2024-01-15

Using Image Translation To Synthesize Amyloid Beta From Structural MRI

Vega Lara, Fernando

UNIVERSITY OF CALGARY

Using Image Translation To Synthesize Amyloid Beta From Structural MRI

by

Fernando Vega Lara

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE

DEGREE OF MASTER OF SCIENCE

GRADUATE PROGRAM IN BIOMEDICAL ENGINEERING

CALGARY, ALBERTA

JANUARY, 2024

Abstract

Amyloid-beta and brain atrophy are known hallmarks of Alzheimer's Disease (AD) and can be quantified with positron emission tomography (PET) and structural magnetic resonance imaging (MRI), respectively. PET uses radiotracers that bind to amyloid-beta, whereas MRI can measure brain morphology. PET scans have limitations including cost, invasiveness (involve injections and ionizing radiation exposure), and have limited accessibility, making PET not practical for screening early-onset AD. Conversely, MRI is a cheaper, less-invasive (free from ionizing radiation), and is more widely available, however, it cannot provide the necessary molecular information. There is a known relationship between amyloid-beta and brain atrophy. This thesis aims to synthesize amyloid-beta PET images from structural MRI using image translation, an advanced form of machine learning. The developed models have reported high-similarity metrics between the real and synthetic PET images and high-degree of accuracy in radiotracer quantification. The results are highly impactful as it enables amyloid-beta measurements form every MRI, for free!

Preface

This thesis is an original work from the author, it is the result of the effort involving all co-authors resulting in one publication where I am the first author.

- Vega, F., Addeh, A., Ganesh, A., Smith, E.E. and MacDonald, M.E. (2023), Image Translation for Estimating Two-Dimensional Axial Amyloid-Beta PET From Structural MRI. J Magn Reson Imaging. https://doi.org/10.1002/jmri.29070 (Vega et al., 2023)

Parts of this thesis are derived from the following international conference proceedings:

- Vega F., Addeh A., & MacDonald M. E. Three-Dimensional Amyloid-Beta PET Synthesis from Structural MRI with Conditional Generative Adversarial Networks. International Society of Magnetic Resonance in Medicine 2024 (In review)

- Vega F., Addeh A., & MacDonald M. E. Amyloid-Beta Axial Plane PET Synthesis from Structural MRI: An Image Translation Approach for Screening Alzheimer's Disease. International Society of Magnetic Resonance in Medicine 2023 arXiv preprint arXiv:2309.00569 (Vega et al., 2023a).

- Vega F., Addeh A., & MacDonald M. E. Amyloid-Beta PET Synthesis from Structural MRI: A Potential Alternative Method for Alzheimer's Disease Screening. Alzheimer's Association International Conference 2023 (Vega et al., 2023b).

- Vega F., Addeh A., & MacDonald M.E. Preliminary Findings of Image Translation Models for the Estimating Amyloid Beta from Structural MRI. Brain & Brain PET 2022, International Society of Cerebral Blood Flow and Metabolism.

Additionally, Chapter 4 of this thesis resulted in a manuscript that will be submitted to the Journal of Magnetic Resonance in Medicine. It should be noted that work outside of the scope of this thesis was also presented and published as conference proceedings. Additional first-author conference proceedings:

- Vega F., Addeh A., & MacDonald M. E Denoising Low-Field Simulated MRI Using Cycle Consistent Generative Adversarial Networks. International Society of Magnetic Resonance in Medicine 2023 (Vega et al., 2023c)

Additional second-author journal papers:

- Addeh A., Vega F., Raj M.P., Williams R. J., Pike G.B., MacDonald M.E. (2023). Direct Machine Learning Reconstruction of Respiratory Variation Waveforms from Resting-State fMRI Data in a Pediatric Population. Neuroimage (Addeh et al., 2023).

Additional co-author international conference proceedings:

- Adila K.L., Munro E., Vega F., Mohite A., Curtis C., Tyndall V. A., MacDonald M.E. Using Machine Learning to Study the Effects of Genetic Predisposition on Brain Aging in the UK Biobank. IEEE International Symposium on Biomedical Imaging 2023.

- Abdoljalil A., Ardial-Lopez K., Vega F., Golestani A., MacDonald M.E. Limitations of the Derived Respiratory Variation Measurements Used in Functional Magnetic Resonance Imaging. IEEE International Symposium on Biomedical Imaging 2023.

- Abdoljalil A., Vega F., MacDonald M.E. Using BOLD-fMRI to Compute Respiration Volume per Time (RVT) and Respiration Variation (RV) with Convolutional Neural Network (CNN) in Children. International Society of Magnetic Resonance in Medicine 2023

- MacDonald M.E., Fikre E., Vega F., Addeh A., Simulation Evidence for use of a Denoising Auto-Encoder (DAE) to Improve Ultra-Low Field (64mT) MRI with High Field (3T) International Society of Magnetic Resonance in Medicine 2022

Finally, three local conference proceedings:

- Vega F., Addeh A., MacDonald M.E. Towards a Non-Invasive Method for Screening Alzheimer's Disease: Image Translation to Estimate Amyloid-Beta Positron from Structural MRI. Podium Presentation at Alberta Biomedical Engineering Conference 2022

- Addeh A., Vega F., MacDonald M.E. Challenges of Respiratory Signals Measurements in fMRI Studies. Poster Presentation at Alberta Biomedical Engineering Conference 2022

- Elmenshawi A., Vega F., MacDonald M.E. Enhancing the Preprocessing Pipeline for Image Translation Models to make Beta-Amyloid Positron Emission Tomography (PET) from Structural Magnetic Resonance Imaging (MRI). University of Calgary, Biomedical Undergraduate Research Symposium 2022

Acknowledgements

I would like to extend my heartfelt gratitude to the following individuals and institutions, whose support was instrumental in the completion of my research and my master's degree.

I am profoundly thankful to my supervisor, Dr. M. Ethan MacDonald, for his exceptional mentorship, unwavering support, and patience throughout my master's journey. His guidance was crucial in keeping me motivated and aiding me through numerous conference and paper deadlines. My gratitude also extends to Dr. Aravind Ganesh and Dr. Eric E. Smith for their invaluable support during paper revisions, and to Dr. Roberto Souza, whose expertise in machine learning greatly enriched my research. Furthermore, I would like to thank Dr. Richard Frayne for serving an as external examiner for this thesis.

This project would not have been possible without the generous donation from Jim Gwynne and the financial support from the Alberta Graduate Excellence Scholarship. I am also grateful to the Hotchkiss Brain Institute and the Biomedical Engineering Department for providing conference travel funding. A special thanks goes to the Open Access Series of Imaging Studies team for granting access to the images crucial for my thesis.

I am equally thankful to my labmates from the Brain Imaging Team and my friends, who created an enjoyable and supportive environment throughout this journey.

Lastly, my deepest appreciation goes to my family, including my mother Alma, father Fernando, and sister Rocio, who have been my pillars of strength and support. Their immense sacrifices and constant encouragement have been fundamental to my success. I also want to express my gratitude to my partner Maria, whose love, motivational messages, and calls were a constant source of motivation.

I would not have been able to do this thesis without all the support from the mentioned people, I look forward to carrying the lessons and experiences from this experience into my future endeavors and wish the best for each of you.

In memory of my grandparents

Rafael Lara Ruiz, Maria de Jesus Sanchez Garcia, and Raul Vega Rios - En Paz Descansen

Table of Contents

List of Figures

List of Tables

List of Symbols

Symbol	Definition
18-AV-45	Florbetapir
AD	Alzheimer's Disease
ADNI	Alzheimer's Disease Neuroimage Initiative (Dataset)
AIBL	Australian Imaging Biomarkers and Lifestyle (Dataset)
ANTs	Advanced Normalization Tools
APP	Amyloid Precursor Protein
ASL	Arterial Spin Labeling
BET	Brain Extraction Tool
CBF	Cerebral Blood Flow
CDR	Clinical Dementia Rating
cGAN	Conditional Generative Adversarial Network
CN	Cognitively Normal
CNN	Convolutional Neural Network
CSF	Cerebral Spinal Fluid
Cycle-GAN	Cycle Consistent Generative Adversarial Networks
FDG	Fluodeoxyglucose
FID	Free Induction Decay
FLAIR	Fluid Attenuated Inversion Recovery
GAN	Generative Adversarial Network
ICBM	International Consortium of Brain Mapping (Brain Atlas)
MAE	Mean Absolute Error
MCI	Mild Cognitive Impairment
MNI	Montreal Neurological Institute (Standard Brain Template)

MP-RAGE	Magnetization Prepared - Rapid Gradient Echo
MRI	Magnetic Resonance Imaging
OASIS-3	Open Access Series of Imaging Studies 3 (Dataset)
PET	Positron Emission Tomography
PiB	Pittsburgh Compound B
PSNR	Peak Signal-to-Noise Ratio
ReLU	Rectified Linear Unit
RF	Radio Frequency
SSIM	Structural Similarity Index Measure
SUVR	Standard Uptake Value Ratio
T1w	T1-weighted
TE	Echo Time
TR	Repetition Time
ViT	Vision Transformer

Epigraph

"Alzheimer's is not about the past - the successes, the accolades, the accomplishments... Alzheimer's is about the present and the struggle, the scrappy brawl, the fight to live with a disease. It's being in the present, the relationships, the experiences, which is the core of life, the courage to live in the soul"

- Greg O'Brien, Lisa Genova, *"On Pluto: Inside the Mind of Alzheimer's"*

Chapter 1

Introduction

1.1 Background

Alzheimer's Disease (AD) is the predominant form of dementia and is increasingly recognized as one of the most costly, lethal, and burdensome diseases (Scheltens et al., 2021b). AD is characterized by a progressive decline in cognitive and functional abilities associated with aging, ultimately leading to death. It was first described by Alois Alzheimer in 1906 after investigating a post-mortem brain that showed brain atrophy and changes in cortical cell clusters (Tanzi and Bertram, 2005; Stelzmann et al., 1995). Later in 1980, amyloid-beta plaques associated with AD progression were quantified (Soria Lopez et al., 2019). Since then, brain atrophy and amyloid-beta have become hallmarks of AD in neuroimaging studies.

Brain atrophy in AD is characterized by the progressive loss of neurons and the connections between them. This degeneration is most pronounced in the hippocampus and cerebral cortex, regions of the brain critical for memory and cognitive functions (Pini et al., 2016a). Neurodegeneration is closely associated with amyloid-beta accumulation: when amyloid-beta aggregates, it disrupts the synapses between neurons and activates an immune response causing further damage to neurons, leading to the characteristic atrophy observed in AD patients primarily affecting the gray and white matter (Hampel et al., 2021). Therefore, amyloid-beta has been extensively studied with respect to AD (Sadigh-Eteghad et al., 2015; Wang et al., 2006; Chen et al., 2017; Liu et al., 2015; Scheltens et al., 2021a; Ghiso and Frangione, 2002).

Given the pathological hallmarks of AD, several detection and diagnostic techniques have been developed over the past few decades. Neuropsychological assessments, which are non-invasive, aim to evaluate the cognitive performance of a subject in various fields, although these tests are not recommended for determining preclinical stages of AD since the changes in behavior and functionality are not often reflected until reaching a probable AD stage (Eppig et al.,

2020). In contrast, lumbar spine punctures that extract cerebral spinal fluid (CSF) to evaluate the concentration of amyloid-beta are highly specific (Shaw et al., 2018). However, they are invasive and can be painful, which may limit their usage in routine screening, as many participants refuse them (Cummings, 2011). Recently, blood tests have shown their ability to detect the concentrations of immune cell types that are related to AD progression; studies have demonstrated their potential to determine if a subject will transition from mild cognitive impairment (MCI to AD. However, such tests have not yet received approval from the U.S. Food and Drug Administration, (Shigemizu et al., 2022), and thus have limited use.

Medical imaging is important for the assessment of neurodegeneration and dementia, as it has the potential to detect early stages of the disease before clinical symptoms manifest (Broderick et al., 2007). Positron Emission Tomography (PET) imaging can measure information about molecular concentrations using radiotracers that bind with specific molecules (Suppiah et al., 2019). For amyloid-beta, there are two widely used radiotracers: 1) 18F-AV-45 (florbetapir) (Camus et al., 2012) and 2) Pittsburgh Compound B (11C-PIB) (Yamin and Teplow, 2017). However, PET imaging has disadvantages, including its high cost (approximately $5000 to $8000 per PET scan) (Mayerhoefer et al., 2020; Lee et al., 2021), invasiveness (injection of radiotracer and blood sampling), and exposure to ionizing radiation (Brix et al., 2009). Such disadvantages limit the use of PET in many hospitals and clinics and make it inappropriate for population screening and detection of early-onset AD.

MRI is another image modality widely used for AD evaluation (Pini et al., 2016b).It can provide many contrast modes, but this work will primarily focus on T1w MRI. MRI has the advantage of being more affordable, non-invasive, free from ionizing radiation, and more widely available compared to PET (Tikkakoski, 2011). However, it cannot measure amyloid-beta concentrations within the brain. Brain atrophy is known to be associated with amyloid-beta accumulation (Chetelat et al., 2010; Tosun et al., 2011), (as seen in Figure 1.1 illustrating a cognitively normal and AD subject with both MRI and PET), meaning that amyloid-beta PET images could potentially be inferred from structural MRI alone using image translation deep learning models. This strategy could have the benefit of reducing the cost per scan, minimizing invasiveness, decreasing radiation exposure, and offering highly valuable molecular information from cheaper structural MRI scans.

Image translation models are a powerful machine learning technology that can generate one image type from another (Isola et al., 2017). In medical imaging, a particularly valuable application of this technology is generating expensive image types from cheaper ones. While studies have attempted to estimate PET images from structural MRI, many have focused on radiotracers that do not target amyloid-beta (Ouyang et al., 2023; Sikka et al., 2021; Lan et al., 2021). Image translation models are often implemented with a Convolutional Neural Network (CNN) (Yamashita et al., 2018). Earlier applications used a U-Net architecture (Ronneberger et al., 2015); however, Conditional Generative Adversarial Networks (cGAN) have recently been shown to outperform the classic U-Net architecture (Isola et al., 2017).

MRI and PET Image Samples from Cognitively Normal and Alzheimer's Disease Subejct
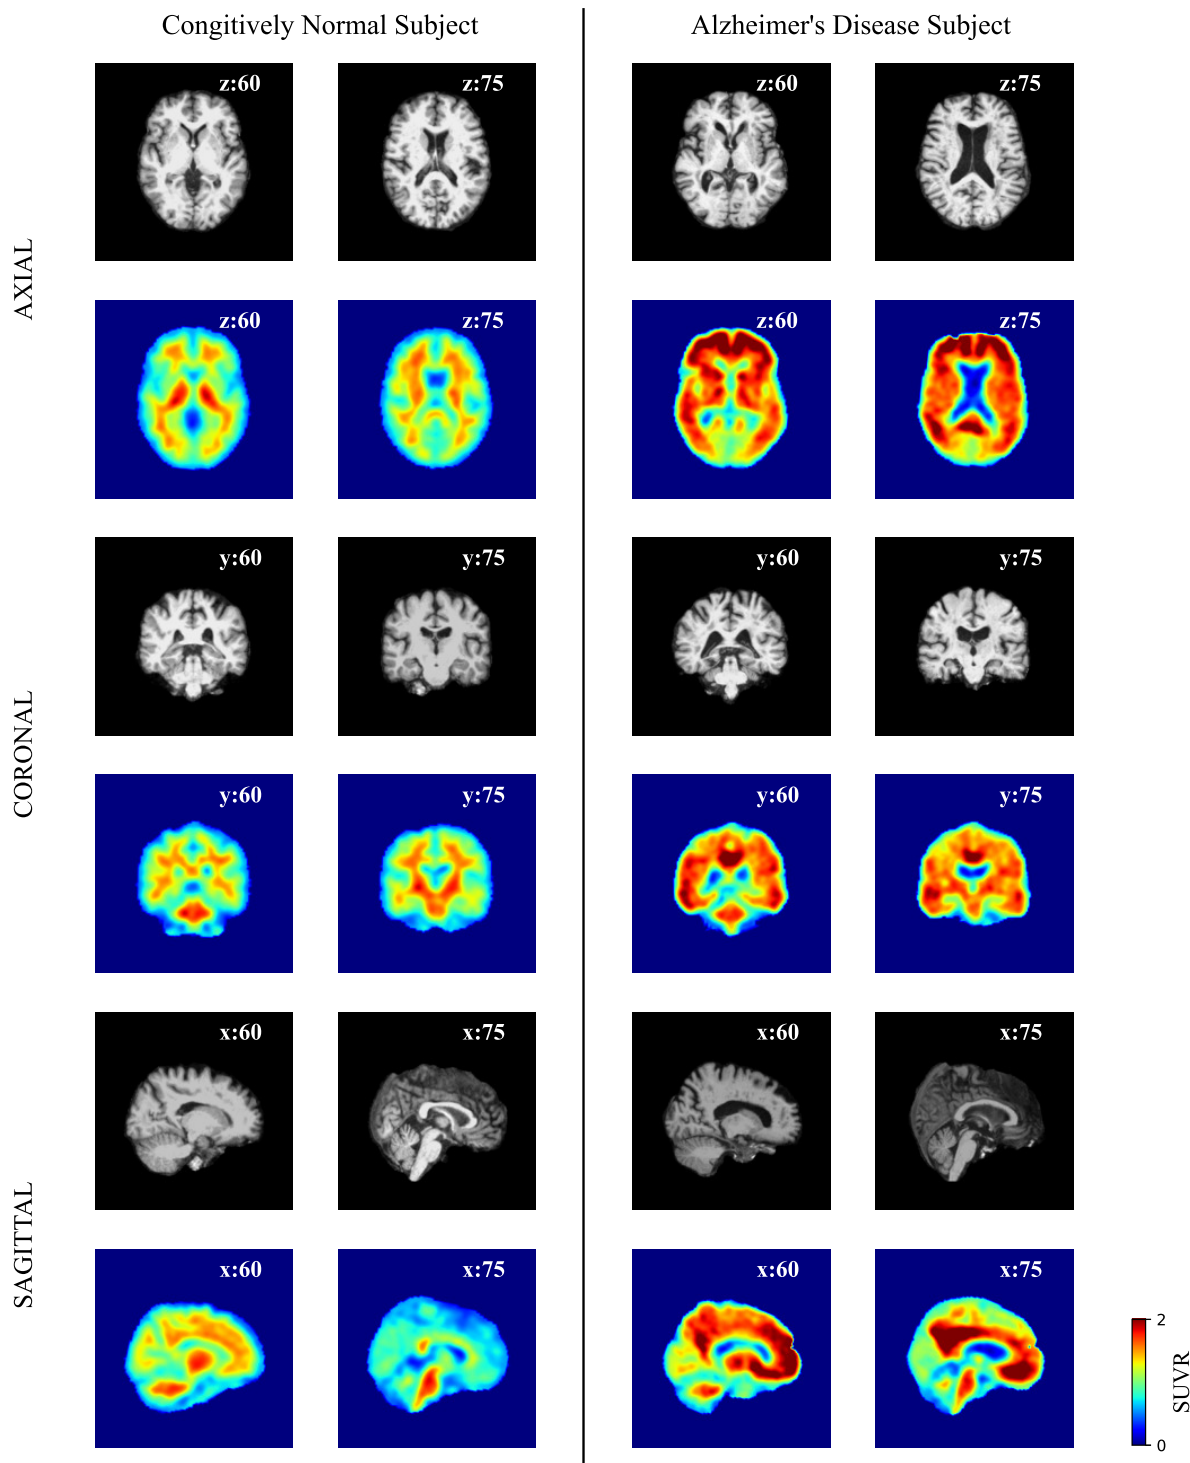
Figure 1.1: MRI and PET images for Cognitively Normal (CN) and Alzheimer's Disease (AD) subjects. Two slices in each orientation are shown with the numbers that correspond to the slice in the Montreal Neurological Institute (MNI) standard brain template. CN subject shows lower levels of brain atrophy (i.e., in the ventricles) with low amyloid-beta burden in comparison with AD subject. AD subject shows more atrophy and higher levels of amyloid-beta.

The cGAN excels in image translation tasks due to its potential to produce synthetic images that are indistinguishable from real ones. They are based on adversarial training where a U-Net (generator) is used to generate images that another CNN architecture (discriminator) classifies as real or synthetic. The generator receives the feedback from the discriminator and, over time, improves the image synthesis. In this way, the generator and discriminator network components are trained against each other until they reach an equilibrium where the most realistic images can be created (Goodfellow et al., 2014; Isola et al., 2017). A more thorough background of CNN and cGAN concepts is provided in Chapter 2, section 2.3.

Medical image translation can be performed with either two-dimensional or three-dimensional images. The two-dimensional approach can use slices from the PET and MRI volume and provide large numbers of images that give more variety to the training process while decreasing the computational load; however, it has been shown that training in two-dimensional slices does not yield consistent results in the direction perpendicular to the slices. The three-dimensional approach often follows a similar training process to the two-dimensional one but uses the whole volume, leading to fewer training samples. While the three-dimensional approach requires more computational resources than the two-dimensional approach, it can provide consistency in all directions of the volume.

To build image translation models, this project used the Open Access Series of Imaging Studies (OASIS-3) (LaMontagne et al., 2019). This dataset comprises over 1000 pairs of T1w MRI and amyloid-beta PiB PET images describing 609 cognitively normal (CN) subjects and 489 at different stages of cognitive decline. Each subject's cognitive state was classified using the clinical dementia rating (CDR) score (Mendez, 2022), which evaluates six different cognitive and behavioral domains such as orientation, memory, problem-solving, judgment, community affairs, home and hobbies performance, and personal care. Then, a score was given for each component that can range on a scale of 0 to 3: CN (CDR=0), questionable dementia (CDR=0.5), MCI (CDR=1), moderate cognitive impairment (CDR=2), and severe cognitive impairment (CDR=3), where a CDR of 0.5 or greater might be associated with AD (Manning and Ducharme, 2010).

The OASIS-3 dataset was curated with basic preprocessing steps for both the MRI and PET images, such as brain extraction for the MRI images and inter-frame motion correction, Centiloid standardization, smoothing, and standard uptake value ratio (SUVR) calculation for the PET images (LaMontagne et al., 2019). To use the images for the image translation model, additional processing steps were performed, such as PET image brain extraction, and image registration, described in Chapter 3, and SUVR normalization, and image normalization as described in Chapter 4.

In summary, this thesis introduces an image translation framework that can generate synthetic amyloid-beta PiB PET images from structural T1w MRI using the OASIS-3 dataset.

1.2 Motivation

It is estimated that dementia has an enormous global economic burden of $1.3 trillion USD in 2019, and this cost is anticipated to increase. Caregivers bear much of this economic burden as well as substantial emotional strife as they witness the patient's mental degeneration. A more timely, safe, and affordable detection technique is urgently needed to achieve earlier detection and treatment, improved care, better management, and potentially enhanced quality of life for both patients and caregivers (Wong, 2020).

AD can be detected by using many techniques that will be discussed in the literature review, such as neuropsychological assessments, lumbar punctures, and neuroimaging. Neuroimaging techniques such as PET and MRI can provide insight into the progression of AD; however, as mentioned previously, regardless of its potential to measure molecular changes within the body, PET scans have several important limitations. MRI scans can provide cheaper and safer images that are widely accesible, however, MRI cannot measure amyloid-beta concentrations.

Developing an image translation model that can synthesize amyloid-beta PET images from structural MRI would make it possible to generate an amyloid-beta PET image for every MRI image. This would significantly reduce the cost, ionizing radiation exposure, and invasiveness of AD diagnostic procedures, enabling broader access to continuous screening of early-onset AD. At the moment of writing this thesis, there is no known peer-reviewed work that has aimed to estimate amyloid-beta PiB PET from structural T1-weighted MRI in this way.

1.3 Hypothesis and Description of Aims

I hypothesize that the structure, signal in gray and white matter, and texture of the brain obtained with structural T1w MRI can be used to infer an individual's amyloid-beta burden represented in PiB PET imaging using image translation deep learning models. This hypothesis is based on the expectation that image translation models can learn from the relationship between changes in the brain reflected in the structural MRI images and amyloid-beta deposition to produce PiB PET images for both cognitively normal subjects and those at various stages of AD. I expect to achieve high degrees of similarity between the synthetic and real PiB PET images (with a structural similarity index measure greater than 0.85)

Developing an image translation model capable of generating amyloid-beta PET images from structural MRI is a problem that has not been addressed before. In order to explore this hypothesis, I took an incremental and rigorous approach with the following specific aims:

1.3.1 Aim 1

The first aim focused on developing a two-dimensional image translation model capable of generating synthetic amyloid-beta PET images from structural MRI. Prior to working with a more complex and computationally intensive three-dimensional model, it was prudent to evaluate the model's performance in a simpler setting to understand the limitations and establish the feasibility of estimating amyloid-beta PET images from structural MRI.

1.3.2 Aim 2

The second aim advanced upon the first aim by developing a three-dimensional image translation model. Building upon the methodology in two-dimensional space, this aim involved a more complex model with fewer training observations and addressed the limitations of the two-dimensional implementation.

1.3.3 Aim 3

The third aim focused on comparing the two-dimensional and three-dimensional implementations. This comparison was performed by training the three-dimensional model developed in Aim 2 with the same training and evaluation cohorts as in Aim 1. The purpose of this aim was to demonstrate the strengths of the three-dimensional approach over the two-dimensional approach and show that the added features in the three-dimensional model improved the work in Aim 1.

1.4 Thesis Outline

This work is divided into 6 chapters:

- Chapter 1 introduces the background of this project with a short description of the role of amyloid-beta and brain atrophy for this thesis, the motivation behind this research, and the specific objectives of the thesis.

- Chapter 2 provides an overview of the foundational concepts regarding amyloid-beta, brain atrophy, medical imaging, and image translation deep learning models, with an analysis on work focused on the estimation of PET images from structural MRI.

- Chapter 3 achieves aim 1, a two-dimensional axial amyloid-beta PET synthesis from structural MRI. The methodology includes the PET image preprocessing steps and implementation of the two-dimensional cGAN architecture.

- Chapter 4 achieves aim 2 by implementing a three-dimensional amyloid-beta PET synthesis from structural MRI. This chapter describes the methodology adaptations and enhancements made to improve upon the two-dimensional approach.

- Chapter 5 achieves aim 3, with a rigorous and fair comparison between the two-dimensional and three-dimensional models, highlighting the strengths of training image translation deep learning models using volumes instead of slices.

- Chapter 6 concludes the thesis with a general discussion, highlighting the contributions of this research to the field of AD diagnosis, considering its limitations, and suggesting directions for future research.

Chapter 2

Literature Review

Chapter 1 briefly presented the background, motivation, hypothesis, and aims of the thesis. Chapter 1 helped the reader to briefly understand the rationale and the work done towards this project without delving into key concepts and technical aspects. This chapter presents a more in-depth review of key concepts that appear throughout the thesis. Firstly, it is necessary to understand the role of amyloid-beta and its synthesis within the brain and how this is related to brain atrophy and Alzheimer's Disease (AD). Subsequently, this chapter will review the commonly used methods for detecting AD and their respective strengths and weaknesses. Next, key technologies related to the work in this thesis will be reviewed, including PET, MRI, and image translation. Finally, a review of current literature is presented, showcasing two articles that are most closely related to the work presented in this thesis.

2.1   The Role of Amyloid-Beta and Brain Atrophy in Alzheimer's Disease

Amyloid-beta accumulation is a manifesting factor along the progression of AD that originates from the proteolytic cleavage of a larger protein called amyloid precursor protein (APP) (O'Brien and Wong, 2011). APP is a protein that plays an important role in biological activities, such as intercellular transport, neuronal development, and repair. Under normal conditions, APP is broken down through cleavage by alpha-secretase and gamma-secretase, producing soluble proteins that are cleared through cerebral spinal fluid (CSF) (O'Brien and Wong, 2011; MacLeod et al., 2015). However, when APP is cleaved by beta-secretase rather than alpha-secretase, it can produce a molecule known as amyloid-beta. Amyloid-beta is a peptide conformed by a chain of aminoacids where lengths of 37 to 49 can be found within the body. However, there are two predominant chains of amyloid-beta that can be generated: less toxic and less hydrophobic amyloid-beta-40 and more toxic and hydrophobic amyloid-beta-42. The key difference between amyloid-beta-40 and amyloid-beta-42 is that the most toxic one has two extra aminoacids that are no-polar, making it more hydrophobic and harder to be cleaned from the brain(MacLeod et al., 2015). The cleaving process is illustrated

in Figure 2.1.

Once formed, amyloid-beta-42 aggregates and is not cleared from the brain as efficiently as amyloid-beta-40. Instead, it accumulates between neurons, disrupting their normal function and leading to synaptic failure and cell death, resulting in atrophy (Hampel et al., 2021; Chen et al., 2017). All mentions of amyloid-beta in this thesis are related specifically to amyloid-beta-42.

