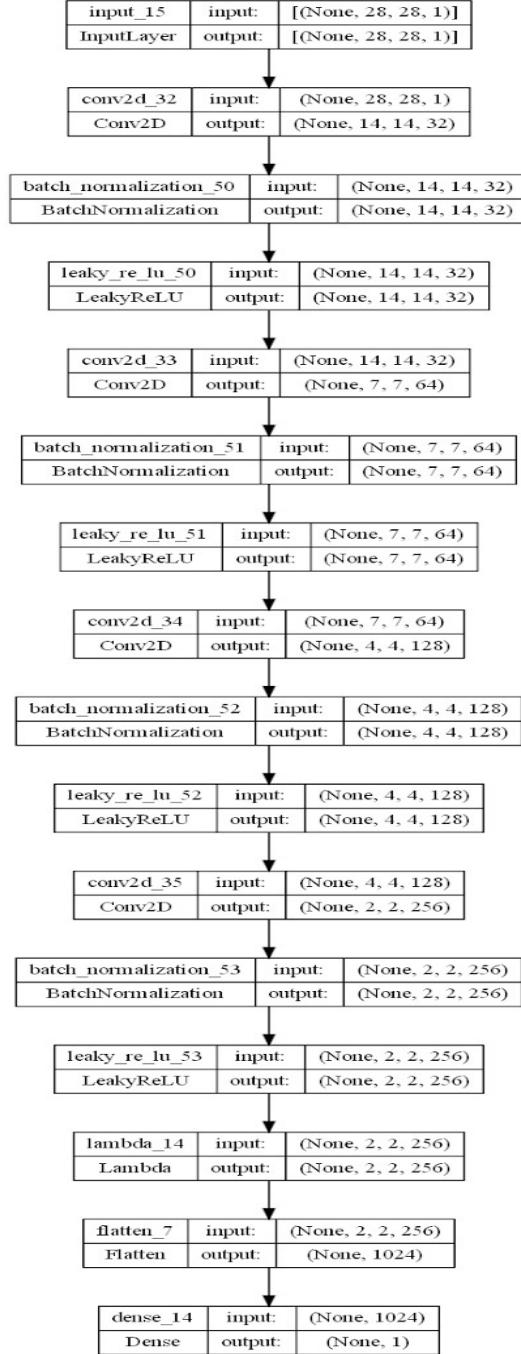


Figure 4.9: Architecture of the Distributed GAN Discriminator

Chapter 5

Results and Analysis

5.1 Evaluating GAN

GANs are a highly difficult topic in a research field that has made remarkable strides in recent years. We fed random noise to a GAN to generate this image. However, there is no exact method for judging the realism of these images. We have no idea how many pixels it should produce. Given this noise vector, it is unclear what pixels we are intended to generate. In addition, the discriminator that distinguishes between true and false in this GAN is never perfect. And frequently overfits to distinguishing between real and fake images for its particular generator. It will probably certainly take into account a huge number of images from our generator. They are deceptive despite their lifelike appearance because they can distinguish certain features. Consequently, the generator may exhibit minute, but occasionally noticeable, characteristics. Consequently, there are no perfect or universal discriminators that can compare two generators and claim which is superior to the other.

5.2 DCGAN

The DCGAN training loop begins with the input of a random seed to the generator. This seed is subsequently used to generate an image. The discriminator is then used to distinguish between genuine and fake images. For each of these models, the loss is calculated, and the gradients are utilized to update the generator and discriminator. Then, we will train our model over 100 epochs with a noise dimension of 100, representing the generator's inputs. The DCGAN was then trained in 64 batches. We will produce 16 images while training our model. The output of the model is (shown in figure 5.1).

The discriminator loss for the real image (blue) and the discriminator loss for the fake images (red) are depicted in one of the subplots in the figure above (orange). A green line represented the generating loss. It is clear after 500 epochs that all three losses are erratic and then stabilizes. The blue line indicates the true discriminator accuracy, whereas the orange line shows the false discriminator accuracy (Shown in figure 5.2).




Figure 5.1: Output Images of the DCGAN




Figure 5.2: Comparison between Generator and Discriminator Loss

5.3 WGAN-GP

During WGAN-GP model training, we were required to train the generator and calculate the generator loss for each batch. The discriminator is then trained, while the loss and gradient penalty are calculated. Once the gradient penalty has been computed, it must be multiplied by a predetermined weight factor. After multiplication, the gradient penalty must be included in the gradient loss, and the discriminator and generator losses must be returned as a dictionary.




Figure 5.3: Output Images of the WGAN

Before beginning model training, we must first train the discriminator. In contrast to the generator, which was trained in a single step, the discriminator is being trained in three steps. We employ a learning rate of 0.0002 and beta values of 0.5 and 0.9 for our generator and discriminator optimizers. Model output is (shown in figure 5.1).

5.4 Distributed GAN

We used 83384 OCT images to train our model, and we created one discriminator and one generator that will be distributed across three separate networks. In each of these neural networks, a dropout layer is typically added.