Amyloid-Beta Creation Process
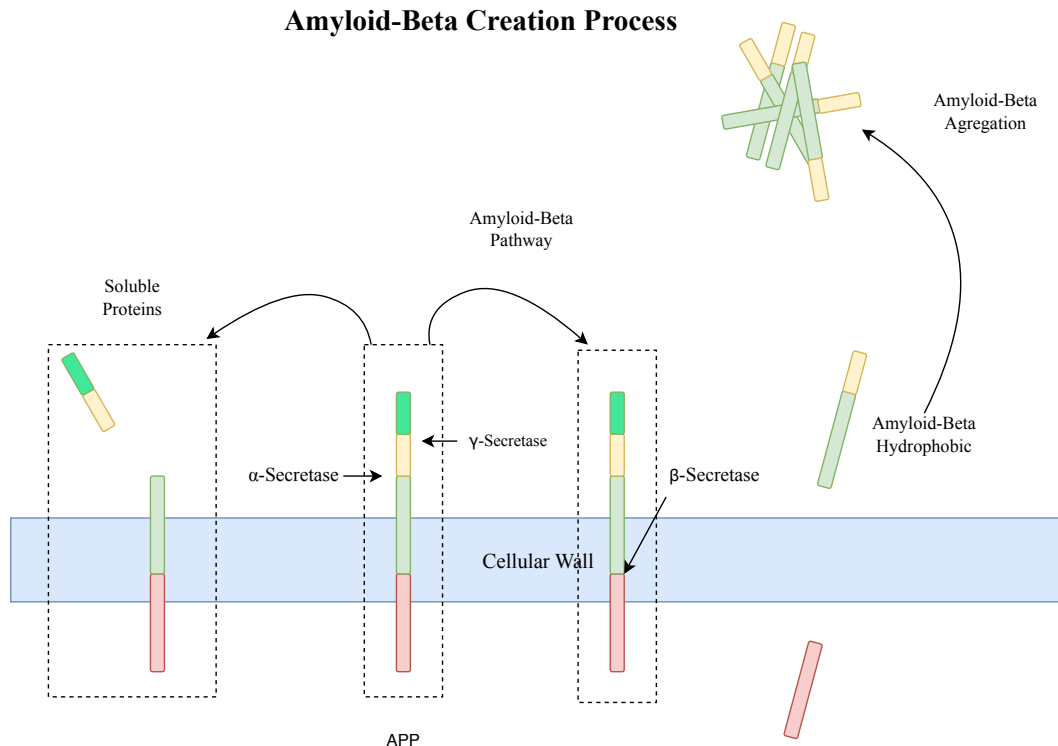
Figure 2.1: An illustration of APP processing and Amyloid-beta accumulation. APP (center, inside dashed box) is synthesized by 2 possible pathways, the alpha-secretase that generates generate soluble proteins (sAPP) and the beta-secretase pathway that generates amyloid-beta. The amyloid-beta aggregates and creates the plaques. Figure adapted from (Gandy, 2005)

As amyloid-beta levels rise, the brain's primary immune cells, known as microglial cells, become activated. These cells, which typically remove damaged cells in healthy brain tissue through phagocytosis, recognize amyloid-beta plaques as harmful entities and initiate the clearance process. Upon activation, microglial cells attempt to phagocytose the amyloid-beta plaques, and in the process, release inflammatory mediators to aid in clearing these plaques. Unfortunately, this response can lead to neuroinflammation, which, if sustained, may contribute to neuronal death and localized brain atrophy.(Heneka et al., 2015).

Brain atrophy is another hallmark of AD and has profound implications on cognitive function (Fox and Schott, 2004). As AD progresses, the brain shrinks noticeably in many areas (Thompson et al., 2003), as shown in Figure 2.2. Usually, brain atrophy alone is not sufficient to diagnose AD, as it is not specific to any one disease; indeed, it can be

present with normal healthy aging, and various other neurological diseases (MacDonald and Pike, 2021)

As the neurons die off, the synapses between them are also lost, causing atrophy and changes in gray and white matter. These changes are not uniform across the brain; certain regions are more vulnerable than others, reflecting the characteristics and symptoms of AD, such as memory loss, difficulty in problem-solving, and challenges in orientation (Whitwell, 2010).
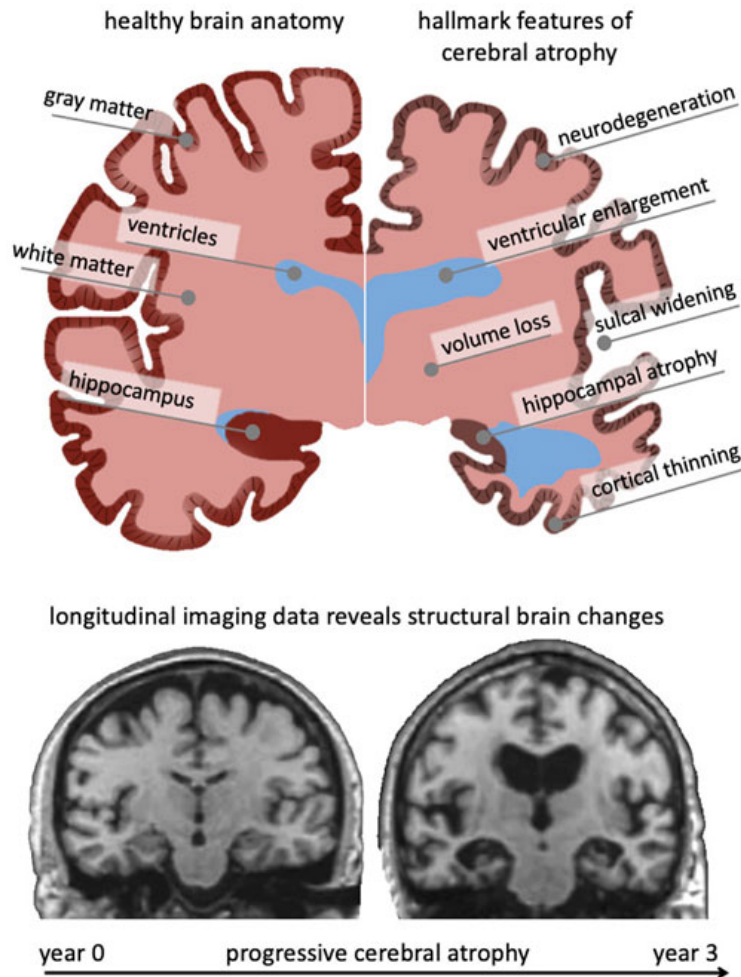
Figure 2.2: Comparison of brain atrophy between a healthy subject (left) and AD subject (right) coronal views. Showing the shrinking of the brain due to AD, Figure retrieved from (Blinkouskaya and Weickenmeier, 2021)

2.1.1 Tau Proteins

Tau is another biomarker of AD progression. While it is not discussed in this thesis, it is still pertinent to describe its importance. Tau proteins are important for supporting internal structural components of the neurons, specifically the microtubules. Microtubules help guide nutrients and molecules from the cell body to the axon and dendrites (Jack et al., 2010). In healthy subjects, tau stabilizes and binds the microtubules. However, in AD subjects, tau detaches

from the microtubules and folds with other tau proteins to form threads known as neurofibrillary tangles within the neurons. These tangles block the neuron's nutrient supply, destabilize the structure of the neuron, and harm the synaptic communication between neurons (Bloom, 2014). Evidence suggests that amyloid-beta toxicity influences tau detaching; rodent model experiments have shown that amyloid-beta promotes a chemical and physical change in tau (Busche and Hyman, 2020).

Given the explained mechanisms that happen in AD, there is a known consensus that describes AD progression as a cascade event that begins with amyloid-beta accumulation in the brain that affects tau which all together provoke changes in brain structure affecting gray and white matter leading to brain atrophy. Finally, manifesting the characteristic clinical symtomps from AD. Figure 2.3 shows the believed hypothesis biomarker progression in AD, the curves were obtained from reviewing existent research, observational studies, and analyzing CSF data, PiB PET data, autopsy studies, structural MRI data, and longitudinal studies (Jack et al., 2010).
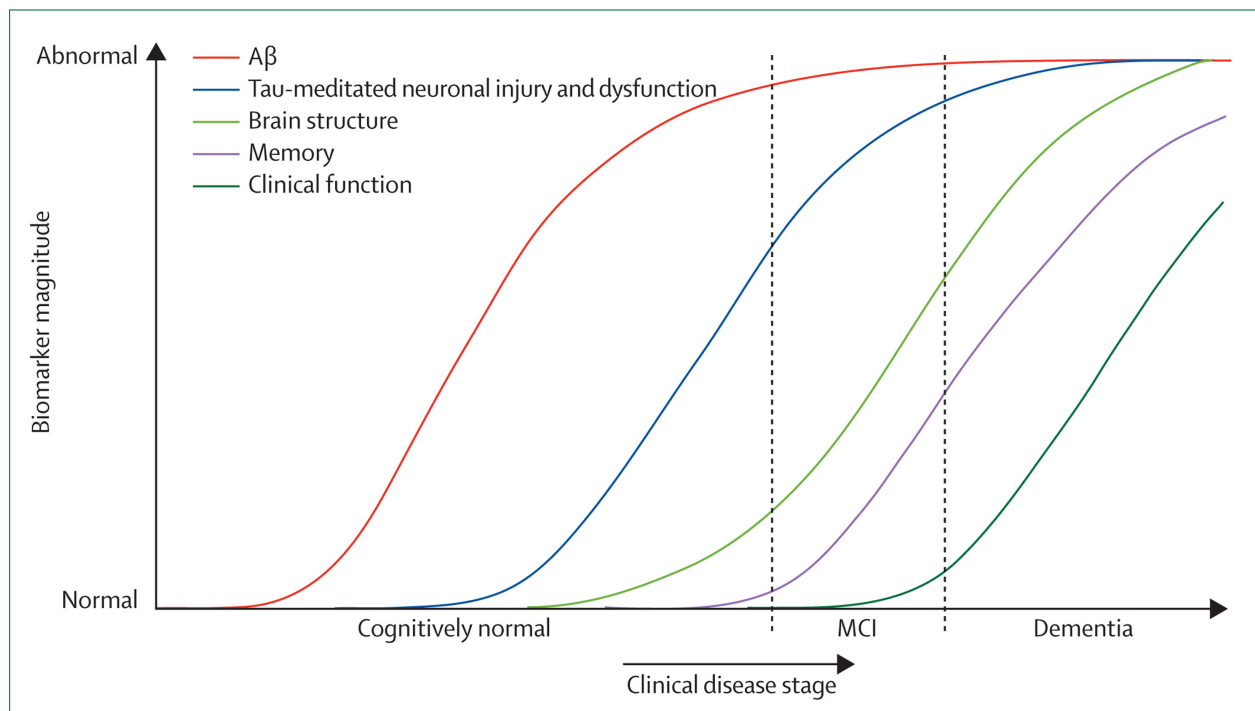
Figure 2.3: Alzheimer's Disease (AD) progression. We see changes in biomarkers, from left to right, amyloid-beta (Aβ) accumulation, tau proteins accumulation, brain atrophy, impact in memory, and impact in clinical function. Figure retrieved from (Jack et al., 2010)

AD is a complex pathology that involves many other factors that were not described in these sections; however, the aforementioned concepts are key concepts to support the work in this thesis.

2.2 Tests to Assess Alzheimer's Disease

Over the past several decades, many AD screening techniques have emerged, the most prominent including: 1) neuropsychological assessments, 2) lumbar punctures for CSF extraction, and 3) neuroimaging (particularly, PET and MRI). Many techniques have emerged due to the complexity of diagnosing AD since it involves many other factors (i.e., genetics, lifestyle), however, these are the most established techniques.

2.2.1 Neuropsychological Tests

A common neuropsychological test for AD is the Clinical Dementia Rating (CDR) score, which evaluates the subject's cognitive performance in six fields: memory, orientation, judgment, problem solving, community affairs, home and hobbies, and personal care. Each domain is rated as follows: 0 (none), 0.5 (questionable), 1 (mild), 2 (moderate) and 3 (severe). A CDR of 0.5 or greater might be associated with AD, depending on the complementary neuropsychological tests that are performed (Manning and Ducharme, 2010). However, these tests lack specificity in diagnosing a patient in the preclinical stages of AD (Chapman et al., 2010; Khan, 2016), as they are not able to determine pathology until clinical symptoms manifest. This limitation can be seen in Figure 2.3, where the changes in memory and clinical function arise in the MCI and dementia stages, respectively.

2.2.2 Lumbar Punctures Test

Lumbar punctures can be used to sample the CSF from the spinal canal to provide valuable insight into the presence of amyloid-beta. In subjects with AD, a decrease of amyloid-beta is expected due to its accumulation in the brain (Khan and Alkon, 2015). Notwithstanding its overall good specificity to detect preclinical AD compared with neuropsychological assessments and low base cost (around $127.45), lumbar punctures are a highly invasive and painful test. They can also have a small risk of provoking undesired effects such as spinal headache, infection, and nerve injury (Hudgins et al., 2017). Even if the risk is low, this technique can induce high levels of anxiety for the participant. Studies have shown that these tests should be coupled with other CSF biomarkers (i.e., tau and amyloid-beta-42 to 40 ratio)(Hampel et al., 2021; Ritchie et al., 2017) in order to increase its specificity for preclinical AD detection. Since lumbar punctures cannot provide localized information regarding amyloid-beta accumulation within the brain (Hansson et al., 2019), it is also recommended to add a molecular neuroimaging test (discussed in the next section) to properly diagnose a subject (Reimand et al., 2020). This large combination of tests not only increases the cost but the risk and negative effects for the participant.

2.2.3   Neuroimaging Tests

Neuroimaging technologies, particularly Positron Emission Tomography (PET) imaging and structural Magnetic Resonance Imaging (MRI), provide a window into the pathophysiology of the brain by examining the deposition of molecules directly in the brain tissue and regional structural atrophy, respectively (Johnson et al., 2012).

Positron Emission Tomography

PET is a neuroimaging technique that uses radiotracers to measure molecular changes within the body. Radiotracers are chemical compounds in which one or more atoms have been replaced by a radioisotope. Once introduced into the body, the radiotracer travels through the bloodstream and binds to specific molecules. The radiotracer emits positrons, and when a positron encounters an electron, an annihilation event occurs, producing a pair of gamma photons that travel in opposite directions. These photons reach opposite detectors of the PET scanner at almost the same time as shown in Figure 2.4.

The scanner calculates the line of response, which is the line connecting the two detection points, to determine the probable origin of the annihilation event. This process often involves the use of probabilistic functions to account for uncertainties. The collected data from all such events are then used to reconstruct a functional image of the tracer distribution in the body. This reconstruction can be performed using various algorithms, such as back projection or iterative reconstruction, to enhance image clarity and resolution (Martínez et al., 2008).

For amyloid-beta, the two radiotracers approved for human use are florbetapir and 11C-PiB. These radiotracers bind to amyloid-beta molecules within the brain to show regional concentrations associated with amyloid-beta deposition (Camus et al., 2012; Yamin and Teplow, 2017). Producing the radiotracers for amyloid-beta is not a simple task, as it requires the use of a cyclotron, a particle accelerator that can produce radioisotopes by bombarding a target material with charged particles. The cost of obtaining and operating a cyclotron is high, ranging from \$2.5 million to \$6.6 million (in addition to the cost of obtaining a PET scan) (Drug and Agency, 2022), thereby limiting the availability of PET imaging in many jurisdictions.

Cost is not the only limitation to performing PET imaging. The use of radiotracers exposes the patients to harmful ionizing radiation from the radioisotope, limiting the number of screenings that a patient can safely undergo (Marti-Climent et al., 2017; Valentin, 2005). Due to radiation safety protocols, a patient cannot be imaged with more than one tracer per day. Therefore, if both amyloid-beta and tau need to be measured, the process is very time-consuming. Furthermore, PET scans typically need to be complemented with a computer tomography or an MRI scan to provide detailed structural information (Ehman et al., 2017). Such considerations make PET imaging difficult for large cohort studies and impractical for screening early AD detection. There is almost no use of PET imaging for AD screening in Canada at this time.

Positron Emission Tomography
Scanner Example
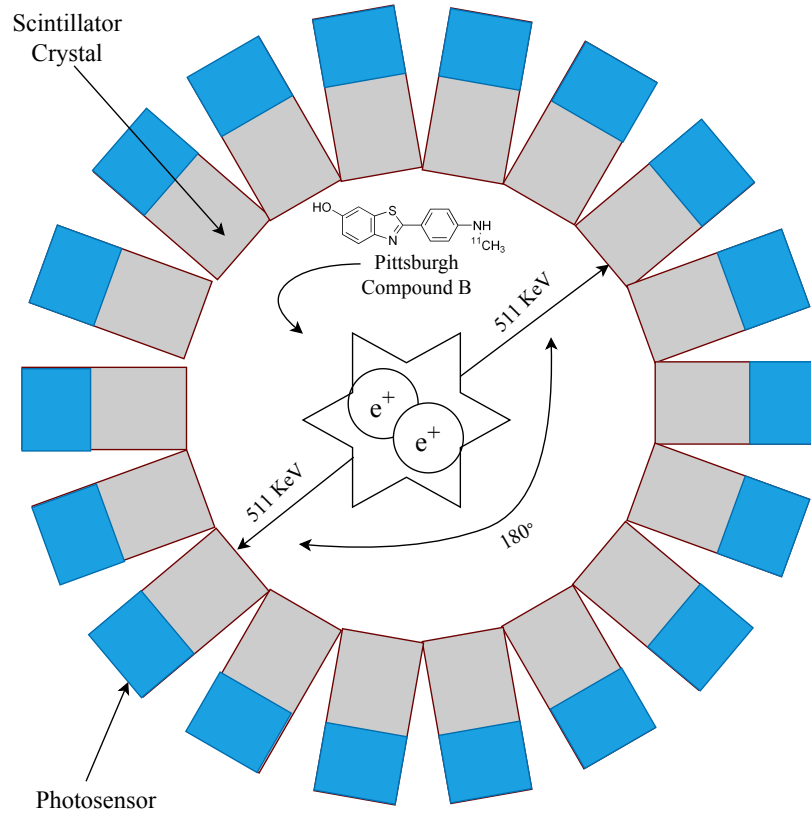
Figure 2.4: Positron Emission Tomography (PET) scanner example where the positron encounters an electron, then annihilates producing a pair of gamma photons that travel opposite directions. These photons hit the detector at almost same time.

Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) is a medical imaging technique that utilizes strong magnetic fields (typically 1.5 or 3 Tesla [T]) and radio frequency (RF) signals to capture detailed images of the body's internal structures and functions. In MRI, hydrogen protons within the body are affected by the magnetic field, denoted as B_0, aligning the protons either parallel or antiparallel to the field to create a net magnetization (M_o) in the direction of B_0 (Z plane). Under the influence of B_0, the protons precess around the axis of the magnetic field at a frequency known as the Larmor frequency (Liang and Lauterbur, 1999), which is proportional to the strength of the magnetic field and can be determined with the following equation:

$$\omega = \gamma * B_0 \tag{2.1}$$

Where omega is the Larmor frequency, gamma is the gyromagnetic ratio, which is 52.75MHz/T for hydrogen, and

B_0 the magnetic field strength.

In order to obtain an image, the protons must be first perturbed with a radiofrequency (RF) pulse at the Larmor frequency. This will induce energy to the protons, causing them to rotate toward the XY plane, thereby changing the direction of M_0. When the RF pulse is switched off, the protons release energy as they realign with the B_0 field. This process, known as relaxation, occurs both longitudinally (along the direction of B_0) and transversely (perpendicular to B_0).

The rate at which M_0 recovers in the Z plane is known as longitudinal relaxation (T1). This relaxation is governed by the following Bloch equation:

$$M_z(t) = M_0 \left(1 - e^{-\frac{t}{T1}} \right) \tag{2.2}$$

Here, $M_z(t)$ represents the longitudinal magnetization at time t, M_0 is the net magnetization, and $T1$ is the longitudinal relaxation time. Equation 2.2 describes how magnetization returns to its equilibrium state along the B_0 direction. The rate of this process varies among different tissues and is the property related to T1w imaging. T1w imaging is commonly used to determine the difference between tissues, particularly between gray and white matter (Plewes and Kucharczyk, 2012).

A second mechanism occurs in MRI that causes a decrease in the transversal magnetization. The spinning protons interact with each other, causing them to dephase in the XY plane. This dephasing leads to an exponential decay of the signal with a time constant called T2. This process, distinct from the influences of external field inhomogeneities (which characterize another type of T2 signal called T2*), generates a signal known as Free Induction Decay (FID) that is captured by the RF coil. The T2 relaxation behavior is depicted as the following equation:

$$M_{xy}(t) = M_0 e^{-\frac{t}{T2}} \tag{2.3}$$

Since T1 and T2 happen at the same time, they can be depicted as Figure 2.5, where the FID is produced during this relaxation process.

Therefore, it is possible to obtain T1w contrast or T2w contrast, and both show different information. For example, in T1w images, the CSF often shows dark while in T2w it shows bright. There are additional mechanisms that affect the T2 signal, however, these will not be explored in this thesis.

Since, this FID signal contains all the anatomy information that is obtained by the RF coil, therefore, additionally steps must be performed to spatially encode the signal. This can be obtained by using a set of three orthogonal gradient coils (x-, y-, and z-gradient). Traditionally, MRI images are collected slice by slice, and to control which slice is being targeted, the z-gradient is used. x- and y- gradients are used to acquire the images in the respective plane.

To record the FID signal, an x-gradient is applied. This alters the Larmor frequency of the protons along the x-
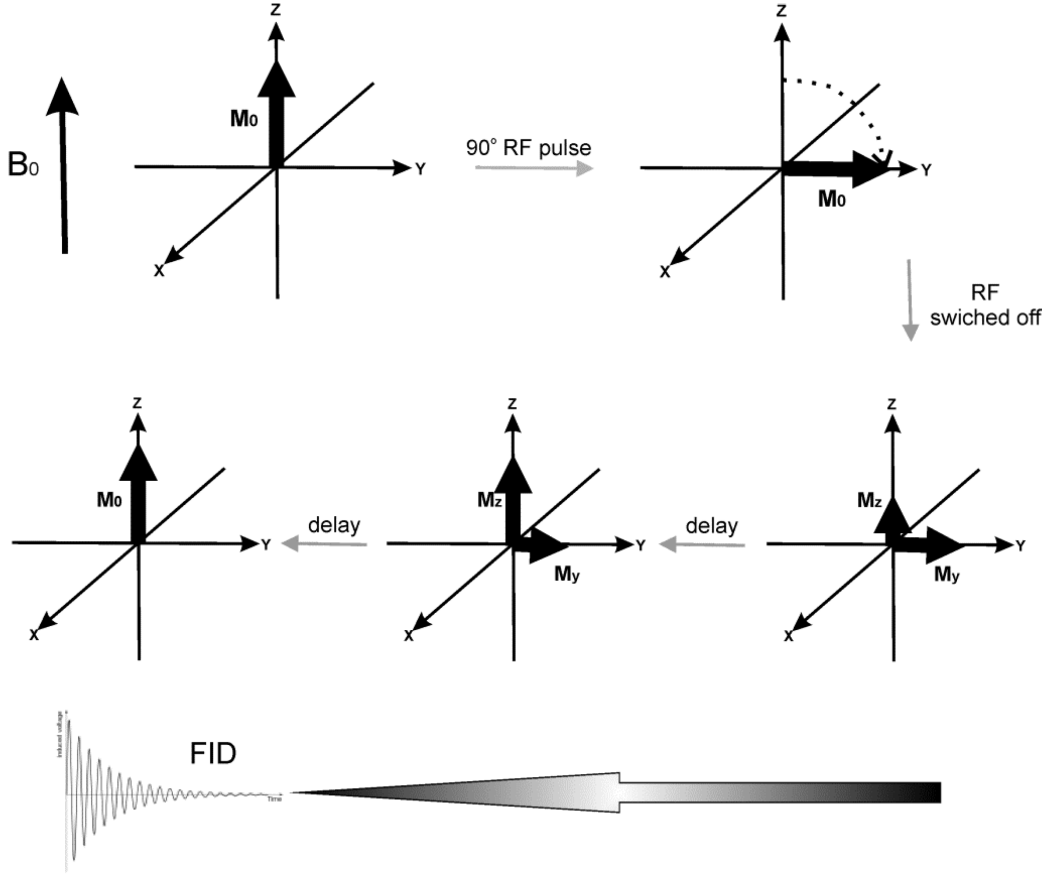
Figure 2.5: Steps leading to the free induction decay (FID) used in MRI. B_0 represents a big magnetic field (i.e. 3T) that affects a hydrogen particle represented as M_0 flag, then a radiofrequency pulse (RF) is applied to M_0 and flips it against B_0. The RF pulse is switched off and it relaxes. Figure retrieved from (Dulińska-Litewka et al., 2019)

direction, encoding the recorded signal as a function of x. This process is called frequency encoding and is described by the following equation.

$$\omega(x) = \gamma B_0 + \gamma G_x x \tag{2.4}$$

Where $\omega(x)$ is the precessional frequency of the hydrogen protons at a particular location in the x-axis in the slice, G_x is the strength of the x-gradient known as frequency encoding gradient, and x is the position of the hydrogen protons along the x-axis.

However, using only frequency encoding is not enough to create an image with contrast between tissues. This is because similar tissues can have similar Larmor frequency, and therefore, no difference will be captured between them. In order to capture the difference, it is necessary to de-phase the signal from those similar tissues, which can be achieved by altering the phase using the y-gradient in a process known as phase-encoding (Liang and Lauterbur, 1999) which is described using the following equation.

$$\phi(y) = -\gamma G_y y T_{pe} \tag{2.5}$$

Where $\phi(y)$ is the phase shift induced in the hydrogen signal, G_y is the strength of the y-gradient known as phase encoding gradient, y is the position of the signal along the y-axis, and T_{pe} is the duration of the phase encoding gradient. This process is repeated a certain amount of times with different amounts of phase encoding to build a two-dimensional image. However, this process entails a long acquisition time, and more sophisticated techniques are now used to accelerate this process.

To generate a complete MRI image, both frequency and phase encoding techniques are integrated to encode spatial information, which, for the purpose of this explanation is in two dimensions. As discussed, while frequency encoding varies the Larmor frequency along the x-axis, phase encoding creates phase shifts along the y-axis. The combination of these methods allows for the precise localization of signal origins within the imaging plane. This integration of both components can be described as spatial encoding, which is captured as a k-space of the MRI signal. The k-space is a function of the spatial frequencies along the x (kx) and y (ky) axes, providing an array that can be transformed back into a spatial image. This process can be described by the following spatial encoding equation:

$$S(k_x, k_y) = \int \int \rho(x,y) e^{-i2\pi(k_x x + k_y y)}, dx, dy \tag{2.6}$$

Where, $S(k_x, k_y)$ is the signal in k-space, $\rho(x,y)$ represents the proton density at point (x,y) and k_x, k_y are spatial frequencies corresponding to the frequency and phase encoding gradients.

However, spatial encoding is just one component that is required to acquire a complete MRI image. Extra considerations such as timing should be taken into account to achieve a desired MRI contrast. The equation that determines how the signal is recorded after the first RF pulse can be obtained from the T1 and T2 equations:

$$M_{xy}(t) = M_0(1 - e^{-\frac{t}{TR}})(e^{-\frac{t}{TE}}) \tag{2.7}$$

Where each subsequent RF pulse has two timing variables: the repetition time (TR) that determines when another RF pulse be repeated, and the echo time (TE) the time at which the FID signal is collected. By changing TR and TE it is possible to modulate the contribution of T1 and T2 relaxation in the FID signal as shown in Figure 2.6.

Traditionally, MRI images were acquired in a slow iterative process that involved selecting a single slice and performing the spatial encoding. To accelerate this process, a technique was created specially for T1w contrast imaging. Magnetization Prepared Rapid Gradient Echo (MP-RAGE) is an optimized sequence that is commonly used to quickly obtain high-resolution T1w images. MP-RAGE enhances the contrast differences between gray and white matter signals, making it attractive for evaluating signal differences in gray and white matter, changes in texture in structures

(a) T1 Relaxation Signals   (b) T2 Relaxation Signals
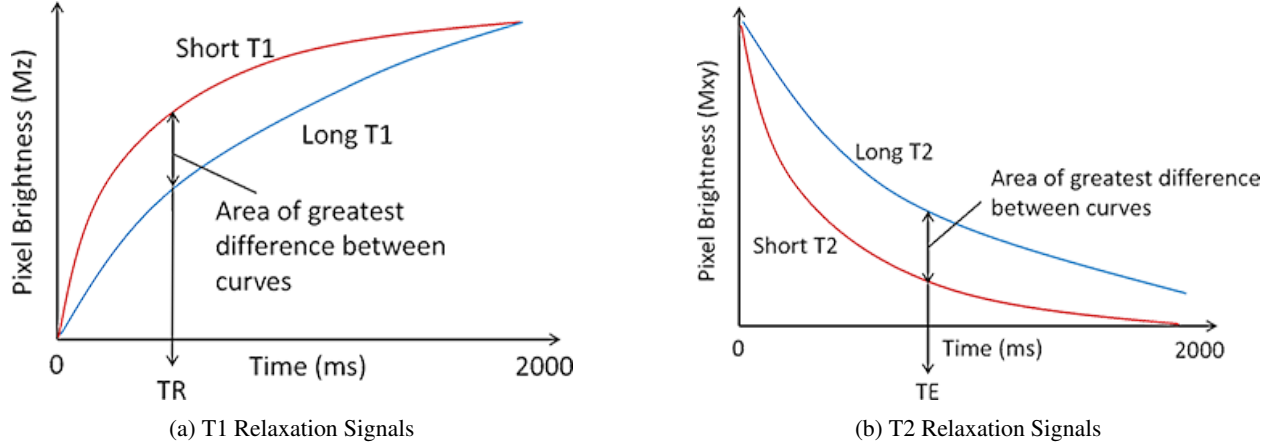
Figure 2.6: T1 longitudinal (left) and T2 transversal (right) relaxation signals showing two tissues that have long and short times, repetition time (TR) controls how much of the T1 signal is being obtained while echo time (TE) controls the T2 signal. Figure retrieved from (Abdulla, 2021)

such as the hippocampus, and cortical brain atrophy (Wang et al., 2014), this sequence is important to assess AD progression as the mentioned differences in signal, changes in texture and morphology are known to be related to this disease(Vemuri and Jack, 2010; Lee et al., 2020). Recent work has shown that such changes in signal intensity in gray and white matter can be used as a biomarker for neurodegeneration **??** where MP-RAGE sequence has proofed to be the most suitable one to measure them. Briefly, MP-RAGE employs a technique called an inversion recovery pulse that inverts the longitudinal magnetization to improve the image quality and contrast between tissues and uses RF pulses with short TR, short TE, and small flip angle to accelerate the image acquisition and preserve the difference in contrast between tissues in the longitudinal plane (Bydder et al., 1998).

For this thesis, T1w images from the OASIS-3 dataset were selected for their ability to provide information about the changes in brain structures and differences in signal intensities in gray and white matter that are associated with amyloid-beta burden and AD progression.

2.3 Deep Learning and Image Translation Models

This thesis makes extensive use of deep learning models to perform image translation. To understand their functionality, the concept of a convolutional neural network (CNN), U-Net, and Generative Adversarial Network (GAN) must be described.

2.3.1 Convolutional Neural Networks

In computer vision, a convolution is a mathematical operation that filters an image using an n×n matrix, referred as a kernel (commonly of the 3×3 size), which traverses (i.e., slides across) the image. As this kernel moves, it can filter

and extract specific features from the image, such as edges, corners, and contours (Albawi et al., 2017).

CNNs utilize a series of kernels to recognize and capture various patterns within an image. The patterns detected by a CNN are often indiscernible the the human eye. These networks are composed of multiple layers that perform convolutional operations, each designed to process different aspects of the input data. The common layers types in CNNs are found as follow:

- **Convolutional layers:** The primary building blocks of CNNs. These layters use filters to perform convolution operations, capturing local features such as edges and textures. The output of this layers is a feature map that represents the presence of specific features in the input image.

- **Activation layers:** Typically following each convolution layer, activaion layers introduce non-linearities into the model, allowing the network to learn complex patterns. The Rectified Linear Unit (ReLU) is a common activation function used in CNNs

- **Pooling layers:** These layers reduce the spatial dimensions of the input image for the next convolutional layer. They help in reducing the computational load, memory usage, and also help in detecting important features that are present regardless of scale changes. For this layer, some CNNs use max pooling or strided convolution.

- **Upsampling layers**: These laters increase the spatial dimensions of the input feature maps. They are the opposite of pooling layers, which reduce the spatial dimensions. Upsampling is commonly used in tasks where fine, detailed information is necessary such as segmentation, super-resoltion, and image generation

- **Normalization layers:** These layers are used to standardize the inputs of the network or its layers. By normalizing the features, these layers help in speeding up the training process and improving the overall performance of the model. They achieve this by reducing the changes in the distribution of the networks activation due to the change in networks parameters during training

- **Dropout layers:** Layers that regularize the network, preventing overfitting during the training process. These layers randomly "shut down" a subset of neurons during training. This means that each neuron, along with its connections, is temporarily removed from the networks with a certain probability. This introduces variability to the network which is often desired in specific applications such as image generation

- **Fully connected layers:** Found at the end of the network, these layers use the features extracted by previous laters to determine the final output. In classification tasks, these layers are used to determine the class of the input image

CNNs can have other types of layers that are often used for specific cases, however, the listed ones are the commonly found in many applications and architectures (Yamashita et al., 2018; Alzubaidi et al., 2021). Figure 2.7. shows a basic

CNN architecture that can be used for classification, such architecture uses some of the layers that were previously described.
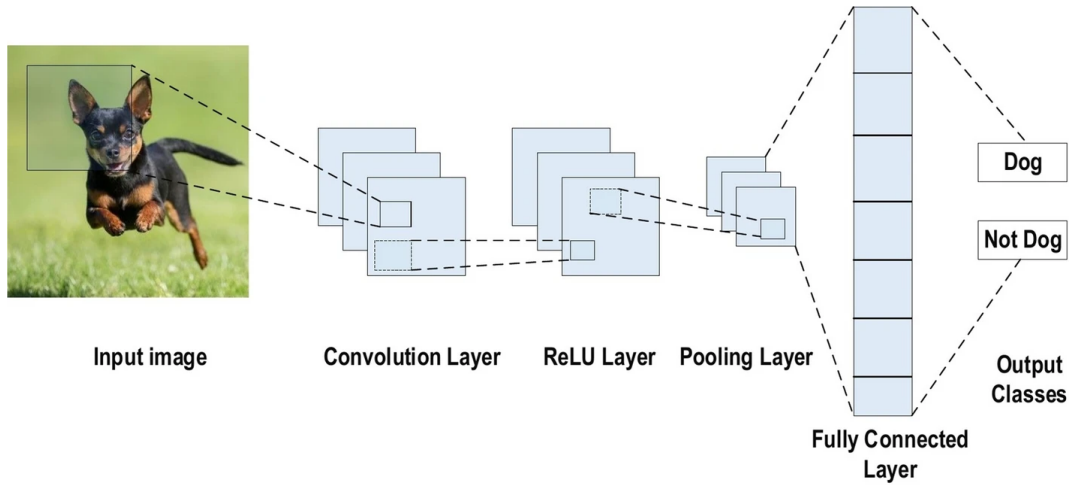
Figure 2.7: Classic CNN architecutre used for classification, it recieves an input image and applies convolutional, activation, pooling, and fully connected layers to reach a decision. Figure retrieved from (Alzubaidi et al., 2021)

2.3.2 U-Net

For image translation, a widely used CNN architecture called the U-Net is often employed (Ronneberger et al., 2015). The U-Net processes an input image through a series of layers that extract features of the image until reaching a bottleneck as an encoder. It also has skipped connections that avoids losing information due to downsampling and upsampling operations, such skipped connections help to bring features that might have been overlocked in the latent space as shown in Figure 2.8 (Ronneberger et al., 2015). The U-Net training process tries to minimize the error between the generated image (synthetic) and the real image using either L_1 or L_2 distance norms that are known as mean absolute error (MAE) and mean squared error, respectively. The U-Net uses the error to adjust itself during training (Goodfellow et al., 2016).

This iterative adjustment ensures that the network emphasizes features in the latent space that bring the generated image closer to the target. While the performance of U-Nets is acceptable, they can generate obvious artifacts and/or blurriness in the synthetic images. Therefore, this thesis makes use of a more sophisticated and powerful approach.

2.3.3 Generative Adversarial Networks

Generative Adversarial Networks (GAN) (Goodfellow et al., 2014) are models capable of generating images that are nearly indistinguishable from the real ones (target image). GAN models have variants such as Conditional Generative Adversarial Networks (cGAN) (Isola et al., 2017) and Cycle-Consistent Generative Adversarial Networks (Cycle-

Classical U-Net Architecture
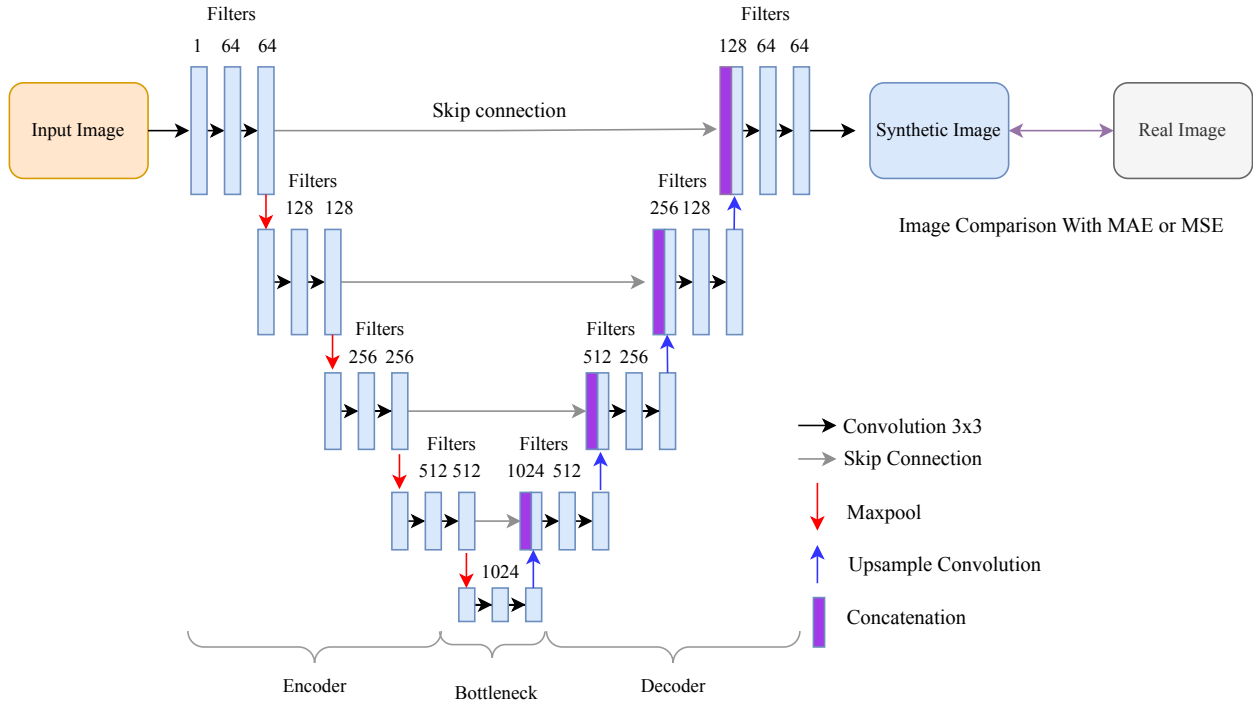
Figure 2.8: Illustration of example U-Net architecture. The U-Net receives an input image where it extracts features while applying convolutional operations and extracts relevant features until reaching a bottleneck as an encoder, then creates a new image using the most relevant features as a decoder, the number of extracted features for each convolutional block is on top of each. The synthetic image is compared with the real image using mean absolute error (MAE) or mean squared error and the network improves. Figure adapted from (Ronneberger et al., 2015)

GAN)(Zhu et al., 2017) that have been extensively used for medical image translation tasks.

In general, GAN, cGAN, and Cycle-GAN models are comprised of multiple CNNs that are trained against each other, namely generator and discriminator networks. The generator can follow a U-Net architecture as described above. It creates synthetic images that minimize the error between real and synthetic image pairs (Isola et al., 2017; Goodfellow et al., 2016). The discriminator follows an encoder architecture that works as an image classification network. The discriminator receives the synthetic images (created by the generator) and the target images. The discriminator's purpose is to learn to classify the synthetic images as real or synthetic by outputting a numerical value between 0 and 1, where 0 means the images are classified as real and 1 means the images are classified as synthetic. The discriminator tries to maximize the probability of correctly classifying the images (Isola et al., 2017; Zhu et al., 2017; Goodfellow et al., 2014). Upon receiving the decision from the discriminator, the generator refines the image synthesis. This iterative process continues until both networks reach an equilibrium.

Conditional Generative Adversarial Network Loss and Objective Function

As described in the GAN section, the cGAN training process is comprised by the generator and discriminator being trained simultaneously against each other. The traditional objective function of the cGAN can be expressed as:

$$L_{cGAN}(G,D) = E_{x,y}[\log D(x,y)] + E_{x,z}[\log(1 - D(x,G(x,z)))] \tag{2.8}$$

Where y is the label image, x is the input, G(x,z) is the synthetic image produced by the generator (G), and z represents model variability that can be induced with a noise vector or latent space variability with dropout layers. $D(x,y)$ is the discriminator's (D) decision after being fed with a true pair (input and label), and $D(x,G(x,z))$ is the discriminator's decision after being fed with a synthetic pair of images (input and synthetic). This function describes the training process of the cGAN where G tries to minimize against an adversarial D that tries to maximize.

Since this traditional objective function only enforces the generator to fool the discriminator, it is necessary to task the generator to create synthetic images that are as close as possible to the label images using L1 distance. This equation can be described as:

$$L_1(G) = E_{x,y,z} \|y - G(x,z)\|_1 \tag{2.9}$$

The final objective function can expressed as:

$$cGAN = \arg \min_G \max_D L_{cGAN}(G,D) + L_1(G) \tag{2.10}$$

Since Equation 2.10 summarized the cGAN training, its workflow is depicted in Figure 2.9

Conditional Generative Adversarial Network Workflow
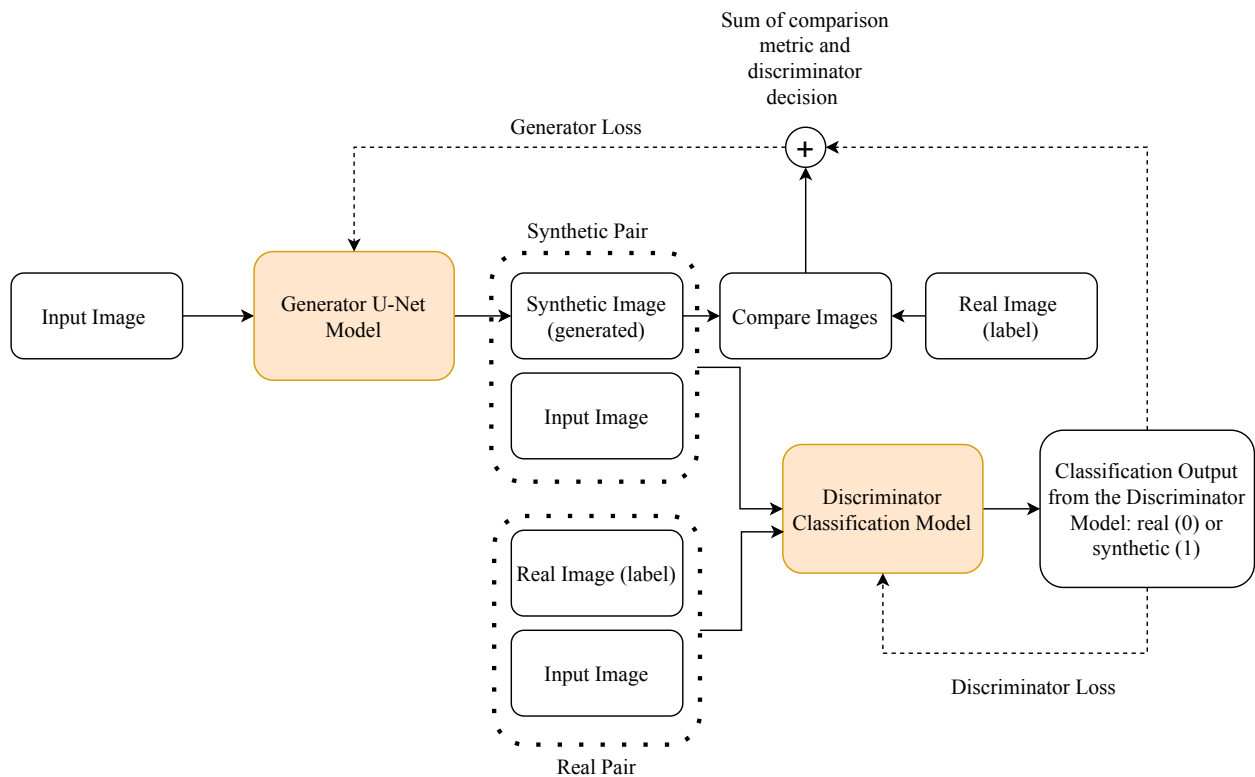
Figure 2.9: Conditional Generative Adversarial Network (cGAN) workflow, showing that the generator receives an input image and generates a synthetic one. The synthetic image pair of input and synthetic (synthetic image), and the real pair of input and label (real image) train the discriminator that learns to classify the synthetic pair as real or synthetic. The discriminator loss is added to the error between synthetic and true images and updates the generator in an iterative process. The dotted line represents the feedback to the networks.

2.4 Example Work Relating T1w with Amyloid-Beta Positivity

Recent findings have shown that there is a relationship between the information that T1w MRI provides with amyloid-beta positivity. One of those works demonstrate that a machine learning model trained with radiomics features from structural MRI images (histogram and texture features) can classify whether a MCI subject has amyloid-beta positivity with an accuracy of at least 71% in the testing cohort (Kim et al., 2021). Another work in the same are which closely relates to this thesis focuses on building a three-dimensional CNN that can determine amyloid-beta positivity from T1w MRI images with gray, white matter, and CSF segmentations. The model was able to obtain accuracy of 76% demonstrating that determining amyloid-beta positivity from T1w MRI is feasible (Chattopadhyay et al., 2023). These works show that T1w MRI images contain valuable information such as changes in brain morphology (atrophy), signal intensities in gray and white matter, and texture which can potentially be used to infer amyloid-beta PET images from single T1w MRI ones. The follow section will describe relevant work that attemted to translate MRI images to PET images.

2.5 Example Work using Image Translation between MRI to PET

GANs holds the potential to produce synthetic images that are nearly identical to the real ones. They have been extensively applied in past research to generate synthetic PET images from MRI using other radiotracer types such as florbetapir or 18F-fluodeoxyglucose (FDG).

2.5.1 Three-Dimensional Self-Attention Conditional GAN with Spectral Normalization for Multimodal Neuroimaging Synthesis

In the work "Three-Dimensional Self-Attention Conditional GAN with Spectral Normalization for Multimodal Neuroimaging Synthesis" (Lan et al., 2021), the authors presented a three-dimensional self-attention cGAN with spectral normalization for multimodal neuroimaging synthesis (SC-GAN) including amyloid-beta florbetapir PET. The SC-GAN receives two MRI images as input (T1w and Flair) and can generate three potential images: amyloid-beta florbetapir PET, mean diffusivity MRI, and fractional anisotropy MRI. The authors utilized the Alzheimer's Disease Neuroimaging Initiative 3 (ADNI-3) database, which provided the 5 image types. The SC-GAN, distinct from a typical cGAN, incorporates advanced techniques such as self-attention, spectral normalization, feature matching loss, brain area root means square error (RMSE) loss, and regularization techniques. These additions can generate high-quality synthetic MRI images. However, the synthesis of amyloid-beta florbetapir PET images required further improvements to be more accurate. For the purpose of this thesis, the focus is on the translation from MRI to PET methods and results.

The authors implemented a cGAN with the following components: Self-attention: This technique emphasizes crucial features during the image feature extraction process with convolutions. Losses: This study employed various losses, including Feature matching loss which compares features from the real and synthetic images obtained from the Discriminator. Brain area RMSE loss masks the brain from the background and compares the real and synthetic images using RMSE. Regularization in the generator and discriminator with mean square error to prevent overfitting. Spectral normalization: Technique that helps to stabilize the cGAN training by adjusting the internal settings of the model to prevent extreme variations. The model was trained with 207 PET images and the images underwent normalization and scaling to a range of 0 to 1.

The authors reached a mean Peak Signal-to-Noise Ratio (PSNR) of 32.14 and a mean Structural Similarity Index Measure (SSIM) of 0.962 for the PET image synthesis, showing that the model was able to generate high-quality synthetic PET images. However, while the paper exemplifies the use of image translation algorithms to generate synthetic PET images, it has several limitations, such as:

- Using min-max scaling in the PET images ensures a stable training process, however, such technique does not allow to revert the scaling. This means that the generated PET images were purely qualitative instead of quantiative yielding data that only ranges between 0 and 1 insted of the original range

- The feature matching loss can be helpful to improve the image synthesis; however, it is known to be prone to overfitting the network as it enforces to approach closer to the data distribution from the real PET images instead of generalizing (Salimans et al., 2016)

- The principal limitation is that the model was trained with a small training sample. Usually in GANs it is desirable to train with a big sample size (more than 400 images), otherwise the model would tend to overfit. This overfiting problem can be seen in the proposed work where the authors only trained for 25 epochs obtaining high similarity metrics, which is hard to obtain in GANsn

- The authors stated that their study was not able to produce accurate amyloid-beta synthetic PET images, suggesting that the radiotracer estimation needs to be improved

One strong point from this project is the technique used for the brain RMSE area which focused the model towards the brain while ignoring the background. As a conclusion, this study demonstrated the potential of PET synthesis from structural MRI. However, methodological enhancements are necessary.

2.5.2 Predicting FDG-PET Images From Multi-Contrast MRI Using Deep Learning in Patients With Brain Neoplasms

In the work, "Predicting FDG-PET Images From Multi-Contrast MRI Using Deep Learning in Patients With Brain Neoplasms" (Ouyang et al., 2023), the authors present a two-dimensional cGAN with self-attention and vision transformer module to predict FDG PET imaging from multi-contrast MRI inputs. The network receives four MRI image contrasts as input (T1, T1 post-contrast, flair, and arterial spin labeling) and generates a synthetic FDG PET image. All data were gathered by the authors. The addition of novel techniques such as self-attention and a vision transformer block helps to produce high-quality PET images.

The authors introduced a variation of cGAN, similar to previously discussed work, where they implemented the following components: Self-attention: As discussed earlier, self-attention can help to learn crucial relationships between features of an input image during the convolution operation for this work, the authors implemented it in the skipped connections of the generator that followed a U-Net configuration. Vision transformer: For this approach, the latent space was used to compare the feature maps generated during the encoder part of the U-Net, then implemented multi-head attention which is an improved version of self-attention that compared pixel-wise features between each other to reveal the most relevant ones. No specifications about the model training were provided. The authors performed five-fold cross-validation to assert the model. The authors also performed clinical blind readings where three clinicians received a set of real and synthetic images and were tasked to assess and classify them. The authors used paired t-tests to evaluate the significant differences between true and synthetic PET evaluations by the clinicians.

As a result, the model was able to score a mean PSNR of 27.84 and SSIM of 0.88, and the paired t-test reported no significant differences between true and synthetic PET cases with a p-value of 0.14. While the paper is another example of image translation to estimate PET from MRI, this work shows evident limitations, including:

- There is no information about the normalization steps for the MRI and PET images, making this study hard to replicate. Likely, they used normalization making the PET non-quantitative

- The authors made use of multi-contrast MRI input images. While this approach might help to produce high-quality PET images, performing four MRI scans on a patient could significantly elevate the cost of this system

- The use of self-attention and vision transformers has shown to be promising for image translation. However, implementing such a computationally heavy technique can make this mode unfeasible in the three-dimensional space

- Lastly, the authors did not show a reconstructed three-dimensional volume showing that the two-dimensional synthetic PET images have continuity across all the dimensions

While this study has some limitations, it reaffirms the value of producing synthetic PET images from MRI.

2.5.3 Chapter Conclusion

This chapter defined key concepts regarding AD, amyloid-beta, techniques for AD detection, PET, MRI, and image translation algorithms. Two articles that show the potential of generating PET imaging from structural MRI are also reviewed to indicate the state of the art, which generally fails in accurate quantification methods. The subsequent three chapters present the experimental portions of this work.

Chapter 3

Image Translation for Estimating Two-Dimensional Axial Amyloid-Beta PET from Structural MRI

3.1   Abstract

Background: Amyloid-beta and brain atrophy are hallmarks of Alzheimer's Disease that can be targeted with Positron Emission Tomography (PET) and MRI, respectively. MRI is cheaper, less invasive, and more available than PET. There is a known relationship between amyloid-beta and brain atrophy, meaning PET images could be inferred from MRI.

Purpose: To build an image translation model using a Conditional Generative Adversarial Network able to synthesize Amyloid-beta PET images from structural MRI

Study Type: Retrospective

Population: Eight Hundred Eighty-Two adults (348-Males/534-Females) with different stages of cognitive decline (control, mild cognitive impairment, moderate cognitive impairment, and severe cognitive impairment). 552 subjects for model training and 331 for testing (80%:20%).

Field Strength/Sequence: 3T, T1-weighted structural (T1w)

Assessment: The testing cohort was used to evaluate the performance of the model using the Structural Similarity

Index Measure (SSIM) and Peak Signal-to-Noise Ratio (PSNR), comparing the likeness of the overall synthetic PET images created from structural MRI with the overall true PET images. SSIM was computed in the overall image to include the luminance, contrast, and structural similarity components. Experienced observers reviewed the images for quality and performance and tried to determine if they could tell the difference between real and synthetic images.

Statistical Tests: Pixel-wise Pearson correlation was significant, and had an R2 greater than 0.96 in example images. From blinded readings, a Pearson Chi-squared test showed that there was no significant difference between the real and synthetic images by the observers (p=0.68)

Results: A high degree of likeness across the evaluation set, which had a mean $SSIM = 0.905$ and $PSNR = 2.685$. The two observers were not able to determine the difference between the real and synthetic images, with accuracies of 54% and 46%, respectively.

Conclusion: Amyloid-beta PET images can be synthesized from structural MRI with a high degree of similarity to the real PET images.

3.2 Keywords

Positron Emission Tomography, Amyloid-beta, Magnetic Resonance Imaging, Image Translation, Generative Adversarial Network

3.3 Introduction

Alzheimer's Disease (AD) is the most common type of dementia with increasing prevalence as a result of aging populations (Wimo et al., 2023). Clinically, AD is characterized by a deterioration of the cognitive processes of the patient, starting with a mild cognitive impairment (MCI) and culminating in dementia (Gaser et al., 2013). Dementia had an enormous global economic burden of $1.3 trillion US in 2019 (Wong, 2020), and the cost is anticipated to continue increasing. Although neuropsychological testing plays an important role in the assessment of dementia within the clinical setting, medical imaging is also important for the assessment of neurodegeneration in dementia and has the potential to detect early stages of the disease prior to the manifestation of clinical symptoms (Johnson et al., 2012).

Amyloid-beta is a monomer that aggregates in the brain resulting in pathological aging and is often considered to be a leading indicator of AD (Chetelat et al., 2010). The synthesis of amyloid precursor protein with hydrophobic fragments of amyloid-beta aggregate and cannot be cleared from the brain. They consequently accumulate between the neurons and lead to synaptic failure, triggering an immune response and death of neurons (Hampel et al., 2021; Chen et al., 2017). The amyloid-beta aggregates were first observed as 'plaques' identified in post-mortem AD patients

more than a hundred years ago in seminal work by Alois Alzheimer between 1906 and 1911. Amyloid-beta has been extensively studied with respect to AD since then (Mantzavinos and Alexiou, 2017; Meyer et al., 2019).

Over the past two decades, several AD screening techniques have emerged, such as primarily neuropsychological assessments, lumbar punctures for cerebrospinal fluid (CSF), and neuroimaging. Neuropsychological assessments aim to evaluate the memory, language impairment, and personality (Chapman et al., 2010); however, these tests lack of specificity to diagnose a patient in the preclinical stage (Chapman et al., 2010; Khan, 2016). A common neuropsychological test for AD is the Clinical Dementia Rating (CDR) score, which evaluates the subject's cognitive performance in six fields: memory, orientation, judgment and problem solving, community affairs, home and hobbies, and personal care. Each domain is rated as follows: 0 (none), 0.5 (questionable), 1 (mild), 2 (moderate) and 3 (severe). A CDR of 0.5 or greater might be associated with AD depending on the complementary neuropsychological tests that are performed (Manning and Ducharme, 2010). Lumbar punctures sample the CSF and evaluate the amyloid-beta concentration in patients (Kumar et al., 2022). Notwithstanding its high specificity to detect AD compared with neuropsychological assessments, lumbar punctures are an extremely invasive test (Khan and Alkon, 2015) and may not give a good indication of the total amyloid-beta deposited in the brain tissue (Hansson et al., 2019). Neuroimaging technologies, particularly positron emission tomography (PET) imaging and structural MRI, provide a window into the pathophysiology of the brain by examining the deposition of molecules directly in brain tissue and regional structural atrophy, respectively (Johnson et al., 2012).

PET imaging can measure molecular concentrations within the body using radiotracers that bind with specific molecules (Forsberg et al., 2008) and target amyloid-beta with radiotracers such as 18F-AV-45 (florbetapir) (Camus et al., 2012) and Pittsburgh Compound B (11C-PiB) (Yamin and Teplow, 2017). However, PET imaging has disadvantages, including: cost (approximately $5000 to $8000 per PET scan) (Wittenberg et al., 2019), invasiveness (requiring injection of a radiotracer, exposure to harmful ionizing radiation, and blood sampling) (Marti-Climent et al., 2017; Valentin, 2005), and limited availability (the technology is not available in many jurisdictions). These limitations make PET infeasible for population screening and difficult to incorporate in large cohort studies.