Instead, we chose batch normalization since it accelerates the learning process. We chose a learning rate of 0.003, a noise size of 128, and a batch size of 64 to train the model. After training our model, we discovered the following results (shown in figure 5.4).




Figure 5.4: Output Images of the Distributed GAN

Chapter 6

Conclusion

In the realm of deep learning, the gravity of a substantial dataset knows no bounds. As a result, the dearth of medical data impedes research in a large field. With the help of state-of-the-art technology like GAN, we can compensate for the data lacking by augmentation all the while protecting the privacy of the sensitive data. Using an OCT image dataset, this study demonstrates that while DCGAN and WGAN contribute to data augmentation, they fail to secure data in a federated environment. Following that we introduced Distributed GAN which inherently conserves data privacy, and surpasses the limitations of federated learning such as mode collapse and data leakage. By utilizing Distributed GAN, we were able to accomplish our two key objectives of data augmentation and protection, despite our inability to produce high-quality images. By experimenting with multiple GANs, we were able to establish the relationship between GANs and data preservation, as well as their pros and downsides.

6.1 Challenges

The most glaring issue we faced was a lack of computational capacity. Due to insufficient GPU and RAM, we were forced to resize our images to $(28*28)$. This negatively impacted our output image quality, and we were compelled to reduce the number of epochs because the training was time-consuming.

6.2 Future Works

In this study, we trained our model using an OCT image dataset, and we hope to expand our research in the future to include non-iid data. In addition, it is impossible to evaluate GAN output beyond the FID score since precise head-to-head comparisons across different discriminators and generators are a long-term aim that is now unattainable. In the near future, we intend to investigate the FID scores of our DCGAN, WGAN, and Distributed GAN autonomously. Moreover, we plan to train multiple discriminators and generators on the central server of our distributed GAN. Due to their ability to be distributed among multiple clients, the client-server retracts the weights from the client side and updates the main server. This modification might solidify the data leakage risk and mode collapse of federated learning.

Bibliography

- [1] R. M. Haralick, K. Shanmugam, and I. H. Dinstein, “Textural features for image classification,” *IEEE Transactions on systems, man, and cybernetics*, no. 6, pp. 610–621, 1973.
- [2] D. Huang, E. A. Swanson, C. P. Lin, *et al.*, “Optical coherence tomography,” *science*, vol. 254, no. 5035, pp. 1178–1181, 1991.
- [3] M. R. Hee, J. A. Izatt, E. A. Swanson, *et al.*, “Optical coherence tomography of the human retina,” *Archives of ophthalmology*, vol. 113, no. 3, pp. 325–332, 1995.
- [4] K. L. Boyer, A. Herzog, and C. Roberts, “Automatic recovery of the optic nervehead geometry in optical coherence tomography,” *IEEE transactions on medical imaging*, vol. 25, no. 5, pp. 553–570, 2006.
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*, Ieee, 2009, pp. 248–255.
- [6] M. Wojtkowski, “High-speed optical coherence tomography: Basics and applications,” *Applied optics*, vol. 49, no. 16, pp. D30–D61, 2010.
- [7] P. Domingos, “A few useful things to know about machine learning,” *Communications of the ACM*, vol. 55, no. 10, pp. 78–87, 2012.
- [8] K. Clark, B. Vendt, K. Smith, *et al.*, “The cancer imaging archive (tcia): Maintaining and operating a public information repository,” *Journal of digital imaging*, vol. 26, no. 6, pp. 1045–1057, 2013.
- [9] W. G. Van Panhuis, P. Paul, C. Emerson, *et al.*, “A systematic review of barriers to data sharing in public health,” *BMC public health*, vol. 14, no. 1, pp. 1–9, 2014.
- [10] M. D. Zeiler and R. Fergus, “Visualizing and understanding convolutional networks,” in *European conference on computer vision*, Springer, 2014, pp. 818–833.
- [11] C. Dong, C. C. Loy, K. He, and X. Tang, “Image super-resolution using deep convolutional networks,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 2, pp. 295–307, 2015.
- [12] A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” *arXiv preprint arXiv:1511.06434*, 2015.