MRI offers hundreds of contrast modes that can aid in detecting the structural changes associated with brain atrophy in AD (Buxton, 2009). MRI has the advantage of being more affordable, accessible, and less invasive than PET, but it cannot measure molecular information (i.e. amyloid-beta concentration). Previous studies have found mutual information relating amyloid-beta PET and structural MRI (Tosun et al., 2011). This means that amyloid-beta PET images could potentially be inferred from the structural MRI images, thereby reducing the cost per scan and minimizing invasiveness (Tosun et al., 2011).

Several recent studies have focused on image translation algorithms that generate synthetic images from another image type (Pang et al., 2021; Isola et al., 2017). Image translation models are a powerful technology that can render images that would not have been captured otherwise, or produce higher cost images from lower cost inputs. For image

translation to be successful, there needs to be sufficient mutual information between the image types (Isola et al., 2017). Image translation algorithms are commonly implemented with Convolutional Neural Networks (CNNs) in a U-Net configuration (Yamashita et al., 2018). U-Nets do not always have enough generative power to produce high quality synthetic images alone, so many recent applications have used Generative Adversarial Networks (GANs) (Goodfellow et al., 2014; Sikka et al., 2021; Armanious et al., 2019) and its variants: Conditional Generative Adversarial Networks (cGANs) (Mirza and Osindero, 2014) and Cycle Consistent Generative Adversarial Networks (Cycle-GAN) (Armanious et al., 2019). These models have achieved the most realistic images, making them indistinguishable from real ones. Variants of GANs have shown promising efficacy in the field of medical imaging, applying them to perform tasks such as denoising (Ran et al., 2019), segmentation (Güven and Talu, 2023), reconstruction (Lv et al., 2021; Yang et al., 2018) and MR motion correction (Armanious et al., 2018; Kustner et al., 2019; Johnson and Drangova, 2019). In general, for medical image translation applications, GANs, cGANs and Cycle-GANs are comprised of multiple CNNs, including both a generator and discriminator that are trained against one another. The generator creates synthetic images that minimize the error between real and synthetic image pairs, and the discriminator classifies the images produced by the generator as real or synthetic to maximize the probability of correctly classifying the images, thereby improving the translation until the synthetic images can deceive the discriminator (Goodfellow et al., 2014). The cGAN better preserves the pixel numerical values as it requires paired data (inputs and labels) that share mutual information (Mirza and Osindero, 2014). Unlike the cGAN, the Cycle-GAN does not require paired data and is known to lack pixel value accuracy compared with the labels. Hence, cGANs will provide better pixel quantification when generating synthetic images than Cycle-GANs.

Most of the computer vision algorithms are focused in producing relative images that do not weight the quantitative information. Therefore, implementing state-of-the-art cGANs to produce synthetic PET images will not yield accurate quantification of the radiotracer for patients that have different levels of amyloid. It is common practice in modern CNNs to use normalization layers to improve model stability and expedite training time. This, however, eliminates the quantitative values of the images, yielding only relative synthetic images. The aim of this study was to train a cGAN model to synthesize two-dimensional axial plane amyloid-beta PET images from T1-weighted structural MRI images.

3.4   Materials and Methods

3.4.1   Dataset and Acquisition Protocol

In this study, the Open Access Series of Imaging Studies (OASIS-3) (LaMontagne et al., 2019) database was utilized, which is a third-generation database gathered by Washington University with a waiver of informed consent. The database includes 609 cognitively normal subjects and 489 individuals at different stages of cognitive decline that

were classified using the Clinical Dementia Rating (CDR) score (Mendez, 2022). There were 939 subjects with pairs of preprocessed T1-weighted (T1w) MRI and Pittsburgh Compound B (11C-PiB) PET imaging. However, 55 subjects were discarded due to corrupted, missing, or repeated data, leaving 884 subjects to build and assert the model.

The T1w images were obtained with three different scanners: Siemens Biograph mMR 3T, Siemens Trio Tim 3T, and Siemens Sonata 1.5T (Siemens, Erlangen, Germany) with the following parameters: repetition time (TR) of 2300, 2400 or 1900 ms; echo times (TE) of 2.95, 3.16, or 3.93 ms; flip angles (α) of 9 8 or 15 and slice thicknesses of 1.2, 1.0, or 1.0 mm, respectively. The amyloid-beta 11C-PiB PET images were obtained with two modern scanners: Siemens Biograph 40 PET/CT scanner and Siemens ECAT HRPlus 962 scanner (Siemens, Erlangen, Germany), using 11C-PiB as a radiolabeled tracer with an administration dose that ranges between 6 to 20 millicuries (mCi) with a 60 minute dynamic PET scan, storing the tracer uptake in either: 24 X 5 seconds frames, 9 X 20 seconds frames, 10 X 1 minute frames, or 9 X 5 minute frames. MRI images were always acquired before PET images, and there was an average difference of 56 days between the two exams. From the 881 subjects remaining after quality assurance (each with 151 axial plane slices per subject that were not blank), a sample of 552 subjects was used for training (166,704 PET and MRI axial slices) while a sample of 331 subjects was used for evaluation (100264 PET and MRI axial slices). The data was stratified to achieve similar female-to-male ratios, CDRs, and ages. A summary of the stratification is shown in Table 3.1, and distributions of the SUVR and MRI intensities are provided in the supplemental materials.

Table 3.1: Demographics of subjects used for model training and evaluating, balanced the proportion of Males (M), Females (F), Clinical Dementia Rating (CDR) and Age with standard deviation.

	Sex	CDR Mean	Age Mean
Model Training	336 F	0.34(0.58)	71.85(9.04)
	214 M	0.24(0.52)	70.02(9.11)
Model Evaluation	200 F	0.35(0.60)	71.08(9.22)
	131 M	0.24(0.53)	69.59(9.14)

3.4.2 Preprocessing

The PET images provided by OASIS-3 were preprocessed with the PET-Unified Pipeline (PUP, Washington University in St. Louis, St. Louis, Missouri, United States of America). These steps comprised of (1) smoothing to achieve a common spatial resolution of 8 mm that minimized inter-scanner differences and to match other MRI/PET databases image resolution such as the Alzheimer's Disease Neuroimaging Initiative and the Australian Imaging Biomarkers and Lifestyle Study, (2) inter-frame motion correction, and (3) Centiloid standardization for obtaining the Standard Uptake Value Ratio (SUVR). The images were converted from 4- dimensional floating point (4dfp) format to Neuroimaging Informatics Technology Initiative (NIfTI) format using the same PUP pipeline that was used in OASIS-3 and were converted from dynamic PET to static PET by summing the last 1800 seconds of the PET images. The summation

window was selected based on the short PiB half-life of 20 minute. The data is summed between 30 and 60 minutes to provide an accurate representation of the PiB retention consistent with previous studies (Shah et al., 2022).

The T1w MR images were preprocessed with FreeSurfer (2018 release version, Harvard, Cambridge, Massachusetts, United States of America), which provided brain extraction masks, anatomical parcellations, and segmentations of the images. Then the images were converted to NIfTI format using mri_convert utility from FreeSurfer.

After data conversion and extraction, the images were reoriented to match the standard orientation using fslreorient2std utility inbuilt in the FMRIB Software Library (FSL, 6.0 version, FMRIB Analysis Group, Oxford, England, United Kingdom). Brain extraction was performed for the PET images with thresholding, erosion, dilatation, and fill holes operations followed by the FSL Brain Extraction Tool (BET). The purpose of these preprocessing steps was to assist the BET tool in accurately extracting the brain from PET images by compensating for streaking artifacts outside the brain. Removing the streaking artifacts is important, as they can otherwise affect the quality of the image translation model and increase the computational load when training the model.

Both PET and MRI images were co-registered to the Montreal Neurological Institute (MNI) ICBM 152 Nonlinear Symmetric 1 mm template (Fonov et al., 2009) using the Advance Normalization Tools (ANTS) by applying a composite transform that comprised translation and affine transforms for both PET and MRI.

The image dimensions were padded to 256X256X256 array size without affecting the resolution to ensure that the image size matched between the PET and MRI images and to ensure that the model will be able to process them efficiently. To work with a less complex 2D image translation model, both PET and MRI images were converted from 3D to 2D using Python, following which only the axial plane images were extracted. Blank slices were subsequently deleted, leaving only 151 slices per subject. It is known that deep learning models perform better when the input data is normalized; therefore, the MRI images were normalized between 0 to 1. In order to preserve the quantification of PET images, which had a wide range between 0 to 18000 in voxel value, the images were all divided by 1000 to retain quantification between images and also feed the model with images of a practical numerical range. This way, the predicted PET images can be scaled by 1000 to yield the original quantified value.

3.4.3   Model Architecture and Training

The implemented model for image translation followed a cGAN architecture and is shown in Figure 1, it similar to the pix2pix implementation (Isola et al., 2017) in Python using the pytorch deep learning library. A total of 90% of the training data was assigned for training the weights and 10% for validation, with a batch size of 128, Adam as the optimizer, and a learning rate of 2e-4.

The generator portion was implemented as a U-Net with the following endcoder configuration: 4x4 convolutional kernel, stride of 2, LeakyReLU with 0.2 slope as activation function, no batch normalization layers, and 7 filters (32,

33

64, 128, 256) until reaching the bottleneck. The decoder was symmetric to the encoder, with transpose operations and 3 dropout layers with 0.2 as a dropout ratio. The decoder was connected to the encoder with skip connections on every convolutional step. The ReLU was selected as an activation function since other activation functions, such as sigmoid or tanh, constrict the pixel range of the synthetic images and LeakyReLU might slow down the training process and result in negative values. It should be noted that no normalization layers were used in this architecture.

The discriminator portion was implemented with a patch architecture consistent with pix2pix that aims to take parts of the image known as patches while applying convolutional operations and dimensionality reductions until obtaining a representation of the complete image in a MXM grid. This produces individual squares that represent a patch of the image which, after evaluating with Binary Cross Entropy with Logit Loss (BCEwithLogitsLoss), have individual values between 0 and 1, where 0 means that the patch has been classified as real and 1 means that the patch has been classified as synthetic image. The generator was pretrained without the discriminator for 100 epochs, then transferred to the GAN for 700 epochs.

The model was trained and run using Narval and Cedar clusters from the Digital Research Alliance of Canada (Canada, 2023) assigned a 40 GB Nvidia A100 GPU node with 32 GB of RAM and 2 CPUs with a training time of approximately 4 weeks. Once trained, the model can make an inference in less than 3 s.
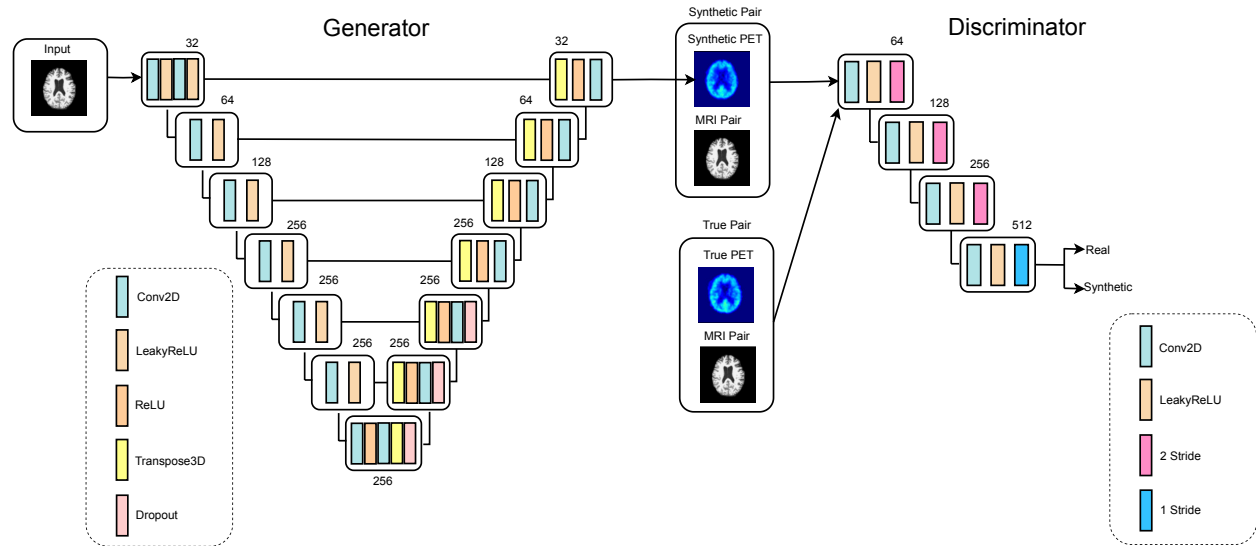
Figure 3.1: Conditional Generative Adversarial Network (CGAN) architecture. The generator is configured as a U-NET architecture that receives the MRI images as input and produces synthetic PET images as output. Subsequently, a pair of MRI & true PET images and a pair of MRI & Synthetic PET is fed into the discriminator to classify the synthetic PET image as true or synthetic. There were a total of 54413313 trainable parameters in the generator and 2765505 in the discriminator.

3.4.4 GAN Losses and Custom Loss

It is desirable to train the model with a specific focus on the brain to prevent the background from dominating the image. Therefore, a custom loss function was implemented that treated regions inside a brain mask differently than the surrounding area. The purpose of the mask was to guide the model towards the information within the brain, considering it to be more important. The mask was computed by thresholding each individual PET image in the batch by 35% of the maximum pixel value of the given image for both true and synthetic PET images during training. A value of 35% was found to eliminate background noise while leaving the brain intact. All images were reviewed for quality assurance.

Therefore, the custom loss function can be described by using L1 loss (i.e., L1 norm equates to the mean absolute error) between the masked real and synthetic PET images as described in the following equation:

$$L_1(G_{\text{mask}}) = E_{(x,y,z)} \|y_{\text{mask}} - G(x,z)_{\text{mask}}\|_1 \tag{3.1}$$

Where y_{mask} is the true PET image masked, $G(x,z)_{\text{mask}}$ is the synthetic PET image masked produced by the generator, based on the inputs x, the MRI image and z, is the latent variable.

Then this equation is added to the equation 2.7, which was discussed in chapter 2, where the final objective function was implemented as follows:

$$\text{cGAN} = \arg\min_G \max_D L_{\text{cGAN}}(G,D) + \lambda_1 L_1(G_{\text{mask}}) + \lambda_2 L_1(G) \tag{3.2}$$

Where the masked L1 loss function is multiplied by a Lambda of 1000 ($\lambda_1 = 1000$) and the non-masked L1 loss function is multiplied by a Lambda of 100 ($\lambda_2 = 100$) to regularize the model and prevent overfitting.

3.4.5 Model Quality Evaluation

A regression and Bland-Altman analysis was performed in a representative sample of 4 subjects with different ages and dementia levels: 2 young (58 and 60 years old) and 2 old (79 and 80 years old), each with one healthy and one AD case. The MNI-ICBM 152 Nonlinear Asymmetric Probability atlas that contains probability maps for gray and white matter with the probabilities for gray and white matter segmentation thresholded to 40% was used to generate gray and white matter masks (Fonov et al., 2009). The masks were applied to the synthetic and real PET images to extract values amyloid-beta in both gray and white matter.

To validate the model, the evaluation set of 331 subjects with 151 slices each were used to evaluate the generalization of the image translation model to unseen images. Two performance metrics were calculated between the synthetic and real PET images to evaluate the quality of the image synthesis: Structural Similarity Index Measure (SSIM) and

Peak Signal to Noise Ratio (PSNR).

SSIM can report a value between 0 to 1, where 0 means that both images are completely different and 1 that both images are the same. SSIM is the product of three comparative measures: luminescence (l), contrast (c), and structure (s) that can be calculated as follows:

$$l(x,y) = \frac{2\mu_x\mu_y + C_1}{\mu_x^2 + \mu_y^2 + C_1} \tag{3.3}$$

$$c(x,y) = \frac{2\sigma_x\sigma_y + C_2}{\sigma_x^2 + \sigma_y^2 + C_2} \tag{3.4}$$

$$s(x,y) = \frac{\sigma_{xy} + C_3}{\sigma_x\sigma_y + C_3} \tag{3.5}$$

Where the standard deviation, σ_x, is calculated as:

$$\sigma_x = \left(\frac{1}{N-1}\sum_{i=1}^{N}(x_i - \mu_x)^2\right)^{\frac{1}{2}} \tag{3.6}$$

And σ_{xy} calculated as:

$$\sigma_{xy} = \frac{1}{N-1}\sum_{i=1}^{N}(x_i - \mu_x)(y_i - \mu_y) \tag{3.7}$$

Where μ_x is the pixel mean of true PET image defined as x, μ_y is the pixel mean of synthetic PET image defined as y, σ_x is the pixel standard deviation of x, σ_y is the pixel standard deviation of y, and σ_{xy} is the covariance of x and y. C1, C2, and C3 are variables used to avoid undefined value when denominator is 0, and were calculated as:

$$C_1 = (K_1L)^2 \tag{3.8}$$

$$C_2 = (K_2L)^2 \tag{3.9}$$

$$C_3 = \frac{C_2}{2} \tag{3.10}$$

$K_1 = 0.01$, $K_2 = 0.03$ and L is the dynamic range for pixel value. Finally, the SSIM was given by:

$$\text{SSIM}(x,y) = [l(x,y)]^\phi \cdot [c(x,y)]^\omega \cdot [s(x,y)]^\gamma \tag{3.11}$$

Where ϕ, ω, and γ control the component weight that has in the SSIM score and are set to 1. Therefore, a high SSIM indicates high similarity in these 3 components.

Since computing SSIM in the complete image leads to inaccurate results, as described in the seminal paper (Zhou

et al., 2004), it is necessary to compute SSIM on different regions of the image. This can be achieved with a gaussian kernel (11X11) that computes SSIM while sliding across the image and reports the mean SSIM that is used to evaluate the overall image quality.

PSNR can be described as the ratio between the maximum possible power of an image and the power of distortion noise that affects image quality. It can report a value between 0 to infinite, where 0 means that the images are completely different and infinite is that the images are the same, it is calculated as follows:

$$\text{PSNR} = 20\log_{10}\left(\frac{L-1}{\text{RMSE}}\right) \tag{3.12}$$

Where L is the number of maximum possible intensities that can be found in the real image, and RMSE is the root mean squared error calculated in between the real and synthetic images. Therefore, it is desirable to maximize both the SSIM and PSNR scores. In this study an acceptable result was considered a $SSIM > 0.85$ and $PSNR > 18$ (Horé and Ziou, 2010).

3.4.6 Blinded Review

To assert that the proposed model can generate synthetic amyloid-beta PET images that are useful and indistinguishable from the real ones, a blinded image reading was performed. Two experienced clinicians evaluated the images. Each was presented with a set of 100 images from 100 subjects in the test-data with a 50% split between real and synthetic images. The readers were only provided with the subject ID, age and sex, and were instructed to respond the following questions for each subject: 1) Is the image assessable? negative or positive, 2) What is the burden of Amyloid? low, moderate or high, 3) What is the burden in gray matter? low, moderate or high 4) What is the burden in white matter? low, moderate or high, and 5) Do you think the image is real or synthetic?

3.4.7 Statistical Analysis

From the example cases where pixel values were measured in gray and white matter, the R^2 was computed to determine the quality of the image sythesis and Bland-Altman plots was used to evaluate the bias. From the blinded review, a Pearson Chi-squared test was performed to inspect for differences between readers' performance in true and synthetic PET images. This test is designed to determine if there is a difference between two sets of categorical data. A p-value less than 0.05 was considered significant. For question 5, the accuracy was computed to evaluate whether the readers could determine the difference between real and synthetic images. The statistical analyses were implemented with Python, version 3.8, using the scipy library, v1.11.2.

3.5 Results

An example of the synthetic amyloid-beta PET images is shown in Figure 3.2, where the full 3D volume is reassembled from the 2D slices. The model produced synthetic PET images with a high degree of similarity compared to the true images. Figure 3.3 shows four representative subjects (2 younger and 2 older) having different levels of CDR and amyloid burden. The model was able to generate high quality amyloid-beta synthetic PET images for subjects that have atrophy driven by AD and atrophy characteristic from aging.

Example Synthetic PET Images
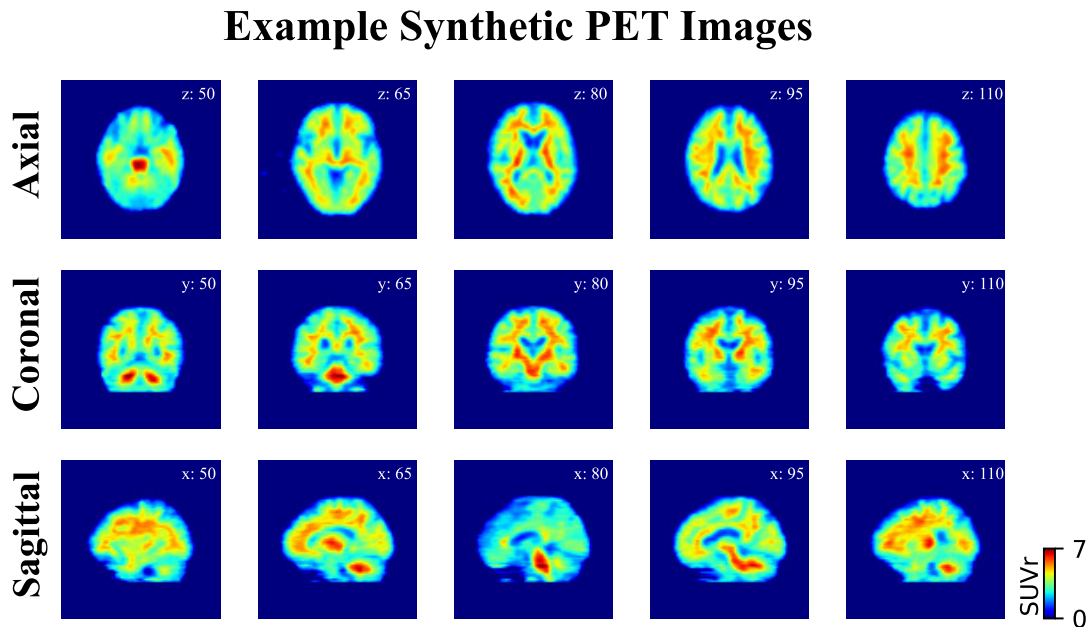
Figure 3.2: Representative example of synthetic amyloid-beta PET images. Five slices in each orientation are shown with the numbers that correspond to the slice in the MNI atlas. The subject is age 64 with a Clinical Dementia Rating (CDR) score of 2. Here the 2D images are all independently estimated and then reassembled into a 3D volume. In general, there is a high degree of continuity in the SUVR between slices, although there are some minor discrepancies towards the base of the brain.

Comparing the estimations between the 58- and 84-year-old subjects, it can be seen that the 84 year old, cognitively normal subject (Figure 3.3, Row 3) has brain atrophy (e.g., reduced brain volume through increased ventricular volume and reduced cortical thickness) due to aging but low amyloid-beta while the CDR 2 subject (Figure 3.3, Row 2) has lower levels of atrophy with mild amyloid-beta burden and the model was able to estimate the amyloid-beta PET images for these cases with a high degree of similarity. The 79-year-old CDR 3 subject (Figure 3.3, Row 4) has the highest levels of atrophy and amyloid-beta burden.

Figure 3.4 shows the SUVR pixel-wise comparison in gray and white matter between true and synthetic PET images (same as those presented in Figure 3.3). The model reached an R2 greater than 0.96 for the 4 example cases.

Comparison Between Amloyid Predictions
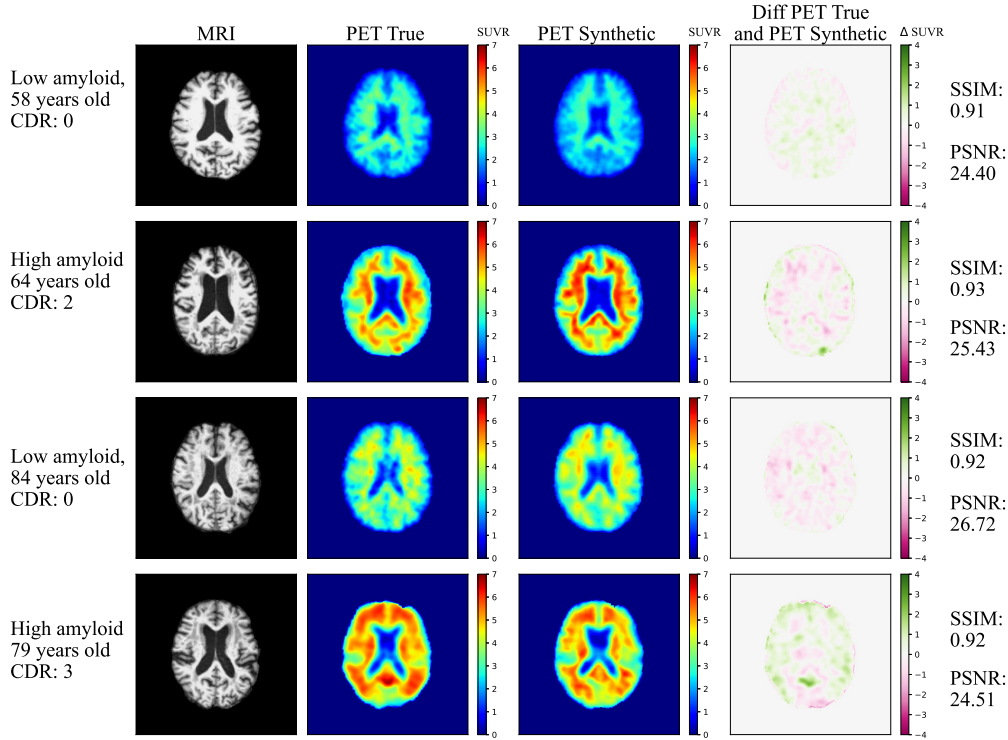
Figure 3.3: Representative model generated synthetic PET images for 4 subjects. There is a high degree of similarity between the true and synthetic PET images. In the panel 4 columns can be seen: the MRI input image (left), the true PET image (center-left), the synthetic PET image (center-right), and the error map between true and synthetic (right). Both true and synthetic PET color bar represents SUVR concentrations, and the error map color bar is the difference in SUVR. The Structural Similarity Index Measure (SSIM) and Peak Signal-to-Noise Ratio (PSNR) are calculated using the 151 slices from each subject, as well as the Clinical Dementia Rating (CDR) score are provided for context.

The Bland-Altman density plots are also shown in Figure 3.4 and reveal a low bias with a mean no greater than 0.372 (8.0%) in the four presented subjects. Given that the SUVR has a typical difference of 100% between healthy and dementia patients, $<10\%$ is considered an acceptable bias.

Histograms of the SSIM and PSNR metrics on the evaluation subjects are reported in Figure 3.5, reaching a mean SSIM averaged across all the subjects of 0.905, SSIM in luminescence of 0.959, SSIM in contrast of 0.931, and SSIM in structure of 1.014, indicating that the model can synthesize PET images from MRI with a high level of similarity in all components. Across the whole population, there was also a high mean PSNR of 22.685, indicating that the model can generate high quality synthetic amyloid-beta PET images. Figure 3.4 shows that the model has a good performance in this evaluation cohort.
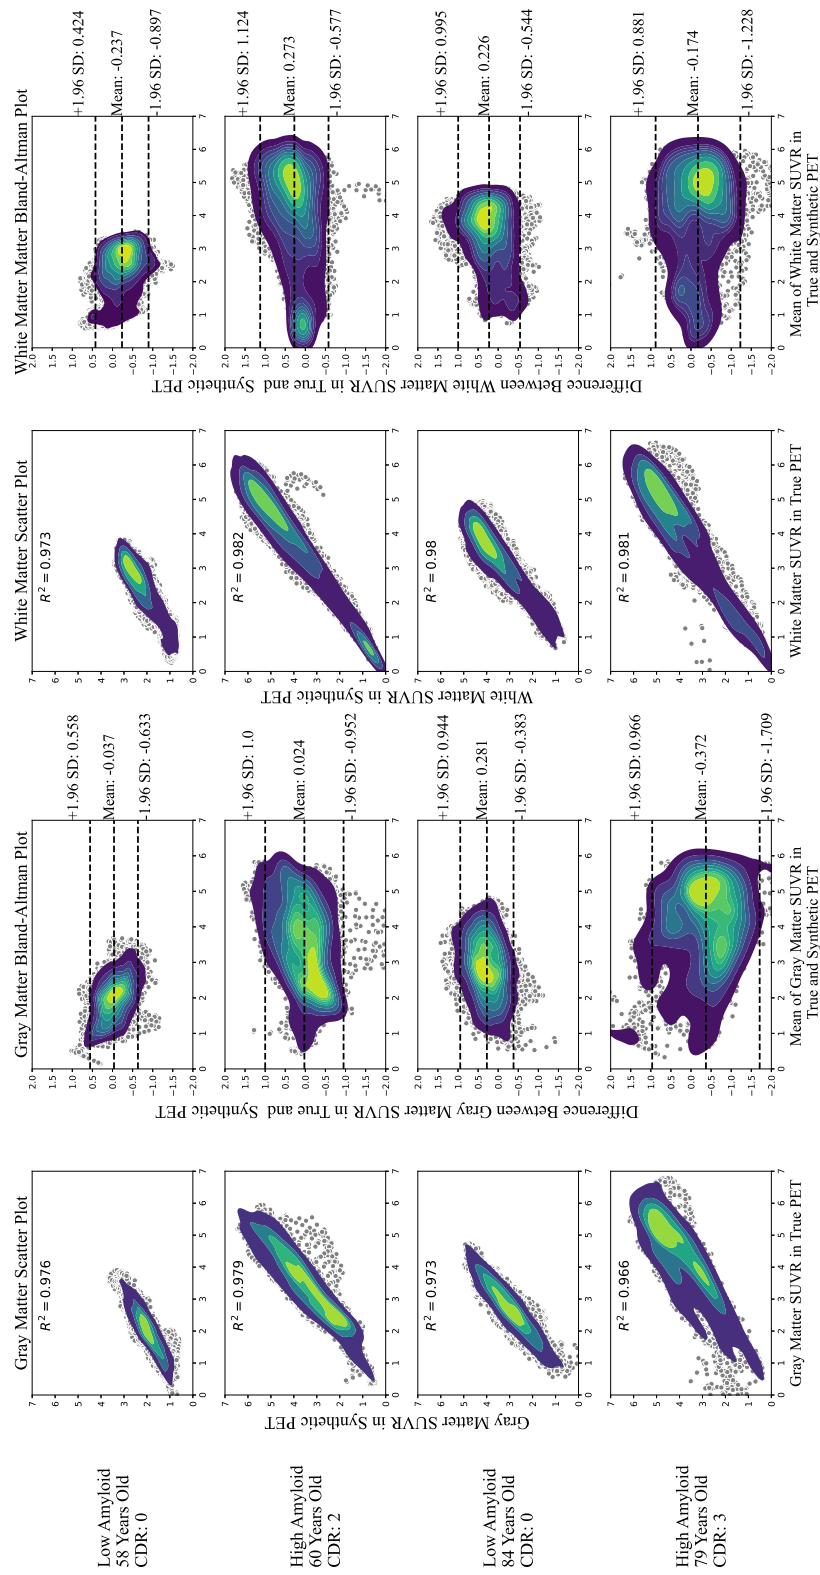
Figure 3.4: Amyloid-beta SUVR comparison in gray matter with regression plot (left) and Bland-Altman plot (center-left), and white matter regression plot (center-right) and Bland-Altman plot (right) with density. Regression plots show that the model creates quantitative estimations across different cases reaching a high R2, the density shows the most repeated SUVR values. Bland-Altman plots show the model bias in the presented subjects, where the middle line is mean and represents the bias while outer lines represent the confidence intervals.
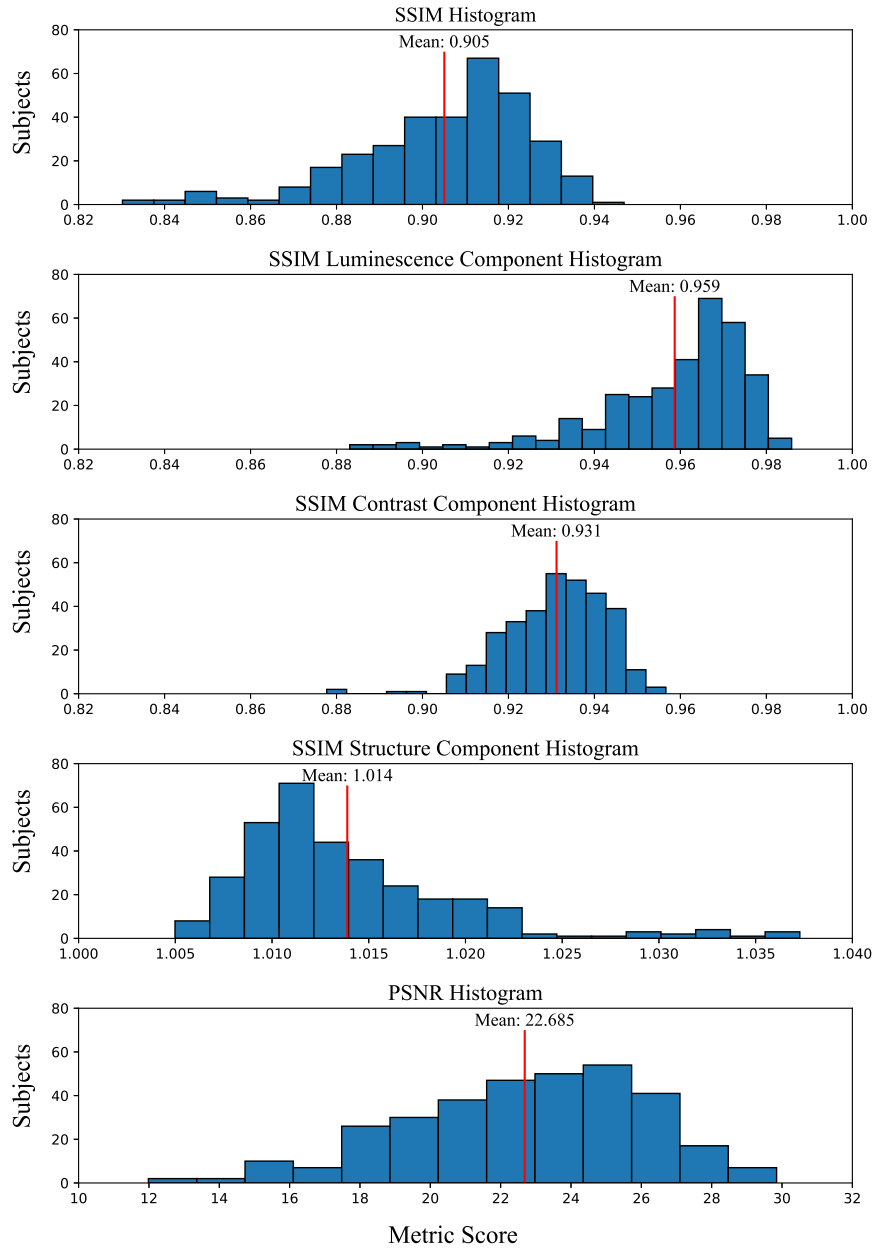
Figure 3.5: Structural Similarity Index Measure (SSIM) and Peak Signal-to-Noise Ratio (PSNR) histogram from 331 evaluated subjects (151 slices each) with a mean SSIM of 0.905, SSIM in luminescence of 0.959, SSIM in contrast of 0.931, and in structure of 1.014 and mean PSNR of 22.685.

In evaluating the blinded expert readings, the Pearson Chi-squared test showed no significant difference ($p = 0.68$) between the real and synthetic images. The accuracy scores for determining whether the images were real or synthetic were 46% and 54%, respectively. These two findings indicated that the readers' performance was not different with either real or synthetic images, and they could not tell the difference between them.

3.6 Discussion

In this work, a pipeline was proposed and validated to translate images from structural MRI to amyloid-beta PET. This work is important as expert readers cannot infer the amyloid-beta burden from structural MRI alone, making the inference extremely valuable as an amyloid-beta estimate can now be achieved from MRI, thereby reducing the cost of the scan, reducing the invasiveness, and making it more accessible than doing traditional PET imaging. In the validation cohort the model was able to reach a high mean SSIM and high mean PSNR, which means the model can generate high quality synthetic PET images. While the subjects shown in Figure 3.3 are four representative cases of the translation quality, they were selected to highlight the model performance at different ages and disease progression.

SSIM and PSNR are metrics commonly used to evaluate the similarity of synthetic images produced by a model against the true images (Armanious et al., 2018; Horé and Ziou, 2010). SSIM is computed with a sliding window (as described in the original paper (Zhou et al., 2004)) and a kernel size of 11, but there was no meaningful change (i.e., ¡0.1% mean SSIM) in the results in using sizes of 7, 15, or 19. To be more rigorous, we separated the SSIM into its three underlying components: luminance, contrast, and structure. In addition to this quantitative numerical evaluation, a blinded reading comparison was also conducted where the readers had similar performance and were not able to tell the difference between the real and synthetic images. All these evaluation approaches indicated a high degree of similarity between the real and synthetic images.

Batch normalization layers were not used in the proposed model and a new loss function was implemented that masked the images to penalize information inside of the brain. Synthesized images presented a high degree of similarity in representative subjects that span the range of the database (old/young, Male/Female, low/high disease severity), and complemented with a quantitative SUVR comparison in gray and white matter, where Figure 3 contains 4 subjects (2 young and 2 old) that are representative of the model performance on different age and disease severity.

Performing image translation to produce quantitatively accurate 11C-PiB amyloid-beta PET images is a challenge in comparison to other studies, which have focused on translation from structural MRI to 18F-fluorodeoxyglucose (FDG) PET imaging (Sikka et al., 2021; Zhang et al., 2022). Furthermore, these studies did not present a quantitative evaluation of SUVR in gray and white matter. The studies also normalized the images and implemented architectures with normalization layers with the purpose of avoiding the common problems in GAN models, such as exploding and vanishing gradients. Therefore, these studies have limited the utility to only estimate relative PET tracer concentrations

across the individual's brain.

It is known that GAN models are prone to artifacts such as inpainting or adding structures that are not in the target image. While these artifacts are commonly seen when training unsupervised GANs such as cycle-GANs, they are rarely observed with supervised machine learning such as the conditional GAN used in this work (Isola et al., 2017). Hence, the trained model created synthetic PET images with no inpainting or structural artifacts.

3.6.1 Limitations

The amyloid-beta estimation needs further improvements to reach higher performance. This is linked to the current state of the field of computer vision that is focused on generating relative images instead of quantitative ones. Such limitations could be addressed by exploring different cost functions that could focus on evaluating contrast and different GAN architectures

The current model was trained using images from OASIS-3 that has specific patient characteristics and experimental paradigms, such as an injected dose ranging between 6 and 20 mCi. This is a classical limitation of machine learning, and can be addressed by incorporating images from other databases that have pairs of MRI and amyloid-beta PET images with different populations and different acquisition schemes.

The OASIS-3 dataset provided smoothed 8 mm spatial resolution PET data to diminish the inter-scan differences and to diminish the differences with other PET/MRI databases. It would be meaningful to explore the effect of using both higher resolution MRI and PET to build future models. Higher resolution imaging would allow for regional comparisons across the cortex, white matter fiber tracts, and subcortical structures.

The OASIS-3 dataset carefully selected subjects to avoid including patients with non-dementia related diseases; therefore, the model was not trained with non-dementia pathologies that might also cause brain atrophy. Such a database that includes MRI and 11C-PiB amyloid-beta PET in, for example, alcoholics, is not yet available, but would be desirable to include for training. Indeed, it might be the case that pathologies with atrophy similar to Alzheimer's Disease synthesize an inaccurate PET image and estimate amyloid-beta where there is none present. This problem can be addressed by training a model with the OASIS-3 dataset and additional pairs of MRI and amyloid-beta PET images from with other pathologies when such data becomes available.

The image preprocessing pipeline makes use of several steps, including: brain extraction and image registration, to ensure that the cGAN could be properly trained as it requires images that are aligned. Sufficient and accurate preprocessing is required for this model to be trained.

The proposed model does not implement normalization layers; this could lead to an exploding gradient problem when there are different image contrast levels. This issue could be addressed by exploring different techniques that could help to suppress the exploding gradients without losing the quantitative information from the images, such as

spectral normalization to normalize the weights of the discriminator instead of the input images (Miyato et al., 2018).

The model was trained in the 2D space using axial plane images, meaning that it is not currently considering other planes that also have meaningful information to assess AD. Furthermore, since the model is training on slices it is also considering some slices that have only a small amount of information. As shown in Figure 3.2, there can be minor discrepancies between neighbouring 2D slices. Future work will aim to implement a 3D cGAN model architecture that will work on the whole volume instead of slices.

Future work could also explore different model architectures that are not GAN based, as they are difficult to train because of the adversarial training that entails finding the equilibrium between generator and discriminator. For example, vision transformers are known to learn based on a self-attention mechanism that might be meaningful to yield quantitatively accurate PET images and are known to outperform CNN based models (Dosovitskiy et al., 2020). For this work to be useful in the clinic, additional validation and seeking of regulatory approval is necessary.

3.6.2   Chapter Conclusion

This chapter showcased a 2D image translation model with elimination of normalization steps that can synthesize amyloid-beta PET images from structural MRI with a high degree of similarity to the true label cases in cognitively normal patients and in patients with various severities of AD dementia. While this work requires several improvements that are addressed in the next chapter, this paper shows promising results and the feasibility to create an amyloid-beta PET image from a structural MRI one.

Chapter 4

Image Translation for Estimating Three-Dimensional Amyloid-Beta PET from Structural MRI

Chapter 3 presented the data preprocessing and implementation of a two-dimensional image translation model. In this Chapter, the methodology is changed with several improvements added to the image translation architecture based on the discussion and limitations of Chapter 3. This Chapter presents the changes in the methodology, including a different data stratification to train a model with more subjects than in the two-dimensional case. In this work, the synthetic PET images achieved higher similarity metrics. Results from this Chapter have been submitted to the upcoming ISMRM meeting and will be submitted for publication, but for the sake of brevity, the introduction and repeated methods similar to Chapter 3 are removed.

4.1 Abstract

Background: Alzheimer's Disease Progression can be defined by amyloid-beta accumulations and brain atrophy. These can be measured with Positron Emission Tomography (PET) and structural MRI, respectively. PET is known to be expensive, invasive, and harmful (by inducing ionizing radiation); meanwhile, MRI is cheaper, less invasive, and ionizing radiation-free, but it cannot measure amyloid-beta. Since there is a known relationship between amyloid-beta and brain atrophy, PET images can potentially be inferred from structural MRI.

Purpose: This study aims to develop a three-dimensional image translation model employing a Conditional Generative Adversarial Network (cGAN) with spectral normalization to accurately synthesize amyloid-beta PET images from

structural MRI.

Study Type: Retrospective Population: 882 adults (346 males/536 females) at different stages of cognitive decline based on the clinical dementia rating score (0:cognitively normal, 0.5:questionable, 1:mild cognitive impairment, 2:moderate cognitive impairment, 3:severe cognitive impairment). 617 pairs of PET/MRI images for model training and 265 for evaluation.

Field Strength/Sequence: 3T, T1-weighted structural (T1w)

Assessment: Structural Similarity Index Measure (SSIM) and Peak Signal-to-Noise Ratio (PSNR) were used in the evaluating cohort to assess the performance of the model. Such metrics compare the similarity of the global synthetic PET images with the global real PET images. SSIM was computed in the global image and reported its components: luminance, contrast, and structure.

Statistical Test: Gray and white matter quantitative assessment was also performed between the synthetic and real PET images using mean absolute error (MAE) for each subject.

Results: The model generated high-quality synthetic PET images across the evaluation cohort, reporting a mean SSIM=0.958 and mean PSNR=28.836. The gray and white matter assessment reported an average MAE¡0.15 in the evaluation cohort.

Conclusion: A Three-dimensional cGAN can generate quantitative synthetic amyloid-beta PET images from structural MRI with a high degree of global similarity and accuracy in contrast estimation.

4.2 Keywords

Amyloid-beta, Conditional Generative Adversarial Networks, Image Translation, Magnetic Resonance Imaging, Positron Emission Tomography

4.3 Introduction

This work advances upon the previous study by introducing a three-dimensional cGAN that is capable of processing entire volumes rather than slices and incorporates a normalization technique that can be reverted, thus maintaining the radiotracer quantification. It also implements spectral normalization to address the exploding gradients problem and stabilize the training process.

Furthermore, a gray and white matter assessment is presented, where the radiotracer estimation in both regions was compared with the true images to show the quantitative error between synthetic and real PET images.

4.4 Materials and Methods

4.4.1 Dataset and Acquisition Protocol

The Open Access Series of Imaging Studies 3 (OASIS-3) (LaMontagne et al., 2019) (Washington University in St Louis, St. Louis, Missouri, United States of America) provided pairs of PET and MRI images. The database included 609 cognitively normal subjects and 489 at different stages of cognitive decline that were classified using the CDR score. OASIS-3 contained 939 subjects with pairs of preprocessed T1w MRI and PiB PET imaging. However, 59 subjects were discarded due to corrupted, missing, or repeated data leaving 884 subject pairs to build and assert the model as discussed in Chapter 3, Materials and Methods section.

As consistent with Chapter 3, the T1w images were obtained with three different scanners: Siemens Biograph mMR 3T, Siemens Trio Tim 3T, and Siemens Sonata 1.5T (Siemens, Erlangen, Germany) with the following parameters: repetition time (TR) of 2300, 2400 or 1900 ms, echo times (TE) of 2.95, 3.16 or 3.93 ms, flip angles (α) of 9 8r 15 and slice thicknesses of 1.2, 1.0 or 1.0 mm, respectively.

The amyloid-beta 11C-PiB PET images were obtained with two modern scanners: Siemens Biograph 40 PET/CT scanner and Siemens ECAT HRPlus 962 scanner (Siemens, Erlangen, Germany), using 11C-PiB as a radiolabeled tracer with an administration dose that ranges between 6 to 20 millicuries (mCi) with a 60 minutes dynamic PET scan, storing the tracer uptake in either: 24 X 5 seconds frames, 9 X 20 seconds frames, 10 X 1 minute frames, 9 X 5 minute frames.

MRI images were consistently acquired before PET images, with an average difference of 56 days between the two exams. From the 882 subjects remaining after quality assurance, a sample of 617 MRI/PET pairs were used for training and 265 pairs for evaluation. The data was stratified in equal proportions of female-to-male ratio, CDR, and age. A summary of the stratification is shown in Table 4.1.

Table 4.1: Three-Dimensional Training and Evaluation Demographics. There are equal proportions of Males (M) and Females (F), Clinical Dementia Rating (CDR), and Age with mean and standard deviation.

	Sex	CDR Mean	Age Mean
Model Training	373 F	0.34(0.52)	69.79(9.12)
	244 M	0.35(0.60)	71.56(9.13)
Model Evaluation	163 F	0.25(0.56)	70.02(9.14)
	102 M	0.24(0.53)	71.58(9.12)

4.4.2 Preprocessing

The PET images were preprocessed by OASIS-3 with the PET-Unified Pipeline (PUP, Washington University in St Louis, St. Louis, Missouri, United States of America). The dynamic PET images underwent thorough smoothing, interframe motion correction, and Centiloid standardization to obtain the Standard Uptake Value Ratio (SUVR). As discussed in Chapter 3, the method to convert from dynamic to static PET is similar, where the 30 to 60 minutes of the dynamic scan were summed to obtain an accurate representation of PiB retention.

SUVR normalization was performed by dividing the static PET images between the mean tracer found in the cerebellum cortex, a standard reference region for PiB imaging (Price et al., 2005; Jack et al., 2009). This is because the tracer's concentration in the cerebellum remains stable throughout the scanning period. SUVR normalization is crucial to account for variability in PiB tracer quantification given that each subject received an injected dose adjusted for factors like weight, body mass, and duration of the scan. The PET can display varying PiB tracer concentrations. Therefore, SUVR normalization standardizes these values, ensuring a consistent representation across all subjects (Price et al., 2005). The measurements of the reference region were provided by OASIS-3.

The T1w MR images were preprocessed with FreeSurfer (2018 release version, Harvard, Cambridge, Massachusetts, United States of America), which provided brain extraction masks. Brain extraction was performed for the PET images as described in Chapter 3: a series of erosion and dilatation operations followed by the FSL BET tool to remove the streaking artifacts produced by the PET reconstruction to avoid including non-brain related information in the modeling process. Finally, the MRI and PET images were co-registered to the MNI mask, as consistent with Chapter 3 methodology.

The MRI images were scaled between 0 to 1 using min-max scaling and the PET images were normalized between 0 to 1 using z-score normalization with a mean of 0 (μ=0) and standard deviation of 49.72 (σ=49.72), unlike min-max scaling, z-score can be denormalized to retain the PET tracer quantification between images while feeding the model with a practical numerical range. This way, the predicted PET images can be denormalized by multiplying by the population standard deviation and adding the mean.

4.4.3 Model Architecture and Training

The implemented model for image translation followed a cGAN architecture and is shown in Figure 4.1, it is similar to Chapter 3 implementation using the monai deep learning framework, with a batch size of 6, Adam as optimizer with a learning rate of 2e-4 and momentum parameters $\beta_1 = 0.5$, $\beta_2 = 0.999$. Those momentum values are commonly used in cGAN training setting, where the generator and discriminator are constantly adapting based on each other's performance. A lower β_1 aids the optimizer to be more responsive to the changes in both networks while a higher β_2 helps to stabilize the variances produced by the adversarial training. In the pix2pix implementation, these values were

empirically obtained.

The generator portion was implemented as a U-NET with the following encoder configuration: 4x4 convolutional kernel, a stride of 2, LeakyReLU with 0.2 slope as activation function, instance normalization, and 7 filters (32, 64, 128, 256) until reaching the bottleneck. The decoder was symmetric to the encoder, with transpose operations, 3 dropout layers with 0.2 as a dropout ratio, and sigmoid as the final activation function to restrict the generated voxels between 0 and 1. The decoder was connected to the encoder with skip connections on every convolutional step, as shown in Figure 4.1.

The discriminator portion was implemented as in Chapter 3 as well, with the addition of spectral normalization on intermediate convolutional blocks which is known to avoid exploding gradients and stabilize training (Miyato et al., 2018), as shown in Figure 4.1.

The model was trained and run using Narval and Cedar clusters from the Digital Research Alliance of Canada (Canada, 2023) assigned a 40 GB Nvidia A100 GPU node with 32 GB of RAM and 2 CPUs with a training time of approximately 2 weeks equivalent to 500 epochs.
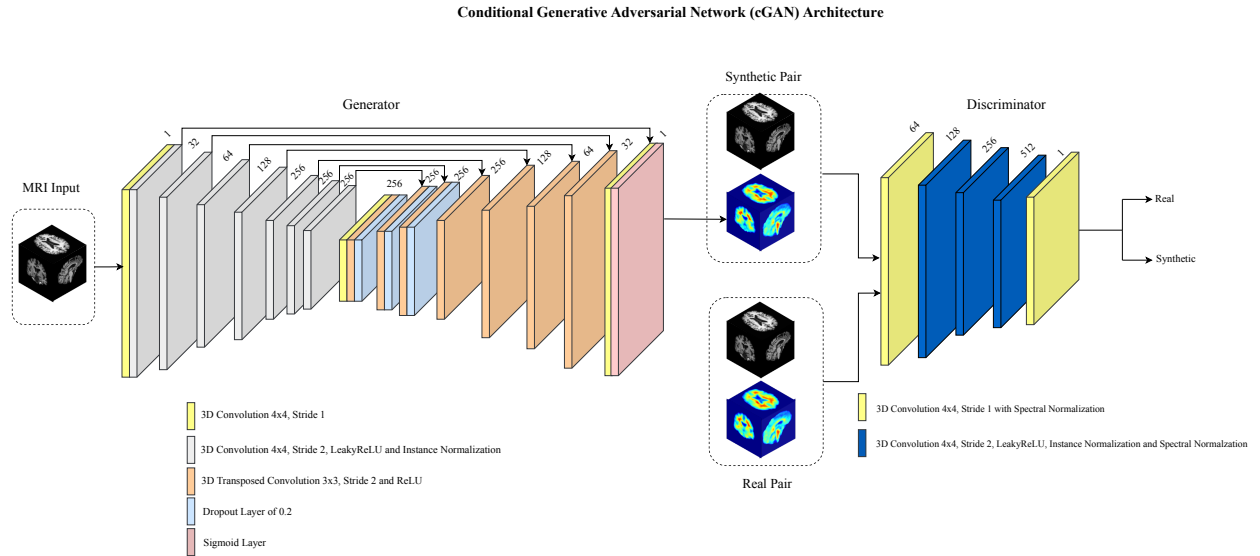
Figure 4.1: Three-dimensional Conditional Generative Adversarial Network (cGAN) architecture. The generator is configured as a three-dimensional U-Net architecture that receives the MRI images as input and produces synthetic PET images as output. Subsequently, a pair of MRI & true PET images and a pair of MRI & Synthetic PET is fed into the discriminator to classify the synthetic PET image as true or synthetic. Every intermediate convolutional layers in the discriminator have spectral normalization.

4.4.4 GAN Losses and Custom Loss

To aid the model in focusing on the image content instead of the background, an improved version of the custom mask loss as in Chapter 3 was implemented, known as brain area mask loss. The purpose of this loss was to guide the model

towards information within the brain and disregarding the background. The brain area mask was obtained by using the torch.masked_select function that returns a masked array that only consists of values greater than 0 using as a reference to the true PET image. The final objective equation is equation 3.2 which is:

$$\text{cGAN} = \arg\min_{G}\max_{D} L_{\text{cGAN}}(G,D) + \lambda_1 L_1(G_{\text{mask}}) + \lambda_2 L_1(G) \tag{4.1}$$

Where the brain area masked L_1 loss function is multiplied by a Lambda of 10 ($\lambda_1 = 10$) and the non-masked L_1 loss function is multiplied by a Lambda of 1 ($\lambda_2 = 1$) to add importance to the brain area mask loss over the non-masked one and prevent overfitting.

4.4.5 Model Quality Evaluation

To validate the image similarity between the synthetic and real PET images on the evaluation cohort, a similar approach was followed to that in Chapter 3. Two performance metrics: Structural Similarity Index Measure (SSIM) and Peak Signal to Noise Ratio (PSNR) were calculated between the synthetic and real PET images to evaluate the quality of the image synthesis. Please see Chapter 3 for a detailed formulation of these metrics.

4.4.6 Gray and White Matter Analysis

Since SSIM and PSNR measure the global image similarity, it is necessary to evaluate whether the image translation model can accurately synthesize the amyloid-beta present in gray and white matter in the PET images. This analysis was performed by extracting both regions from the synthetic and real PET images with the MNI ICBM probabilistic mask. After extracting each individual's gray and white matter, Mean Absolute Error (MAE) was computed between the real and synthetic gray and white matter in the evaluating cohort.

4.5 Results

Representative three-dimensional Synthetic amyloid-beta PET images are shown in Figures 4.2, 4.3 and 4.4, for the axial, sagittal, and coronal views, respectively. Two representative subjects (CN and AD) with different levels of CDR, amyloid-beta burden, and age (69 and 79 years of age) were selected. The model was able to generate high-quality three-dimensional amyloid-beta synthetic PET images for healthy and AD subjects. The 69 year old CN subject shows more structural integrity in the brain (i.e., ventricles) and less amyloid-beta concentrations in all views, while the 79 year old AD subject shows a more pronounced atrophy (i.e., ventricles) with more tracer concentrations, especially in the frontal lobe as shown in Figures 4.2 and 4.4. The model was able to generate the respective synthetic amyloid-beta PET images for both cases while demonstrating consistency across slices with low error.

Comparison Between Amyloid-Beta PET Predictions Axial View

63 Years Old, Cognitively Normal, CDR 0

79 Years Old, Alzheimer's Disease, CDR 2

Figure 4.2: Axial Views from three-dimensional volume of MRI input, true PET label, synthetic PET and difference map between real and synthetic PET for cognitively normal (CN) and Alzheimer's Disease (AD) subject. It can bee seen that AD subject has higher amyloid-beta retention and atrophy around the in the frontal lobe and around the ventricles compared with CN subject. The synthetic PET images highly resemble the true PET images.