- [13] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*, Springer, 2015, pp. 234–241.
- [14] F. Milletari, N. Navab, and S.-A. Ahmadi, “V-net: Fully convolutional neural networks for volumetric medical image segmentation,” in *2016 fourth international conference on 3D vision (3DV)*, IEEE, 2016, pp. 565–571.
- [15] H.-C. Shin, H. R. Roth, M. Gao, *et al.*, “Deep convolutional neural networks for computer-aided detection: Cnn architectures, dataset characteristics and transfer learning,” *IEEE transactions on medical imaging*, vol. 35, no. 5, pp. 1285–1298, 2016.
- [16] V. Tresp, J. M. Overhage, M. Bundschuh, S. Rabizadeh, P. A. Fasching, and S. Yu, “Going digital: A survey on digitalization and large-scale data analytics in healthcare,” *Proceedings of the IEEE*, vol. 104, no. 11, pp. 2180–2206, 2016.
- [17] Z. A. Ali, K. Karimi Galougahi, A. Maehara, *et al.*, “Intracoronary optical coherence tomography 2018: Current status and future directions,” *JACC: Cardiovascular Interventions*, vol. 10, no. 24, pp. 2473–2487, 2017.
- [18] T. Klein and R. Huber, “High-speed oct light sources and systems,” *Biomedical optics express*, vol. 8, no. 2, pp. 828–859, 2017.
- [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017.
- [20] C. Ledig, L. Theis, F. Huszár, *et al.*, “Photo-realistic single image super-resolution using a generative adversarial network,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4681–4690.
- [21] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Artificial intelligence and statistics*, PMLR, 2017, pp. 1273–1282.
- [22] C. Nøhr, L. Parv, P. Kink, *et al.*, “Nationwide citizen access to their health data: Analysing and comparing experiences in denmark, estonia and australia,” *BMC health services research*, vol. 17, no. 1, pp. 1–11, 2017.
- [23] I. O. Tolstikhin, S. Gelly, O. Bousquet, C.-J. Simon-Gabriel, and B. Schölkopf, “Adagan: Boosting generative models,” *Advances in neural information processing systems*, vol. 30, 2017.
- [24] D. S. Kermany, M. Goldbaum, W. Cai, *et al.*, “Identifying medical diagnoses and treatable diseases by image-based deep learning,” *Cell*, vol. 172, no. 5, pp. 1122–1131, 2018.
- [25] D. S. Kermany, M. Goldbaum, W. Cai, *et al.*, “Identifying medical diagnoses and treatable diseases by image-based deep learning,” *Cell*, vol. 172, no. 5, pp. 1122–1131, 2018.
- [26] Y. Ma, X. Chen, W. Zhu, X. Cheng, D. Xiang, and F. Shi, “Speckle noise reduction in optical coherence tomography images based on edge-sensitive cgan,” *Biomedical optics express*, vol. 9, no. 11, pp. 5129–5146, 2018.

- [27] Y. Xue, T. Xu, H. Zhang, L. R. Long, and X. Huang, “Segan: Adversarial network with multi-scale l1 loss for medical image segmentation,” *Neuroinformatics*, vol. 16, no. 3, pp. 383–392, 2018.
- [28] J. R. Zech, M. A. Badgeley, M. Liu, A. B. Costa, J. J. Titano, and E. K. Oermann, “Confounding variables can degrade generalization performance of radiological deep learning models,” *arXiv preprint arXiv:1807.00431*, 2018.
- [29] D.-K. Hwang, C.-C. Hsu, K.-J. Chang, *et al.*, “Artificial intelligence-based decision-making for age-related macular degeneration,” *Theranostics*, vol. 9, no. 1, p. 232, 2019.
- [30] X. Yi, E. Walia, and P. Babyn, “Generative adversarial network in medical imaging: A review,” *Medical image analysis*, vol. 58, p. 101 552, 2019.
- [31] Q. Chang, H. Qu, Y. Zhang, *et al.*, “Synthetic learning: Learn from distributed asynchronous discriminator gan without sharing medical image data,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 13 856–13 866.
- [32] Z. Dong, J. Men, Z. Yang, *et al.*, “Flynet 2.0: Drosophila heart 3d (2d+ time) segmentation in optical coherence microscopy images using a convolutional long short-term memory neural network,” *Biomedical Optics Express*, vol. 11, no. 3, pp. 1568–1579, 2020.
- [33] I. Goodfellow, J. Pouget-Abadie, M. Mirza, *et al.*, “R1,” *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [34] T. Hassan, M. U. Akram, N. Werghi, and M. N. Nazir, “Rag-fw: A hybrid convolutional framework for the automated extraction of retinal lesions and lesion-influenced grading of human retinal pathology,” *IEEE journal of biomedical and health informatics*, vol. 25, no. 1, pp. 108–120, 2020.
- [35] L. Huang, Y. Yin, Z. Fu, S. Zhang, H. Deng, and D. Liu, “Loadaboost: Loss-based adaboost federated machine learning with reduced computational complexity on iid and non-iid intensive care data,” *Plos one*, vol. 15, no. 4, e0230706, 2020.
- [36] L. Ma, R. Shuai, X. Ran, W. Liu, and C. Ye, “Combining dc-gan with resnet for blood cell image classification,” *Medical & biological engineering & computing*, vol. 58, no. 6, pp. 1251–1264, 2020.
- [37] H. Qu, Y. Zhang, Q. Chang, Z. Yan, C. Chen, and D. Metaxas, “Learn distributed gan with temporary discriminators,” in *European Conference on Computer Vision*, Springer, 2020, pp. 175–192.
- [38] X. Li, Y. Jiang, J. J. Rodriguez-Andina, H. Luo, S. Yin, and O. Kaynak, “When medical images meet generative adversarial network: Recent development and research opportunities,” *Discover Artificial Intelligence*, vol. 1, no. 1, pp. 1–20, 2021.
- [39] X. Zhou, S. Qiu, P. S. Joshi, *et al.*, “Enhancing magnetic resonance imaging-driven alzheimer’s disease classification performance using generative adversarial learning,” *Alzheimer’s research & therapy*, vol. 13, no. 1, pp. 1–11, 2021.