Comparison Between Amyloid-Beta PET Predictions Coronal View

63 Years Old, Cognitively Normal, CDR 0

79 Years Old, Alzheimer's Disease, CDR 2
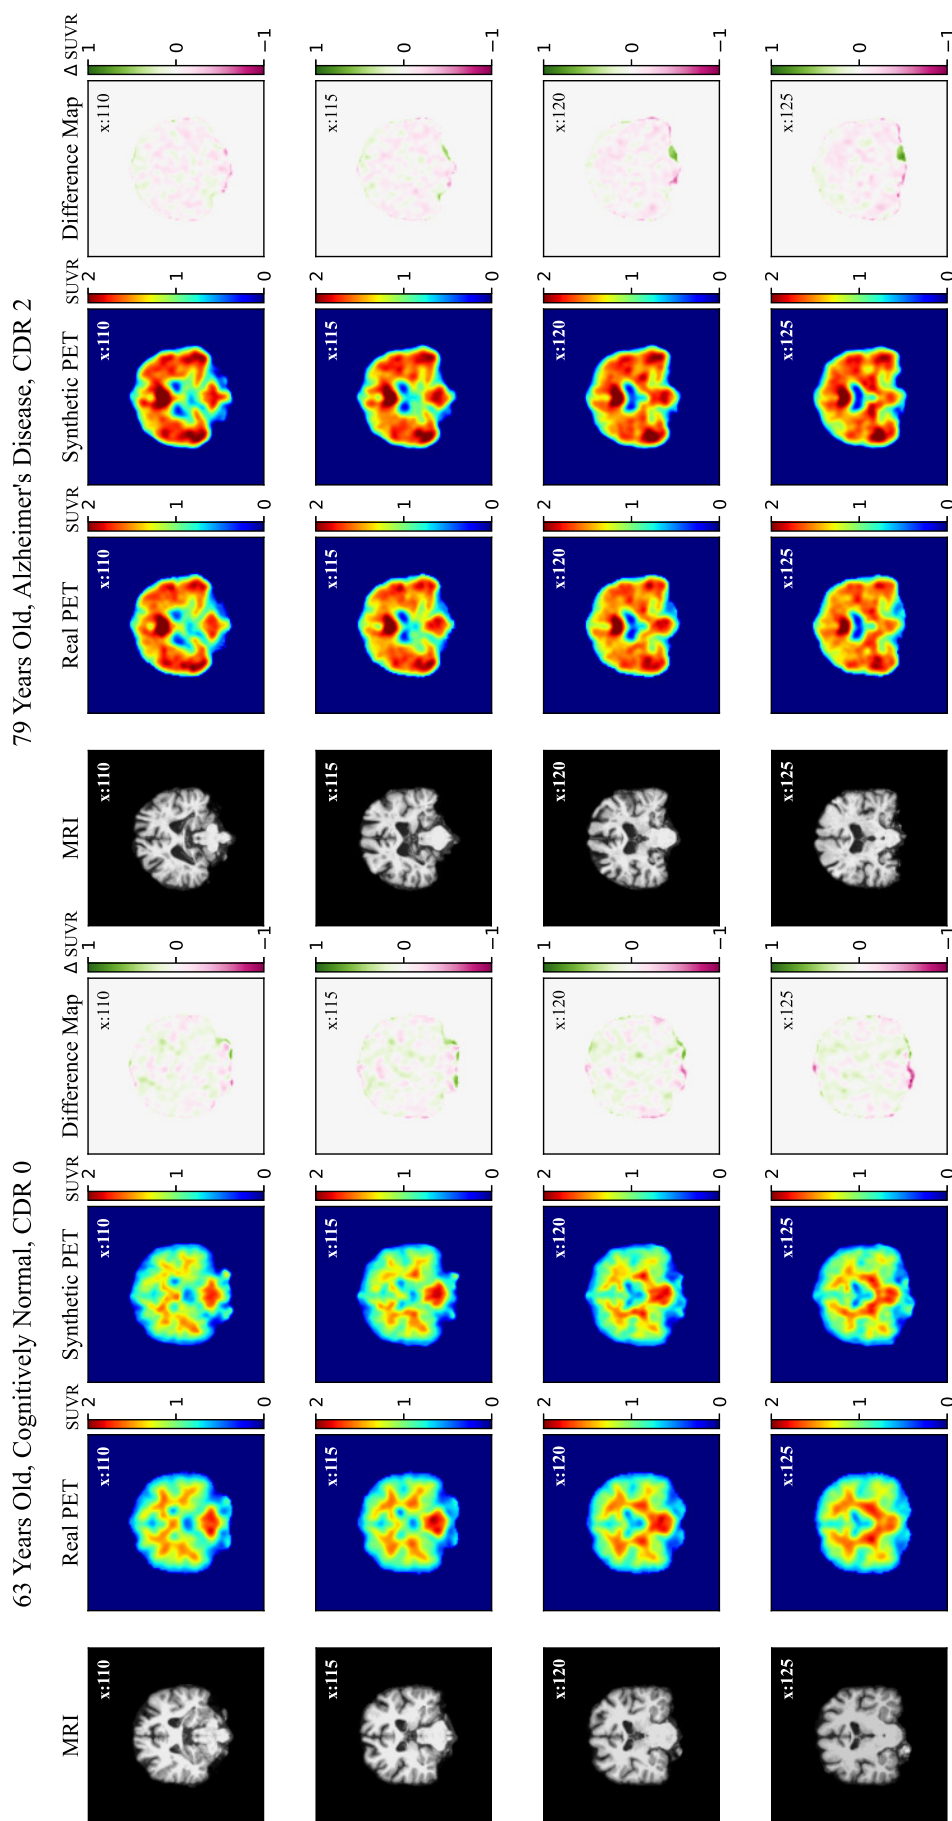
Figure 4.3: Coronal Views from three-dimensional volume of MRI input, true PET label, synthetic PET and difference map between real and synthetic PET for cognitively normal (CN) and Alzheimer's Disease (AD) subject. Unlike the CN subject, the AD subject show higher amyloid-beta concentrations concentration around the gray and white matter. The model is able to create a synthetic PET image that looks close to the true PET both in structure and SUVR.

Comparison Between Amyloid-Beta PET Predictions Sagittal View

63 Years Old, Cognitively Normal, CDR 0

79 Years Old, Alzheimer's Disease, CDR 2

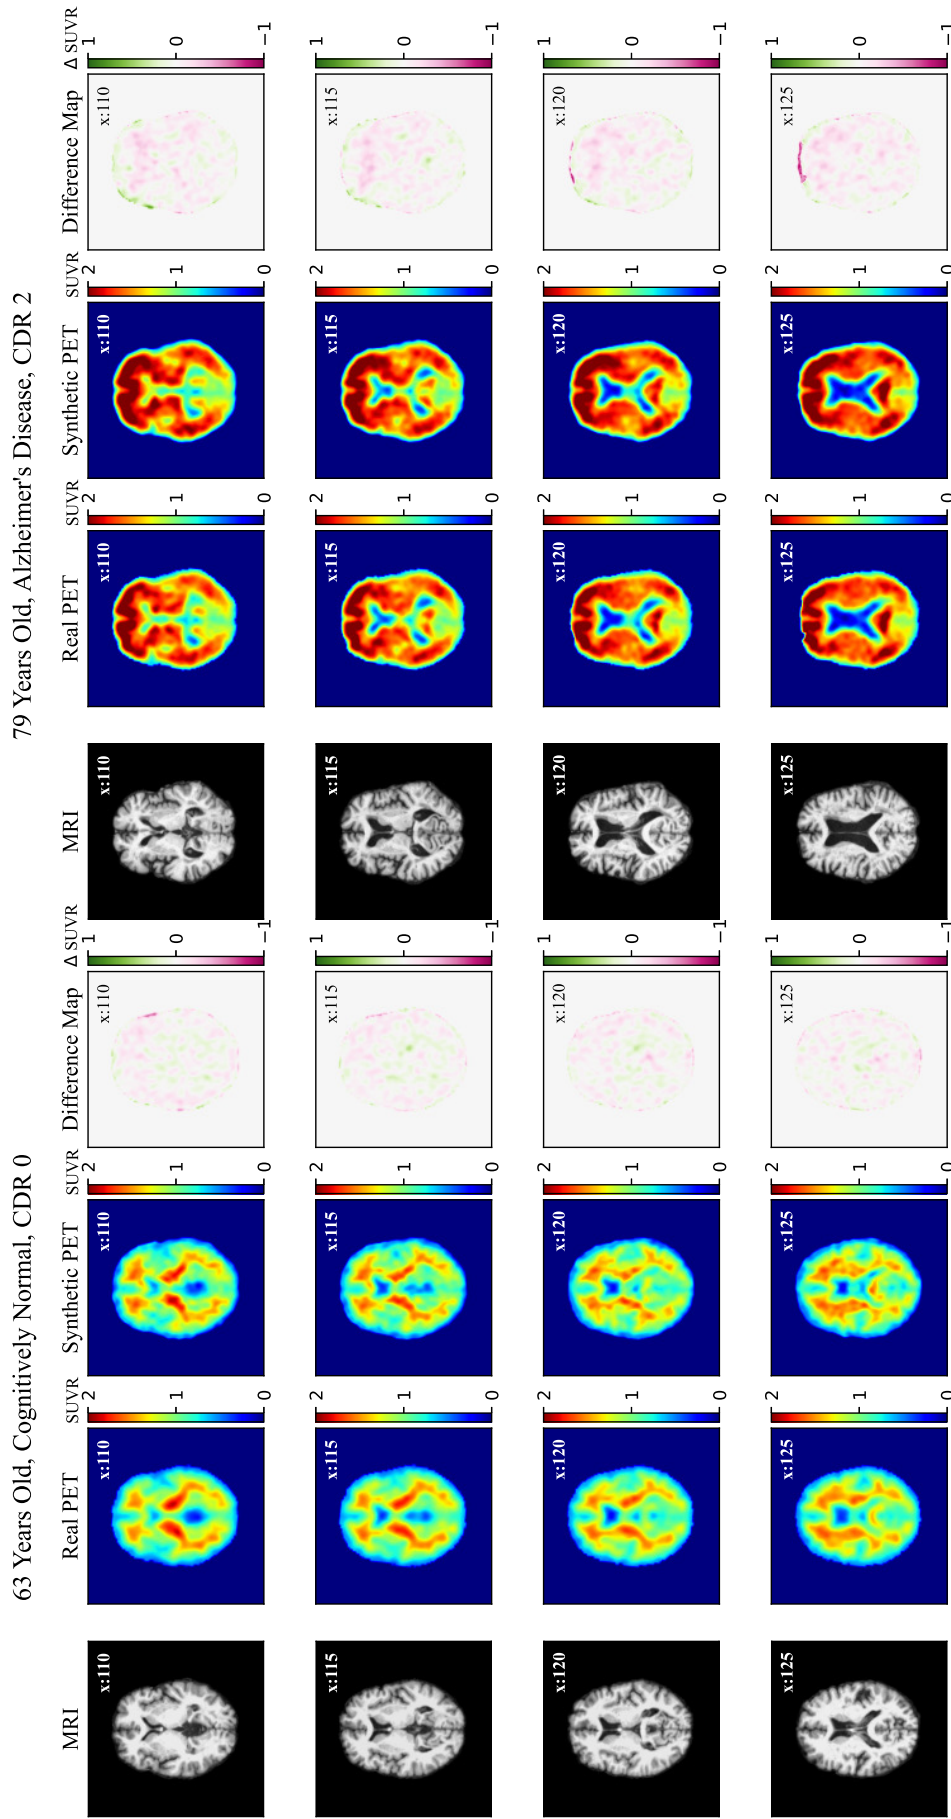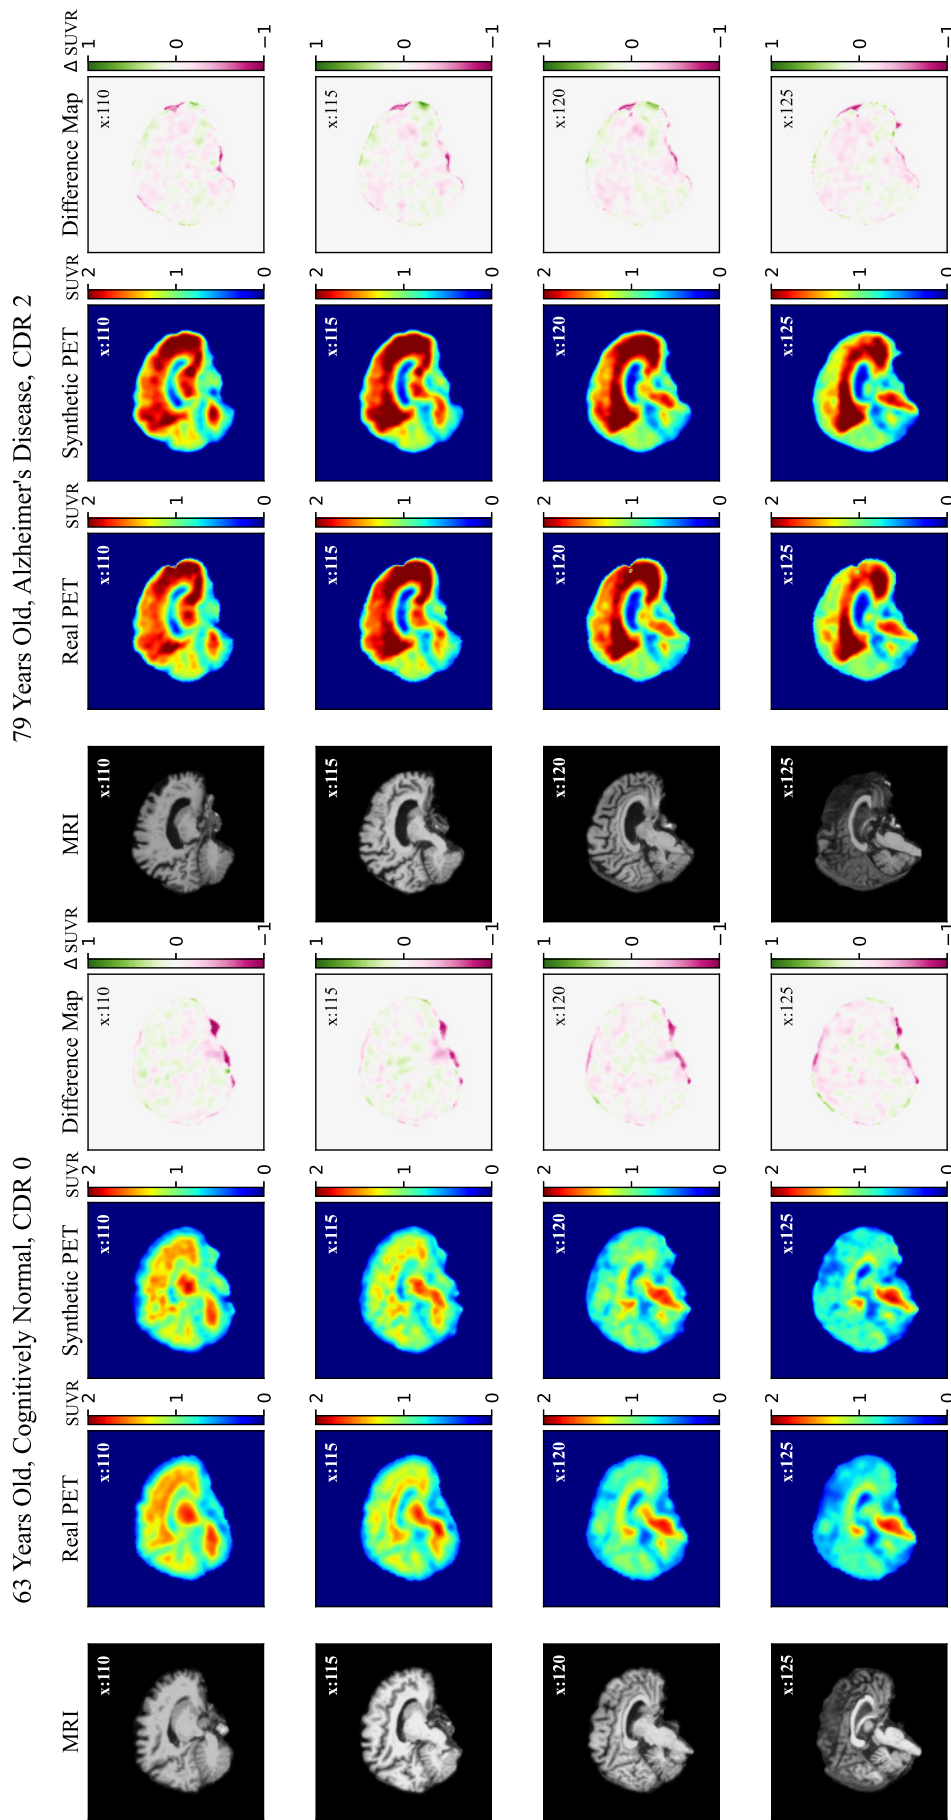Figure 4.4: Sagittal Views from three-dimensional volume of MRI input, true PET label, synthetic PET and difference map between real and synthetic PET for cognitively normal (CN) and Alzheimer's Disease (AD) subject. The Alzheimer's Disease (AD) subject shows a higher concentration in the frontal lobe than the CN subject, the synthetic PET image has minor discrepancies with the true PET in regions with low levels of amyloid.

The difference map shows the error between synthetic and real amyloid-beta PET, the maps show that synthetic PET images highly resemble the true PET ones and while there are subtle discrepancies between the images, the radiotracer localization is being properly generated for all cases. This localized estimation can be seen well in the sagittal views, where the model was able to generate the tracer in the most important regions, particularly the frontal and parietal lobes.

Histograms of the SSIM and PSNR metrics on the evaluating cohort are reported in Figure 4.5, reaching a mean SSIM averaged across all the subjects of 0.958, SSIM in luminescence of 0.985, SSIM in contrast of 0.965, and SSIM in structure of 1.008, indicating that the model can synthesize amyloid-beta PET images from T1w MRI with a high level of similarity in all components. Across the whole population, there was also a high mean PSNR of 28.836, indicating that the model can generate high-quality global three-dimensional synthetic amyloid-beta PET images.

4.5.1   Gray and White Matter Analysis Results

The gray and white matter analysis reports that the average MAE in gray and white matter for all the subjects in the evaluating cohort is less than 0.15, as shown in Figure 4.6, where the violin plots represent the data distribution of MAEs in gray and white matter for all the evaluating 265 subjects cohort. The plots indicate that the error between synthetic and real PET is low since the average error was less than 3% in gray and white matter. Therefore, the model can create synthetic amyloid-beta PET images that are highly accurate in SUVR compared to the real ones.

Model Histogram of SSIM Components and PSNR Across 265 Subjects
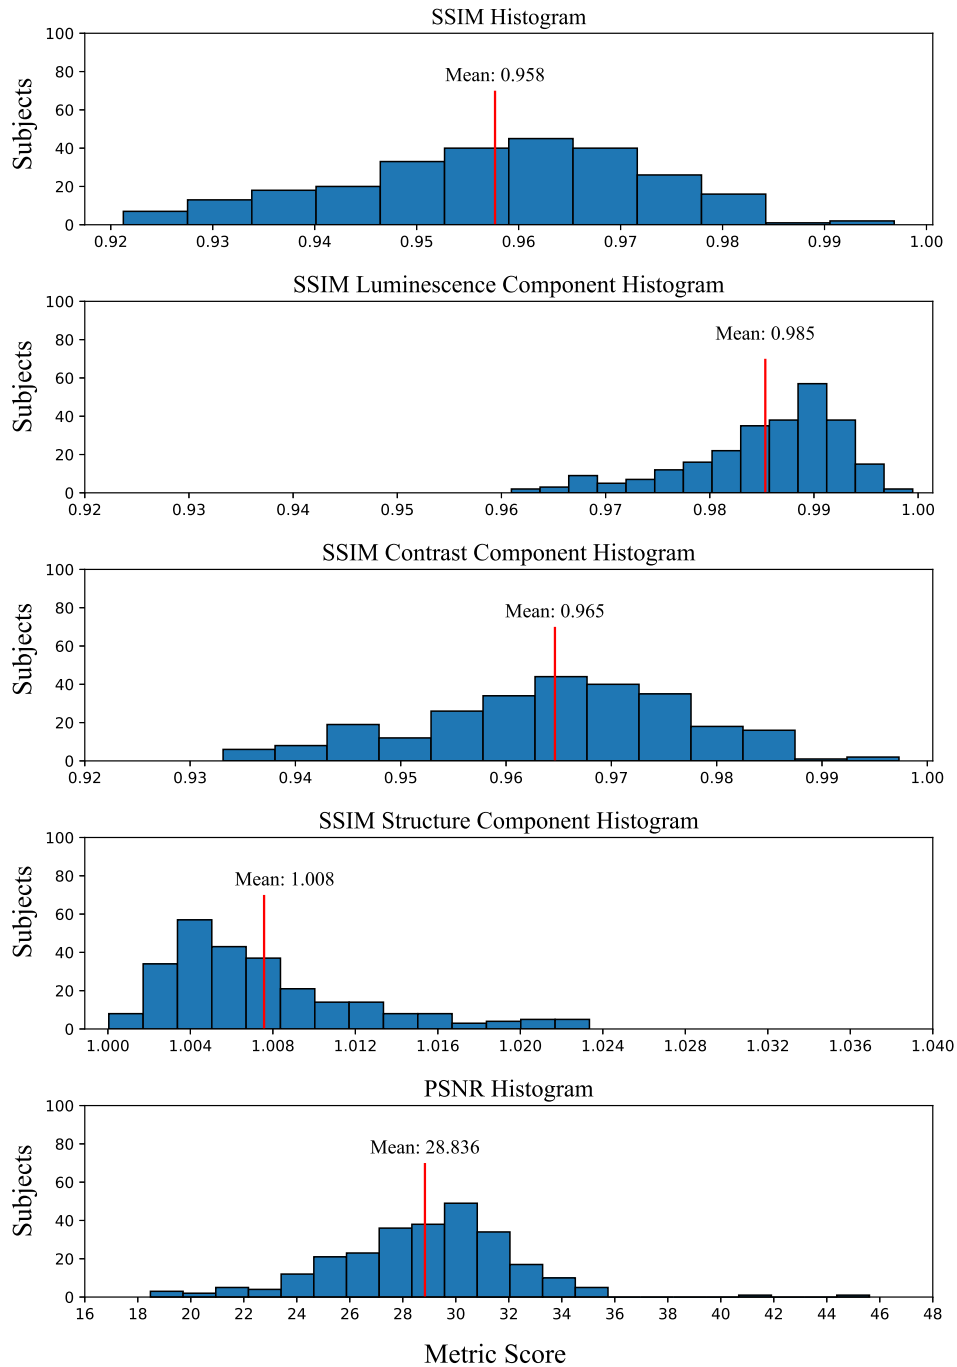
Figure 4.5: Structural Similarity Index Measure (SSIM) and Peak Signal-to-Noise Ratio (PSNR) histogram from 265 evaluated subjects with a mean SSIM of 0.958, SSIM in luminescence of 0.985, SSIM in contrast of 0.965, and in structure of 1.008 and mean PSNR of 28.836.

**Violin plot of Mean Absolute Error Between Synthetic and Real PET images
in Gray and White Matter**
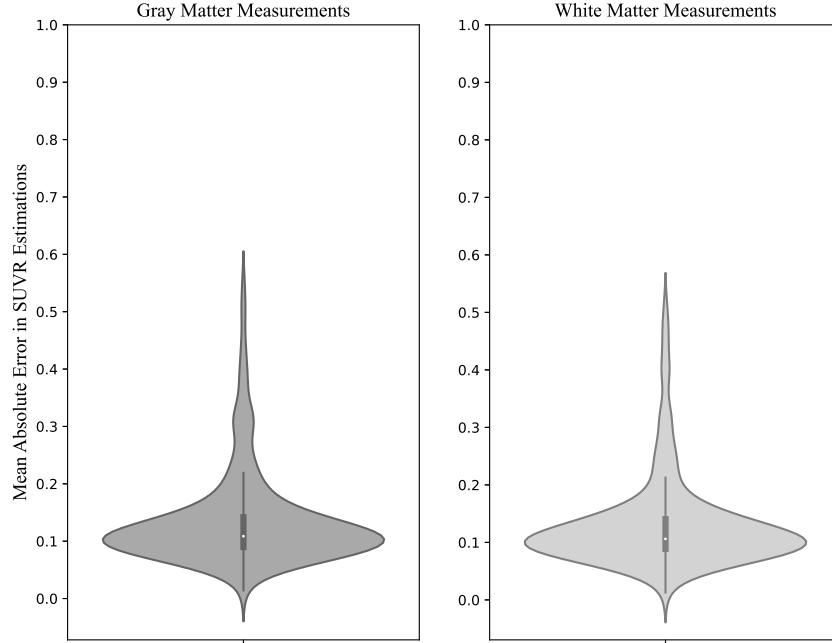
Figure 4.6: Gray (left) and white (right) matter mean absolute errors (MAEs) across evaluating cohort. The violin plots describe the distribution of MAEs where the average report a 0.135 for gray matter and 0.130 for white matter. Such results indicate that the model can generate highly accurate synthetic amyloid-beta PET images.

4.6 Discussion

In this work, a pipeline was proposed and validated to translate structural MRI into amyloid-beta PET images in three-dimensional. This work is important as it shows that it is possible to generate extremely valuable amyloid-beta images from cheaper MRI ones, hence reducing the cost of the scan, reducing the invasiveness, and offering a more accessible technology than doing traditional PET imaging. In the evaluation cohort, the model was able to reach a high mean SSIM and high mean PSNR, which means the model can generate high-quality synthetic 3D PET images that outperformed metrics reported in past two-dimensional implementation (Vega et al., 2023).

In comparison with the two-dimensional implementation in Chapter 3, the PET images underwent z-score normalization to restrict the PET images values during training. Z-score normalization can be denormalized, and it does not affect the tracer quantification and improves the model performance. Instance normalization layers were implemented on the generator and discriminator (Figure 4.1) instead of batch normalization since it is dependent on batch size, where larger batch size improve the model stability, however, in a three-dimensional implementation large batch sizes

is not feasible so instance normalization can be utilized regardless of batch size (Armanious et al., 2019; Ulyanov et al., 2016). Spectral normalization was also implemented, as it is known to improve training stability in GANs by avoiding exploding gradients during training of the discriminator. This technique was not implemented in the past two-dimensional work in Chapter 3. Therefore, the applied changes to the proposed 3D implementation show an overall improvement in the metrics compared with past work.

Performing image translation to produce quantitatively accurate three-dimensional 11C-PiB amyloid-beta PET images is a challenge that has not been addressed in previous studies Ouyang et al. (2023); Lan et al. (2021). The studies also implemented min-max scaling of the PET images that cannot be denormalized when generating the synthetic images. Therefore, these studies have limited the utility to only estimate relative PET tracer concentrations across the individual's brain.

4.6.1 Limitations

While the amyloid-beta estimations show high global SSIM and PSNR with low MAE in gray and white matter, there is still room for improvement in such estimations. As can be seen in Figure 4.6, there are outliers present where the MAE in both gray and white matter was high in comparison to the average. This can potentially be addressed by incorporating the gray and white matter masks into the training and penalizing the error between synthetic and real.

Since the OASIS-3 dataset selected subjects that only had dementia-related diseases, the model was not trained with subjects that affect the integrity of the brain structure due to other pathologies or conditions (i.e., stroke, alcoholism) (Martín-González et al., 2020; Luijten et al., 2022) that have no effects on the amyloid-beta within the brain. Indeed, it might be the case that an MRI image from a subject with such conditions might synthesize an inaccurate amyloid-beta PET image. This problem can be addressed by training a model with pairs of MRI and PET images from those types of subjects when a dataset that contains that data becomes available.

4.6.2 Future Work

Future work could explore the use of multi-contrast MRI inputs to estimate a single amyloid-beta PET image such as fluid attenuated inversion recovery (FLAIR). It could also be meaningful to explore different model architectures that are not GAN based, as GANs are difficult to train due to the adversarial training procedure. For example, vision transformers are known to learn based on a multi-head attention mechanism that might be meaningful to yield quantitatively accurate PET images and is known to outperform CNN based models for some applications(Dosovitskiy et al., 2020). Diffusion models are also known to excel at image translation; work such as the palette diffusion image translation architecture has shown promising results in this field (Saharia et al., 2021).

For this work to be useful in the clinic, additional validation and seeking of regulatory approval will be necessary.

4.6.3 Chapter Conclusion

A three-dimensional cGAN architecture with spectral normalization can synthesize amyloid-beta PET images from structural MRI with a high degree of similarity to the true label cases in cognitively normal patients and in patients with various severities of AD dementia.

Chapter 5

Comparison Of Two-Dimensional and Three-Dimensional Image Translation Models to Synthesize Amyloid-Beta PET from structural MRI.

Chapter 3 presented an image translation model that synthesizes two-dimensional axial amyloid-beta PET images from structural MRI. While the two-dimensional model reached high SSIM and PSNR metrics, it had room for improvement as discussed in Chapter 3, section 3.6. In Chapter 4, an image translation model that synthesizes three-dimensional amyloid-beta PET images from structural MRI was presented. This model addressed limitations found in the two-dimensional approach. In this chapter, a comparison is performed between the two presented models and discussed.

5.1 Abstract

With the aim of obtaining more accurate diagnostic tools for AD, image translation models to estimate amyloid-beta PET from structural MRI offer promising advancements. These models can be trained either in two- or three-dimensions. Training two-dimensional axial image translation models has some advantages such as providing high numbers of training samples, giving more variability to the model and using less computational resources. However, these models struggle to reconstruct a three-dimensional volume from only axial slices, as inconsistencies are bound to happen due to training on individual slices instead of the whole volume image. Three-dimensional image

translation models address this limitation, as they can be trained using volumes instead of slices, thereby producing more consistent results. This study shows the impact of using the same training and evaluating cohorts to train two- and three-dimensional image translation models. The results on the evaluating cohort were compared between the two model types, showing that the three-dimensional model outperformed the two-dimensional one, producing more consistent volumes and improving the SSIM by 6% and PSNR by 26.3%. Therefore, this analysis demonstrates the superior efficacy of three-dimensional models in synthesizing more accurate and consistent amyloid-beta PET images, showing a significant advancement in AD detection.

5.2   Introduction

This chapter achieves the third aim of the thesis by performing a rigorous and fair comparison between the two-dimensional and three-dimensional models.

Performing image translation to estimate amyloid-beta PET from structural MRI is a task that has not been properly addressed before. Past work in the field of translating from MRI to another PET tracer has shown promising results with two-dimensional (Ouyang et al., 2023; Sun et al., 2022; Vega et al., 2023) and three-dimensional (Lan et al., 2021; Hussein et al., 2022) models; however, none of these works perform a comparison between the two- and three-dimension models.

Two-dimensional axial image translation models have some advantages, such as being trained with more images. In Chapter 3, section 3.4.1 explained that each subject's image volume was divided into 151 axial slices, providing 100264 PET and MRI axial slices for training that added more samples to the training process. It also allowed for the reduction of the data as a number of images in the volume are outside the brain.

However, training with two-dimensional space also has some disadvantages. Although the model was able to yield high-quality synthetic PET images, it had inconsistencies along the coronal and sagittal views when reconstructing the three-dimensional volume back from the axial slices, as displayed in Figure 3.2. This is bound to happen, as the model has no context of whether a set of slices belongs to the same subject or not and therefore generates slightly different amyloid-beta estimations on neighboring slices for the same subject. Such a limitation affected the SSIM and PSNR metrics reported in Figure 3.5 and gave the motivation to explore three-dimensional models.

Three-dimensional image translation models have the advantage of being trained with volumes instead of slices, addressing the limitations found in the two-dimensional approach by producing consistent results across the axial, coronal, and sagittal views. Thereby, the model is able to reach higher similarity metrics with the only drawback being that it requires more computational resources in order to be trained.

The three-dimensional approach also improved some of the methodology as described in Chapter 4, sections 4.2.2 and 4.2.3 while being trained on a different stratified cohort to increase the number of subjects. This chapter focuses

on training a three-dimensional model with the same data stratification cohort as described in Chapter 3, but using the enhancements described in Chapter 4, with the purpose of doing a rigorous and fair comparison of the models.

5.3 Materials and Methods

After establishing the context behind this study, it is necessary to describe the specific materials and methods employed. This study used the OASIS-3 dataset described in Chapters 3 and 4 to train the three-dimensional image translation model. To make this comparison possible and more rigorous, the image translation model was trained using the same stratified data as in Chapter 3, Table 3.1 where 552 pairs of PET and MRI images were used for training and 331 pairs for evaluation. It should be noted that for this model, the images were not split in axial slices, instead the full volume was used.

5.3.1 Preprocessing

The preprocessing steps were performed as in Chapter 4 section 4.2.2, where all the PET images were SUVR normalized with the cerebellum cortex mean SUVR as the reference region. The T1w MRI images were preprocessed with FreeSurfer (2018 release version, Harvard Cambridge, Massachusetts, United States of America), which provided the brain extractions. The PET images were brain extracted as described with Chapter 3, section 3.4.2. Finally, the PET and MRI images were co-registered to the MNI standard template.

To train the image translation model with a practical data range, the PET images were normalized between 0 and 1 using z-score normalization with a mean of 0 and standard deviation of 49.72 and the MRI images were scaled between 0 to 1 using min-max scaling. The reason behind this preprocessing step for the PET images was so they can be denormalized back by multiplying by the standard deviation and adding the mean.

5.3.2 Model Training

These models were trained following the same architecture and settings as described in chapter 4, section 4.2.3, and Figure 4.1. Using the same settings is crucial to conduct a fair comparison between the two-dimensional and three-dimensional approach and show the impact that the methodological improvements in Chapter 4 had on the similarity metrics.

5.3.3 Model Quality Evaluation and Comparison

To validate the model, the same steps as in Chapters 3 and 4 were performed where the SSIM and PSNR were calculated in the evaluation cohort, to make this comparison, three-dimensional SSIM and PSRN was calculated for

both models, where the three-dimensional volumes were generated from the two-dimensional model axial images. Chapter 3 details the formulation of these metrics. Then, the same subject used in Figure 3.2 was generated in this three-dimensional implementation and compared with the two-dimensional one across the axial, coronal, and sagittal views.

5.4   Results

Figure 5.1 shows a comparison between the two-dimensional and three-dimensional improved synthetic PET image in the axial view. Since the two-dimensional model was trained using this view, both models generate synthetic PET images that closely resemble the real PET images.

Figure 5.2 shows the coronal plane comparison, which demonstrates the limitations of the two-dimensional model in generating this plane from slices alone. Such inconsistencies are evidenced in the difference map where the reconstructed synthetic PET views show more pronounced error in comparison with the three-dimensional synthetic PET image views that show a closer resemblance with the real PET image views.

Lastly, the sagittal view comparison is shown in Figure 5.3 where the reconstructed synthetic PET views discrepancies are more evident, especially in slices 65, 80, and 95. The three-dimensional synthetic PET image achieved lower error in this view as well.

The models were also compared with the SSIM and PSNR metrics to assess the image quality and fidelity. The three-dimensional model reached higher metrics in comparison with the two-dimensional model. Table 5.1 shows a summary of SSIM and PSNR in both models. The three-dimensional model reported a mean SSIM of 0.96 and a mean PSNR of 28.64 in the validation cohort while the two-dimensional model reported a mean SSIM of 0.868 and a mean PSNR of 22.677 in the same validation cohort. These results indicate that the three-dimensional model improved the mean SSIM by 10% and mean PSNR by 26.3% compared to the two-dimensional one.

Figure 5.4 shows a detailed distribution of the SSIM in all its components and PSNR obtained by the three-dimensional model in the validation cohort. It can also be denoted that this model reported a SSIM range from 0.92 to 0.98 with PSNR range from 18 to 35, while the two-dimensional model reported a SSIM range from 0.83 to 0.94 with PSNR range from 12 to 29, meaning that the three-dimensional model images outperformed the two-dimensional model.

Table 5.1: Overall structural similarity index (SSIM) and peak signal-to-noise ratio (PSNR) comparison between two- and three- dimensional image translation models, both models were trained using the same data stratification as described in chapter 3, section 3.4.1

	SSIM	PSNR
Two-Dimensional Model	0.868	22.685
Three-Dimensional Model	0.96	28.648

Comparison Between Two-Dimensional and Three-Dimensional Model Axial View

Figure 5.1: Comparison between two-dimensional and three-dimensional models in the axial plane. Five images are presented, from left to right, the real PET image, the synthetic two-dimensional PET image, the synthetic three-dimensional PET image, the difference map between synthetic two-dimensional and real PET, the difference map between synthetic three-dimensional and real PET image, and the difference map between synthetic two-dimensional and three-dimensional PET. In this plane, both models generate images with high degree of similarity with the true images because the two-dimensional model was trained using axial plane images.

Comparison Between Two-Dimensional and Three-Dimensional Model Coronal View
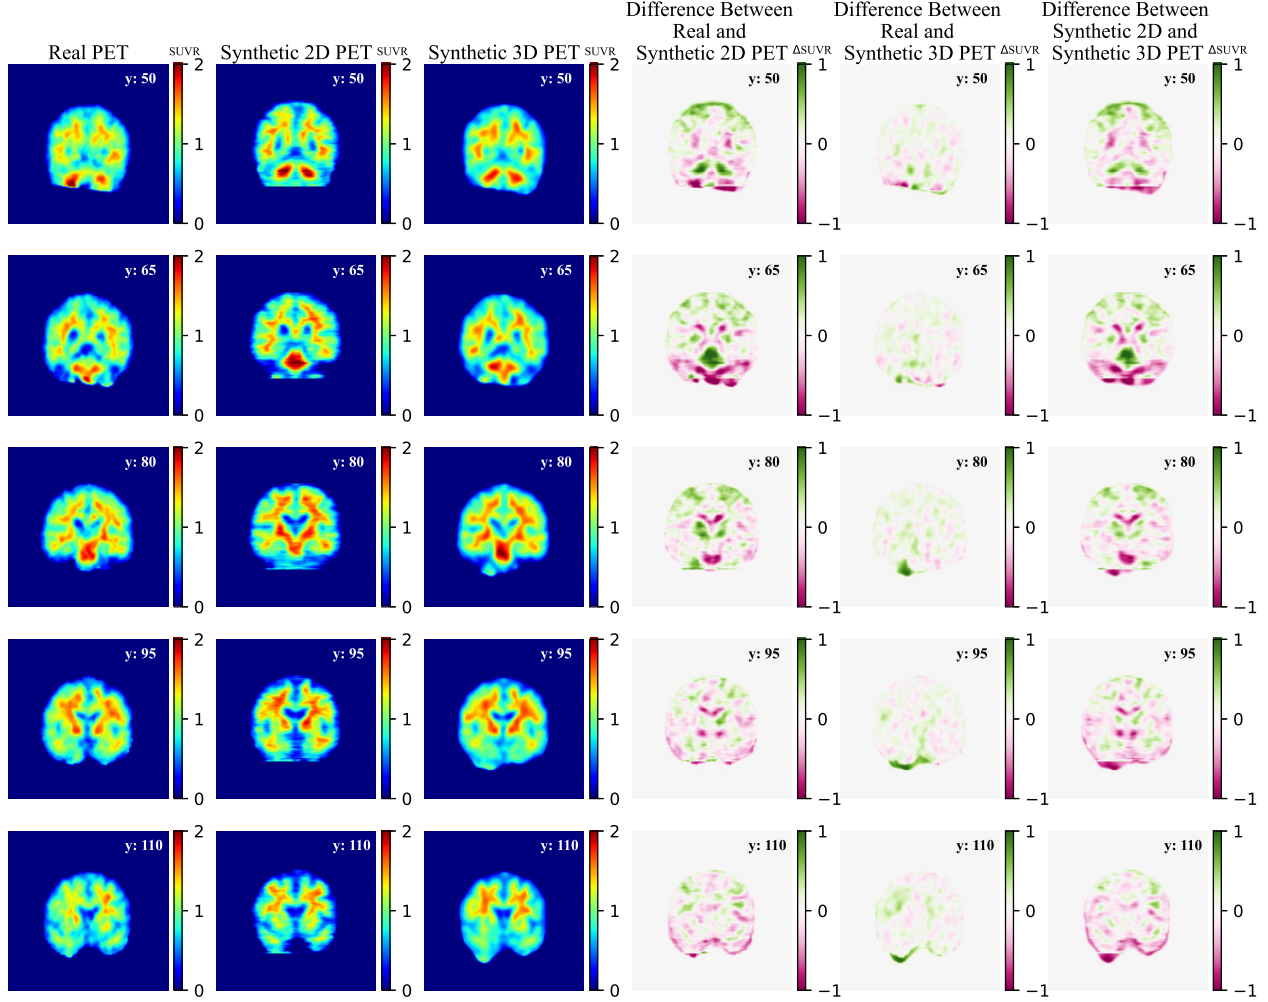
Figure 5.2: Comparison between two-dimensional and three-dimensional models in the coronal plane. Five images are presented, from left to right, the real PET image, the synthetic two-dimensional PET image, the synthetic three-dimensional PET image, the difference map between synthetic two-dimensional and real PET, the difference map between synthetic three-dimensional and real PET image, and the difference map between synthetic two-dimensional and three-dimensional PET. The two-dimensional synthetic PET images show inconsistencies in this plane as seen in its corresponding difference map (i.e., the inferior portion), while the three-dimensional synthetic PET images show better consistency.

Comparison Between Two-Dimensional and Three-Dimensional Model Sagittal View
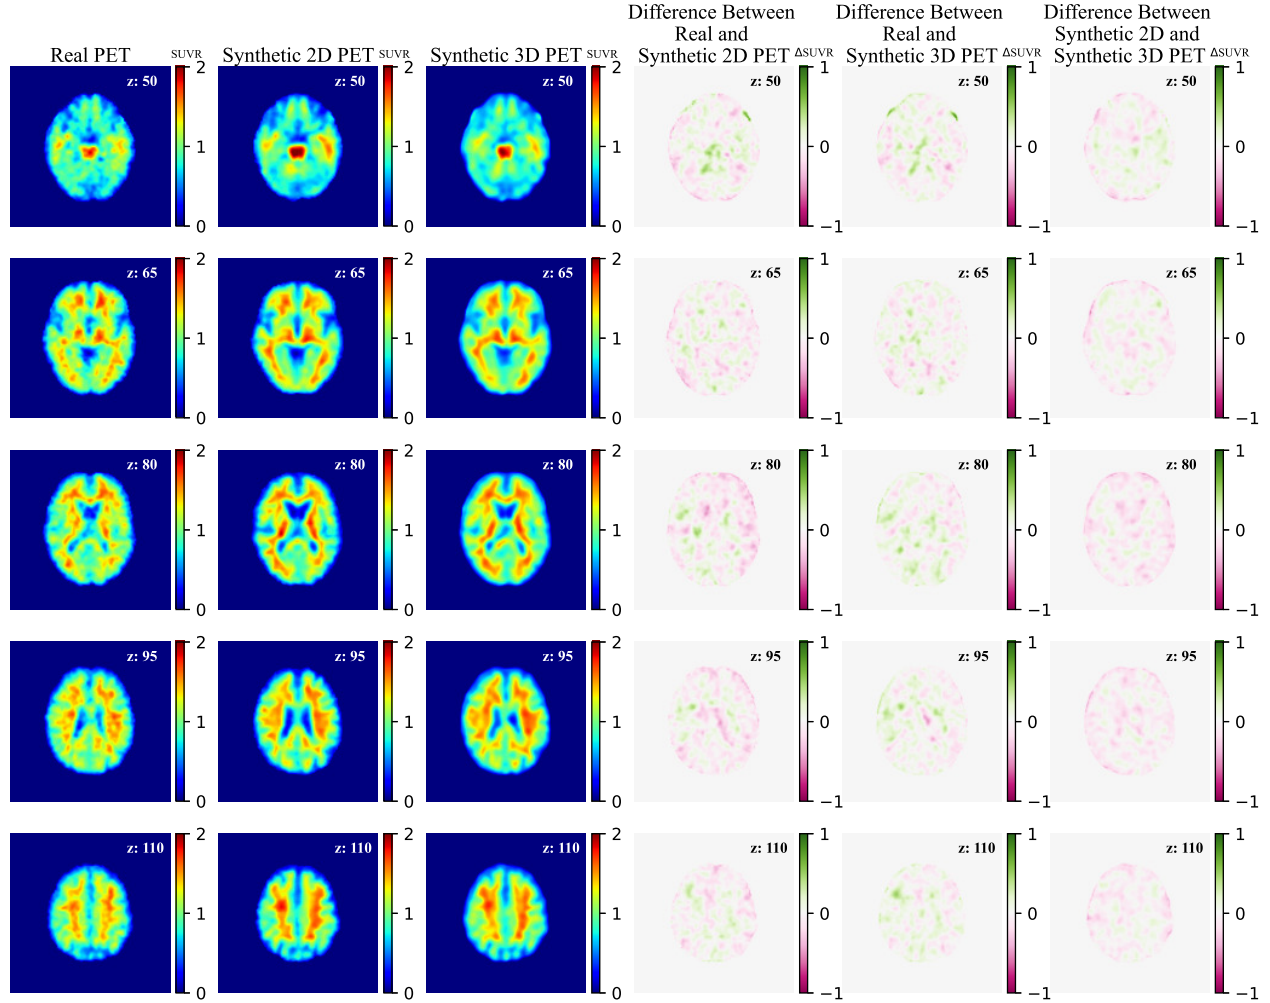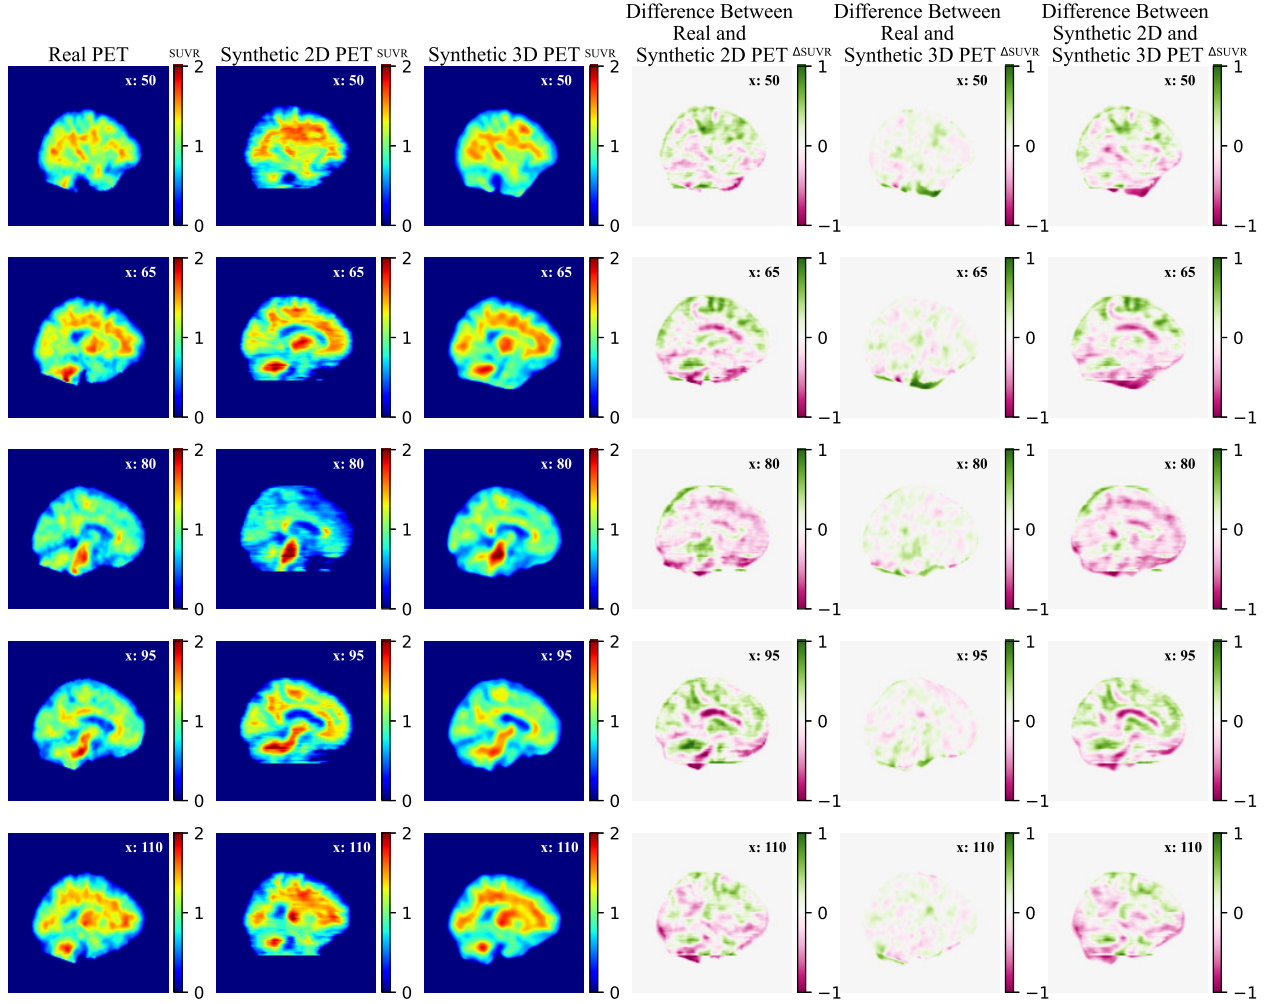

Figure 5.3: Comparison between two-dimensional and three-dimensional models in the sagittal plane. Five images are presented, from left to right, the real PET image, the synthetic two-dimensional PET image, the synthetic three-dimensional PET image, the difference map between synthetic two-dimensional and real PET, the difference map between synthetic three-dimensional and real PET image, and the difference map between synthetic two-dimensional and three-dimensional PET. As discussed, the three-dimensional synthetic PET images are closer to the true PET images in comparison with the two-dimensional synthetic PET reconstructions. It can also be seen that the inferior image shows a greater error in the two-dimensional samples. Such comparison shows the benefits of training in the three-dimensional space.

**Comparison Between Three-Dimensional and Two-Dimensional Models
on SSIM Components and PSNR Across 331 Subjects**
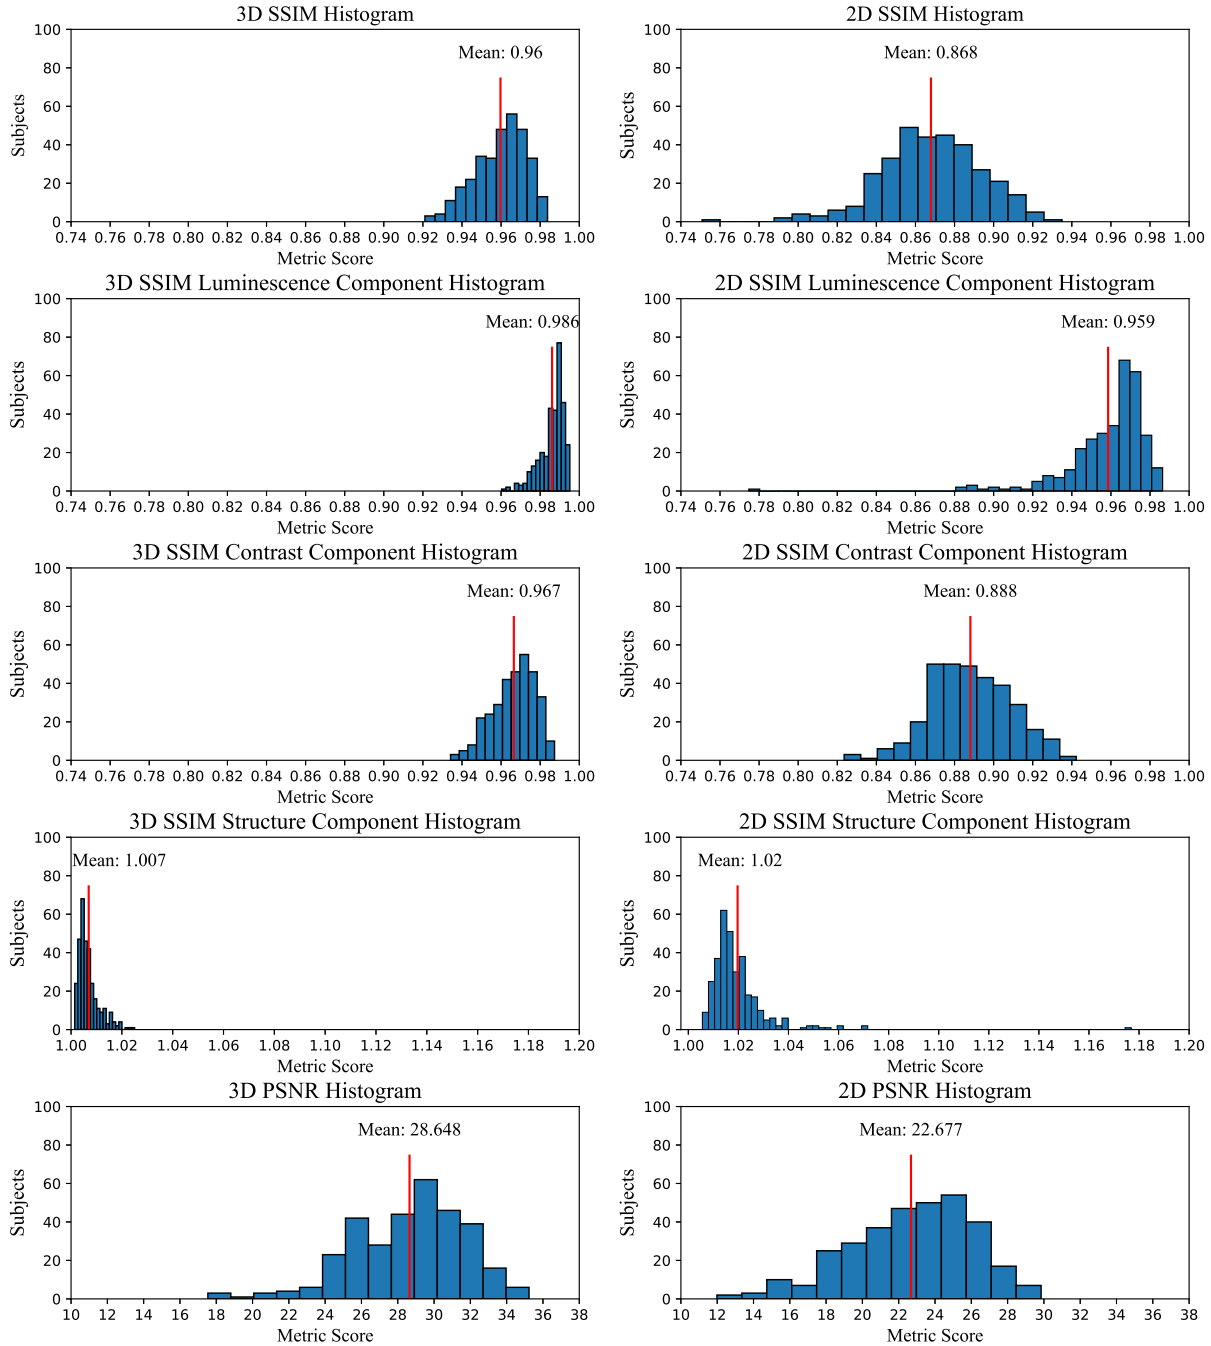
Figure 5.4: Comparison between three-dimensional and two-dimensional models across the evaluation cohort of 331 subjects. For each model structural similarity index measure (SSIM) and peak signal-to-noise ratio (PSNR) was computed, the SSIM and PSNR for the two-dimensional model were obtained from Chapter 4 results. The mean SSIM for the three-dimensional model scored 10% higher than the two-dimensional one while the mean PSR for the three-dimensional mode scores 26.3% higher as well. Such results show that performing image translation from T1w MRI to amyloid-beta PET image is more effective when training in three-dimensional space.

5.5 Discussion

In this chapter, a comparison between the two model architectures proposed in this thesis was performed. The three-dimensional model was trained using the same stratified cohort as in the two-dimensional case presented in Chapter 3 to perform a rigorous comparison. The reported metrics indicate that the three-dimensional approach improved work presented in Chapter 3, an encouraging result as it means that the changes in methodology discussed in Chapter 4 were appropriate.

While a strength of two-dimensional model is that it can be trained with more observations than the three-dimensional approach, this analysis demonstrated that training with slices might not be the most optimal approach. The inconsistency problem across the slices could be addressed by adding a method that indicates to the network that each set of 151 slices belongs to one single subject. It is also unknown if changing the batch size to 151 might also help to improve the consistency.

For future work for comparison analyses, it could be meaningful to train a two-dimensional model using the same architecture improvements proposed in the three-dimensional approach and analyze the impact of the given architecture in two-dimensional space before comparing again with the three-dimensional one.

5.6 Chapter Conclusion

This chapter demonstrates the strengths of using three-dimensional models to generate synthetic amyloid-beta PET images from structural MRI over the two-dimensional model approach, this analysis is important to understand the effect of each model and make sure that the changes in architecture positively impact the results.

Chapter 6

Summary

In this thesis, I have presented 2D and 3D pipelines that can translate T1w structural MRI into amyloid-beta PET images. The work shows that generating expensive, more-invasive, and limited accessible amyloid-beta PET images from cheaper, less-invasive, and more widely available MRI is feasible. The results demonstrate that the models were able to generate high-quality synthetic images that resemble the real images based on the reported experienced reviewers, SSIM, and PSNR metrics.

The motivation for using image translation to estimate amyloid-beta PET images from structural MRI is driven by the limitations and availability of PET imaging relative to MRI. This thesis supports the hypothesis that it is possible to synthesize amyloid-beta PiB PET images from T1w MRI, showing that it was achieved.

In this final chapter, I summarize the key findings presented in this thesis, put the findings in context with the current literature, and discuss limitations of the work overall. Additionally, I discuss possible improvements and future avenues that could extend this line of research.

6.1 Summary of Key Findings

Three aims were proposed in the introduction chapter. Aim 1 and aim 2 undertake the development of an image translation pipelines in two- and three- dimensions, respectively. The two-dimensional model aimed to perform image translation with non-normalized PET images, elimination of normalization layers, and a custom mask loss to disregard the background. The three-dimensional model advanced the work done in two-dimensions by incorporating SUVR normalization and z-score normalization to improve model stability, instance normalization layers, spectral normalization shown in, and an improved custom mask loss.

Aim 3 was to conduct a comparison of two- and three- dimensional models, where the same three-dimensional model as aim 2 is trained with the training and evaluation cohorts as in aim 1 to show that the implemented changes in

the three-dimensional implementation addressed the limitations from the two-dimensional and improved the performance. The results from aim 3 show that the three-dimensional approach outperformed the two-dimensional approach. It can be observed that the figure comparison between axial, coronal, and sagittal view show that the three-dimensional model can generate consistent synthetic amyloid-beta PET images. The similarity metrics also reported an improvement in comparison with the two-dimensional approach: the mean SSIM increased 10% and PSNR increased by 26.3%, indicating that the model was able to create higher-quality synthetic amyloid-beta PET images.

6.2 Context with Existing Literature

The field of image translation to estimate PET images from MRI ones is constantly evolving. There has been an increase in literature from the year 2021 to 2023 (the period over which I conducted this research) where researchers have explored different techniques to produce high-quality synthetic PET images. However, none of these projects has demonstrated PiB PET image estimation from structural MRI.

There has been past work that relates the use of T1w MRI to predict amyloid-beta positivity which has been discussed in Chapter 2, where one of them used a three-dimensional CNN to determine whether a subject had positive amyloid-beta using only T1w MRI images. Such work complement the results shown in this thesis, where there changes in structure, texture and signal differences in gray and white matter can provide meaningful information to predict amyloid-beta PET images. Existent work also supports that T1w and T2w MRI relates amyloid-beta PiB PET in different regions of the brain, meaning that this thesis with further work can strongly support such relationship (Yasuno et al., 2017).

In Chapter 2, two papers in this field were discussed, one aiming at estimating FDG PET from multi-contrast MRI and the other on estimating amyloid-beta florbetapir PET from MRI. These two papers can be compared with Chapters 3 and 4, respectively. The FDG PET from multi-contrast MRI work has shown promising results when incorporating multiple MRI images from the same subject to estimate three-dimensional florbetapir amyloid-beta PET; however, this work fell short in demonstrating consistency results across all the views. Chapter 3 of this thesis, on the other hand, shows that it is possible to generate a volume from the axial slices with only minor inconsistencies while reaching higher SSIM and PSNR metrics.

The estimation of three-dimensional amyloid-beta florbetapir PET from MRI is highly related to this thesis, and the authors were able to reach high SSIM and PSNR. However, the synthetic PET images in this work were qualitative instead of quantitative. Chapter 4 addressed this limitation by showing an analysis in gray and white matter SUVR estimation between the synthetic and real PET images, showing an MAE of less than 0.15 while using only a single T1w MRI to estimate amyloid-beta PET and a less complex architecture.

This thesis work adds to the advancements in image translation to estimate quantitative PET from MRI.

6.3 Key Limitations

Although the results are impressive, certain limitations should be acknowledged and explained, which will lead to suggestions for future work.

There were subjects with uncertain dementia (CDR 0.5) that were used as inputs to the network. Such subjects can alter the training process to become less stable, as they might contain high levels of amyloid-beta with a low degree of structural atrophy and changes in gray and white matter (as reflected in Figure 2.3). However, they were necessary, as without them, the training would use a lower sample size that might make the model prone to over-fitting.

In regards to model architecture, cGAN requires the PET images to be perfectly aligned with the MRI pair to properly work, this is a potential limitation that requires extensive work to ensure the quality of the co-registration. The adversarial component of the network makes the training process difficult. Furthermore, finding the equilibrium between generator and discriminator is not an easy task, and it can potentially be replaced with newer architectures. Furthermore, the current model was tested with a holdout cohort due to long training times (in the order of weeks), however, it would be meaningful to perform K-fold cross-validation. Cross-validation is a popular technique in deep learning that involves training and testing the model with different data cohorts, it shows the model robustness against different data distributions. (Bates et al., 2021)

Additionally, this work only used T1w MRI images. While this helps to make the networks lighter and less complex, the results could potentially be improved by adding more information to the network, such as cerebral blood flow (CBF) maps with arterial spin labeling (ASL) MRI images to measure the vascular integrity of the brain. Indeed, studies have shown that changes in CBF can be used to assess preclinical AD, as shown in Figure 6.1. Some of the discussed papers also made use of FLAIR imaging (Lan et al., 2021; Ouyang et al., 2023), which could also be used along T1w and ASL images to predict amyloid-beta PET images with higher similarity to the real ones and research has found a relationship between T1w and T2w with amyloid-beta accumulation represented in PiB PET (Yasuno et al., 2017).

Training only with the OASIS-3 dataset might be considered a limitation, as the model might only be representative of one single database. since the preprocessing steps for PET imaging vary among different datasets. Training only in the OASIS-3 dataset entails that the model might only be representative of one single database since the preprocessing steps for PET imaging vary among different datasets. For example, when converting from dynamic to static, the ADNI (Gaser et al., 2013) PET images can produce them by either summing the last 30 to 60 minutes (as performed in this thesis) or averaging all the scans. These two techniques are valid, however, they can generate different SUVR for the same image. It would be necessary to perform data harmonization to account for the preprocessing protocol differences. Since the dataset curated the data to avoid incorporating non-dementia related subjects, it is known that there are other conditions that can cause changes in brain structure and affect gray and white matter (i.e., stroke,
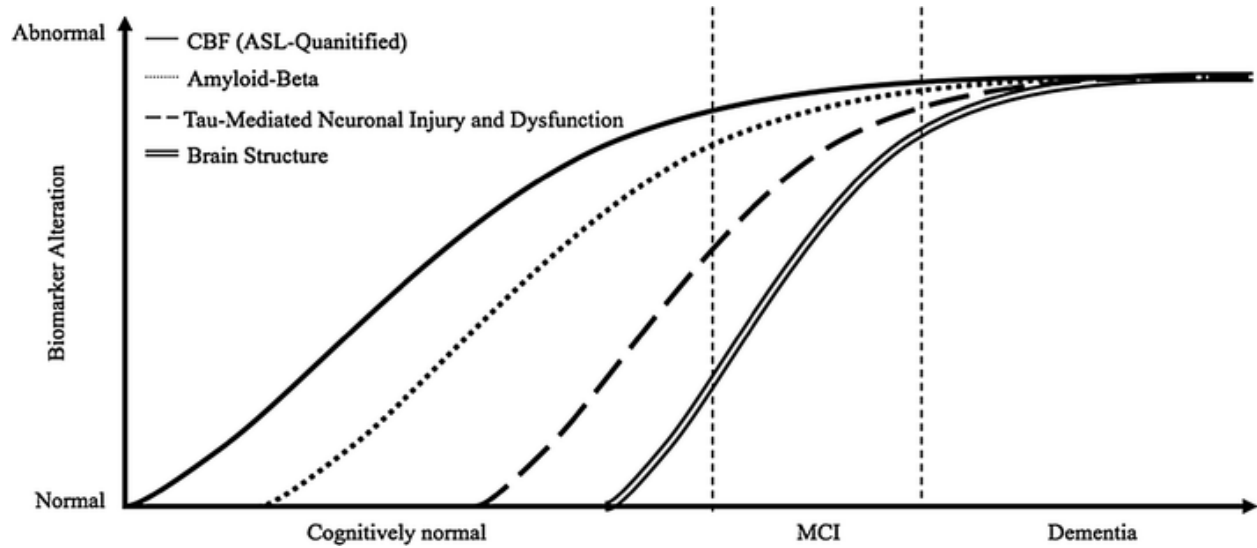
Figure 6.1: Updated Alzheimer's Disease (AD) progression curves based on different biomarkers. It can be observed that cerebral blood flow (CBF) abnormalities can precede AD development earlier than amyloid-beta and tau. Figure retrieved from (Hays et al., 2016)

traumatic brain injury, substance abuse) the developed model is only limited to a strictly controlled data.

Another limitation that was discussed in Chapter 3, section 3.6.1 is that the amyloid-beta PET images had a spatial resolution of 8 mm and it would be meaningful to explore the effect of using a higher resolution PET scan in future models since it would allow a further comparison across subcortical structures and white matter fiber tracts.

6.4 Future Work

With respect to the model architecture, it could be meaningful to explore a cGAN with a self-attention mechanism and multi-contrast MRI consistent to work done (Ouyang et al., 2023; Lan et al., 2021). However, these additions might compromise training time and require further preprocessing considerations. Unpaired training might help in this regard, although this would also require further research to avoid common problems such as: creating non-existent features (hallucinating), inpainting, and contrast inaccuracy. More advanced model schemes such as vision transformers (ViT) could also be explored, as they are known to learn based on multi-head attention, which compares each pixel individually and obtains a representation of the most significant pixels to reconstruct an image. Furthermore, integrating aspects of diffusion models, such as palette architecture (Saharia et al., 2021), might also be another useful alternative, as these have been shown to create more realistic images than cGAN in structure and contrast.

Adding more MRI image contrasts might help to improve the image synthesis, as discussed above, recent findings establish a relationship between T2w, ASL and functional MRI with amyloid-beta. Incorporating such images as part of the input of the network can allow to produce higher quality images, it will also help to evaluate the effect of using

different contrast have in the model performance. Therefore, it would be possible to investigate the contribution of each contrast type towards creating an amyloid-beta PET image.

Since there are minor discrepancies in the generate amyloid-beta PET images compared with the real ones, adding a multi-contrast input can help to reduce them, Figure 6.2 illustrates a detailed comparison of these discrepancies for the shown subjects in Chapter 4. Once more, however, adding more images will entail more computational resources making the training process slower.
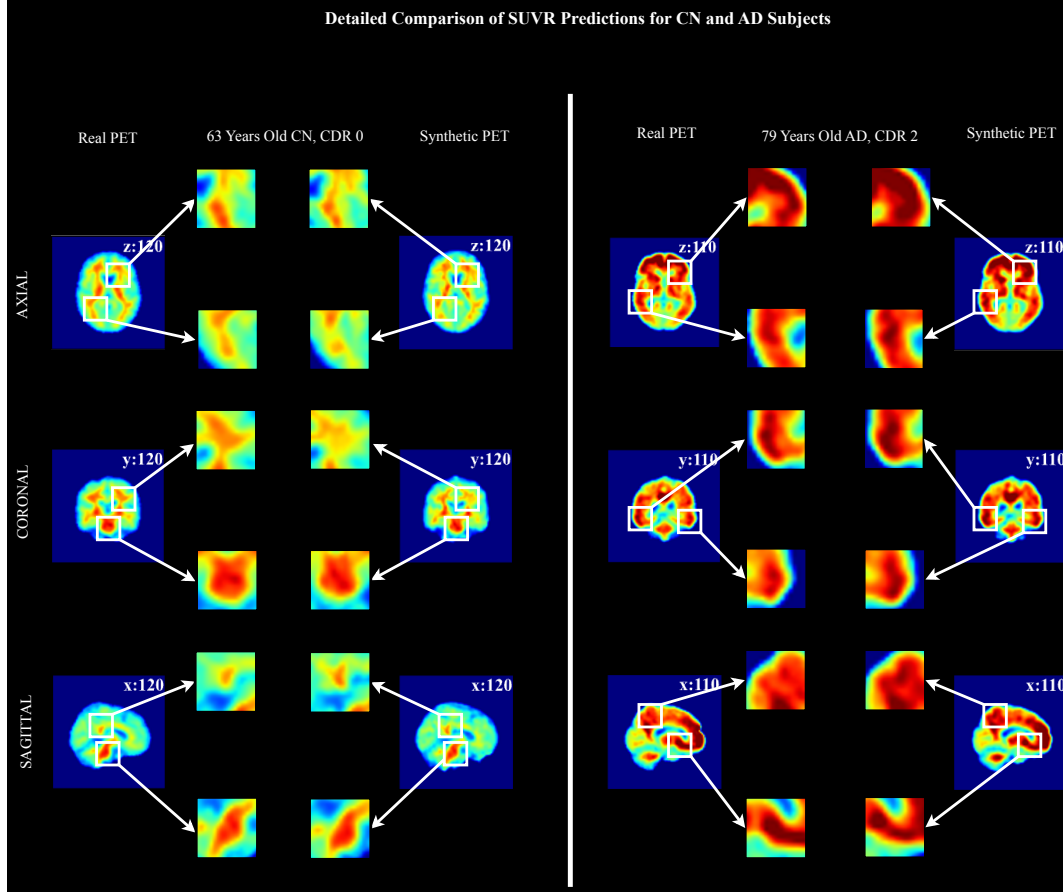
Figure 6.2: Detailed comparison in the standard uptake valuer Ratio (SUVR) prediction for cognitively normal (CN) and Alzheimer's Disease (AD) subjects. It can be seen that there are minor discrepancies between the predictions of these two subjects, where the relevant ones were highlighted.

It will be necessary to test the model with an unseen dataset. As previously discussed, if the results are not satisfactory, incorporating a set of images from ADNI and AIBL (Fowler et al., 2021) datasets may help to improve the model to generalize. However, in order to do this, it will be necessary to explore techniques that can harmonize the PET images to avoid adding images with different preprocessing protocols and unify the pipeline across all the images (Akamatsu et al., 2023).

Since tau PET was not available in OASIS-3 dataset until mid-2022, this architecture can potentially be used to

predict tau PET images from structural MRI. For instance, it would be possible to apply a technique called transfer learning to apply the trained parameters from the amyloid-beta model into this new model which can help to improve image synthesis, since it has been stated that tau and amyloid-beta are closely related. Furthermore, Figure 6.1 shows that tau can be highly correlated with structural changes in brain (i.e., brain atrophy and changes in gray and white matter) (Jack et al., 2010; Schäfer et al., 2021) therefore, it might be possible to generate highly accurate synthetic AV-1451 tau PET images scoring high similarity metrics (SSIM¿0.9)

The incorporation of paired amyloid-beta structural and structural MRI of participants with atrophy inducing disease but without amyloid-beta pathology (i.e. alcoholism, stroke, substance abuse) would be very helpful in validating the model's robustness, and if necessary, incorporate such studies to account for such events.

Finally, for this work to be applied in a clinical setting, rigorous validation is needed to seek regulatory approval.

6.5   Conclusion

I can conclude that image translation can synthesize amyloid-beta PET images from structural MRI with a high degree of similarity to the true PET images for cognitively normal and AD subjects.

Bibliography

Abdulla, S. (2021). T1, t2 and pd weighted imaging.

Addeh, A., Vega, F., Medi, P. R., Williams, R. J., Pike, G. B., and MacDonald, M. E. (2023). Direct machine learning reconstruction of respiratory variation waveforms from resting state fmri data in a pediatric population. *NeuroImage*, 269:119904.

Akamatsu, G., Tsutsui, Y., Daisaki, H., Mitsumoto, K., Baba, S., and Sasaki, M. (2023). A review of harmonization strategies for quantitative pet. *Annals of nuclear medicine*, 37(2):71–88.

Albawi, S., Mohammed, T. A., and Al-Zawi, S. (2017). Understanding of a convolutional neural network. In *2017 International Conference on Engineering and Technology (ICET)*, pages 1–6.

Alzubaidi, L., Zhang, J., Humaidi, A. J., Al-dujaili, A., Duan, Y., Al-Shamma, O., Santamaría, J., Fadhel, M., Al-Amidie, M., and Farhan, L. (2021). Review of deep learning: concepts, cnn architectures, challenges, applications, future directions. *Journal of Big Data*, 8:53.

Armanious, K., Jiang, C., Abdulatif, S., Küstner, T., Gatidis, S., and Yang, B. (2019). Unsupervised medical image translation using cycle-medgan. In *2019 27th European Signal Processing Conference (EUSIPCO)*, pages 1–5.

Armanious, K., Jiang, C., Fischer, M., et al. (2018). Medgan: Medical image translation using gans. *arXiv*. arXiv:1806.06397.

Bates, S., Hastie, T., and Tibshirani, R. (2021). Cross-validation: what does it estimate and how well does it do it? *arXiv e-prints*, page arXiv:2104.00673.

Blinkouskaya, Y. and Weickenmeier, J. (2021). Brain shape changes associated with cerebral atrophy in healthy aging and alzheimer's disease. *Frontiers in Mechanical Engineering*, 7.

Bloom, G. S. (2014). Amyloid-beta and tau: The trigger and bullet in alzheimer disease pathogenesis. *JAMA Neurology*, 71(4):505.

Brix, G., Nekolla, E., Nosske, D., et al. (2009). Risks and safety aspects related to pet/mr examinations. *European Journal of Nuclear Medicine and Molecular Imaging*, 36(Suppl 1):131–138.

Broderick, P. A., Rahni, D. N., and Kolodny, E. H. (2007). *Bioimaging in Neurodegeneration*, chapter II, pages 69–107. Springer Science & Business Media.

Busche, M. A. and Hyman, B. T. (2020). Synergy between amyloid-beta and tau in alzheimer's disease. *Nature Neuroscience*, 23(10).

Buxton, R. (2009). *Introduction to functional magnetic resonance imaging: principles and techniques*. Cambridge, second edition.

Bydder, G., Hajnal, J., and Young, I. (1998). Mri: Use of the inversion recovery pulse sequence. *Clinical Radiology*, 53(3):159–176.

Camus, V., Payoux, P., Barré, L., et al. (2012). Using pet with 18f-av-45 (florbetapir) to quantify brain amyloid load in a clinical environment. *European Journal of Nuclear Medicine and Molecular Imaging*, 39(4):621–631.

Canada, D. (2023). Digital research alliance of canada.

Chapman, R., Mapstone, M., Porsteinsson, A., et al. (2010). Diagnosis of alzheimer's disease using neuropsychological testing improved by multivariate analyses. *J Clin Exp Neuropsychol*, 32(8):793–808.

Chattopadhyay, T., Ozarkar, S. S., Buwa, K., Thomopoulos, S. I., Thompson, P. M., and Initiative, A. D. N. (2023). Predicting brain amyloid positivity from t1 weighted brain mri and mri-derived gray matter, white matter and csf maps using transfer learning on 3d cnns. *bioRxiv*.

Chen, G.-f., Xu, T.-h., Yan, Y., et al. (2017). Amyloid beta: structure, biology and structure-based therapeutic development. *Acta Pharmacologica Sinica*, 38(9):1205–1235.

Chetelat, G., Villemagne, V., Bourgeat, P., et al. (2010). Relationship between atrophy and beta-amyloid deposition in alzheimer disease. *Ann Neurol*, 67(3):317–324.

Cummings, J. (2011). Biomarkers in alzheimer's disease drug development. *Alzheimer's & Dementia*, 7:e13–e44.

Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv*. arXiv:2010.11929.

Drug, C. and Agency, H. T. (2022). The implementation considerations of pet-ct.

Dulińska-Litewka, J., Łazarczyk, A., Hałubiec, P., Szafrański, O., Karnas, K., and Karewicz, A. (2019). Superparamagnetic iron oxide nanoparticles—current and prospective medical applications. *Materials*, 12:617.

Ehman, E., Johnson, G., Villanueva-Meyer, J., Cha, S., Leynes, A. P., Larson, P., and Hope, T. (2017). Pet/mri: Where might it replace pet/ct? *Journal of Magnetic Resonance Imaging*, 46(5):1247–1262.

Eppig, J. et al. (2020). Neuropsychological contributions to the diagnosis of mild cognitive impairment associated with alzheimer's disease. In Libon, D. J. et al., editors, *Vascular Disease, Alzheimer's Disease, and Mild Cognitive Impairment: Advancing an Integrated Approach*. Oxford Academic, New York. Online edn, accessed 9 Nov. 2023.

Fonov, V., Evans, A., McKinstry, R., Almli, C., and Collins, D. (2009). Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage*, 47:S102.

Forsberg, A., Engler, H., Almkvist, O., et al. (2008). Pet imaging of amyloid deposition in patients with mild cognitive impairment. *Neurobiol Aging*, 29(10):1456–1465.

Fowler, C., Rainey-Smith, S., Bird, S., and Bomke, J. (2021). Fifteen years of the australian imaging, biomarkers and lifestyle (aibl) study: Progress and observations from 2,359 older adults spanning the spectrum from cognitive normality to alzheimer's disease. *Journal of Alzheimer's Disease Reports*, 5(1):443–468.

Fox, N. C. and Schott, J. M. (2004). Imaging cerebral atrophy: normal ageing to alzheimer's disease. *Lancet (London, England)*, 363(9406):392–394.

Gandy, S. (2005). The role of cerebral amyloid beta accumulation in common forms of alzheimer disease. *The Journal of clinical investigation*, 115(5):1121–1129.

Gaser, C., Franke, K., Klöppel, S., Koutsouleris, N., Sauer, H., and Initiative, A. D. N. (2013). Brainage in mild cognitive impaired patients: predicting the conversion to alzheimer's disease. *PloS one*, 8(6):e67346.

Ghiso, J. and Frangione, B. (2002). Amyloidosis and alzheimer's disease. *Advanced Drug Delivery Reviews*, 54(12):1539–1551. Current treatments and therapeutic targets in Alzheimer's disease.

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. http://www.deeplearningbook.org.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al. (2014). Generative adversarial networks. *arXiv*. arXiv:1406.2661.

Güven, S. and Talu, M. (2023). Brain mri high resolution image creation and segmentation with the new gan method. *Biomedical Signal Processing and Control*, 80:104246.

Hampel, H., Hardy, J., Blennow, K., et al. (2021). The amyloid-β pathway in alzheimer's disease. *Molecular Psychiatry*, 26(10):5481–5503.

Hansson, O., Lehmann, S., Otto, M., Zetterberg, H., and Lewczuk, P. (2019). Advantages and disadvantages of the use of the csf amyloid beta (abeta) 42/40 ratio in the diagnosis of alzheimer's disease. *Alzheimers Res Ther*, 11(1):34.

Hays, C. C., Zlatar, Z., and Wierenga, C. (2016). The utility of cerebral blood flow as a biomarker of preclinical alzheimer's disease. *Cellular and Molecular Neurobiology*, 36:167–179.

Heneka, M. T., Carson, M. J., El Khoury, J., Landreth, G. E., Brosseron, F., Feinstein, D. L., Jacobs, A. H., Wyss-Coray, T., Vitorica, J., Ransohoff, R. M., Herrup, K., Frautschy, S. A., Finsen, B., Brown, G. C., Verkhratsky, A., Yamanaka, K., Koistinaho, J., Latz, E., Halle, A., Petzold, G. C., et al. (2015). Neuroinflammation in alzheimer's disease. *The Lancet. Neurology*, 14(4):388–405.

Horé, A. and Ziou, D. (2010). Image quality metrics: Psnr vs. ssim. In *2010 20th International Conference on Pattern Recognition*, pages 2366–2369.

Hudgins, P., Fountain, A., Chapman, P., and Shah, L. (2017). Difficult lumbar puncture: Pitfalls and tips from the trenches. *American Journal of Neuroradiology*, 38(7).

Hussein, R., Shin, D., Zhao, M., Guo, J., Davidzon, G., Moseley, M., and Zaharchuk, G. (2022). Brain MRI-to-PET Synthesis using 3D Convolutional Attention Networks. *arXiv e-prints*, page arXiv:2211.12082.

Isola, P., Zhu, J., Zhou, T., and Efros, A. (2017). Image-to-image translation with conditional adversarial networks. In *Proc Cvpr Ieee*, pages 5967–5976.

Jack, C. R., Knopman, D. S., Jagust, W. J., Shaw, L. M., Aisen, P. S., Weiner, M. W., Petersen, R. C., and Trojanowski, J. Q. (2010). Hypothetical model of dynamic biomarkers of the alzheimer's pathological cascade. *The Lancet Neurology*, 9(1):119–128.

Jack, C. R., J., Lowe, V. J., Weigand, S. D., Wiste, H. J., Senjem, M. L., Knopman, D. S., Shiung, M. M., Gunter, J. L., Boeve, B. F., Kemp, B. J., Weiner, M., and Petersen, R. C. (2009). Serial pib and mri in normal, mild cognitive impairment and alzheimer's disease: implications for sequence of pathological events in alzheimer's disease. *Brain : a journal of neurology*, 132(Pt 5):1355–1365.

Johnson, K., Fox, N., Sperling, R., and Klunk, W. (2012). Brain imaging in alzheimer disease. *Cold Spring Harb Perspect Med*, 2(4):a006213.

Johnson, P. and Drangova, M. (2019). Conditional generative adversarial network for 3d rigid-body motion correction in mri. *Magn Reson Med*, 82(3):901–910.

Khan, T. (2016). *Biomarkers in Alzheimer's Disease Diagnosis*. Elsevier.

Khan, T. and Alkon, D. (2015). Alzheimer's disease cerebrospinal fluid and neuroimaging biomarkers: Diagnostic accuracy and relationship to drug efficacy. *J Alzheimers Dis*, 46(4):817–836.

Kim, J. P., Kim, J., Jang, H., Kim, J., Kang, S. H., Kim, J. S., Lee, J., Na, D., Kim, H. J., Seo, S., and Park, H. (2021). Predicting amyloid positivity in patients with mild cognitive impairment using a radiomics approach. *Scientific Reports*, 11:6954.

Kumar, A., Sidhu, J., Goyal, A., and Tsao, J. (2022). *Alzheimer Disease*. StatPearls, Treasure Island (FL).

Kustner, T., Armanious, K., Yang, J., Yang, B., Schick, F., and Gatidis, S. (2019). Retrospective correction of motion-affected mr images using deep learning frameworks. *Magn Reson Med*, 82(4):1527–1540.

LaMontagne, P., Benzinger, T., Morris, J., et al. (2019). Oasis-3: Longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. *medRxiv*.

Lan, H., Initiative, A. D. N., Toga, A. W., and Sepehrband, F. (2021). Three-dimensional self-attention conditional gan with spectral normalization for multimodal neuroimaging synthesis. *Magnetic resonance in medicine*, 86(3):1718–1733.

Lee, S., Lee, H., and Kim, K. W. (2020). Magnetic resonance imaging texture predicts progression to dementia due to alzheimer disease earlier than hippocampal volume. *Journal of Psychiatry & Neuroscience : JPN*, 45(1):7–14.

Lee, Y., Youn, H., Jeong, H., et al. (2021). Cost-effectiveness of using amyloid positron emission tomography in individuals with mild cognitive impairment. *Cost Effectiveness and Resource Allocation*, 19:50.

Liang, Z.-P. and Lauterbur, P. C. (1999). *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. IEEE Press.

Liu, Y.-H., Wang, Y., Xiang, Y., Zhou, H., Giunta, B., Mañucat-Tan, N., Tan, J., Zhou, X.-F., and Wang, Y. (2015). Clearance of amyloid-beta in alzheimer's disease: Shifting the action site from center to periphery. *Molecular Neurobiology*, 51:1–7.

Luijten, S. P. R., Compagne, K., van Es, A. V., Roos, Y., Majoie, C., van Oostenbrugge, R. V., van Zwam, W. V., Dippel, D., Wolters, F., van der Lugt, A., and Bos, D. (2022). Brain atrophy and endovascular treatment effect in acute ischemic stroke: a secondary analysis of the mr clean trial. *International Journal of Stroke*, 17(8):881–888.

Lv, J., Wang, C., and Yang, G. (2021). Pic-gan: A parallel imaging coupled generative adversarial network for accelerated multi-channel mri reconstruction. *Diagnostics*, 11.

MacDonald, M. E. and Pike, G. (2021). Mri of healthy brain aging: A review. *NMR in Biomedicine*, 34:e4564.

MacLeod, R., Hillert, E. K., Cameron, R. T., and Baillie, G. S. (2015). The role and therapeutic targeting of α-, β-and γ-secretase in alzheimer's disease. *Future science OA*, 1(3):FSO11.

Manning, C. and Ducharme, J. (2010). *Chapter 6 - Dementia Syndromes in the Older Adult*, pages 155–178. Academic Press, San Diego.

Mantzavinos, V. and Alexiou, A. (2017). Biomarkers for alzheimer's disease diagnosis. *Current Alzheimer Research*, 14(11):1149–1154.

Marti-Climent, J., Prieto, E., Moran, V., et al. (2017). Effective dose estimation for oncological and neurological pet/ct procedures. *EJNMMI Research*, 7(1):37.

Martín-González, C., Romero-Acevedo, L., Fernández-Rodríguez, C., Medina-Vega, L., García-Rodríguez, A., Ortega-Toledo, P., González-Navarrete, L., Vera-Delgado, V., and González-Reimers, E. (2020). Brain-derived neurotrophic factor among patients with alcoholism. *CNS Spectrums*, 26(4):400–405.

Martínez, M. J., Ziegler, S. I., and Beyer, T. (2008). Pet and pet/ct: Basic principles and instrumentation. In Dresel, S., editor, *PET in Oncology*, volume 170 of *Recent Results in Cancer Research*. Springer, Berlin, Heidelberg.

Mayerhoefer, M., Prosch, H., Beer, L., et al. (2020). Pet/mri versus pet/ct in oncology: A prospective single-center study of 330 examinations focusing on implications for patient management and cost considerations. *European Journal of Nuclear Medicine and Molecular Imaging*, 47:51–60.

Mendez, M. (2022). Chapter 16 - general mental status scales, rating instruments, and behavior inventories. In Mendez, M., editor, *The Mental Status Examination Handbook*, pages 181–199. Elsevier, Philadelphia.

Meyer, P.-F., McSweeney, M., Gonneaud, J., and Villeneuve, S. (2019). Ad molecular: Pet amyloid imaging across the alzheimer's disease spectrum: From disease mechanisms to prevention. *Progress in Molecular Biology and Translational Science*, 165:63–106.

Mirza, M. and Osindero, S. (2014). Conditional generative adversarial nets. *arXiv*. arXiv:1411.1784.

Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. (2018). Spectral normalization for generative adversarial networks. *arXiv*. arXiv:1802.05957.

O'Brien, R. J. and Wong, P. C. (2011). Amyloid precursor protein processing and alzheimer's disease. *Annual review of neuroscience*, 34:185–204.

Ouyang, J., Chen, K., Duarte Armindo, R., Davidzon, G., Hawk, K., Moradi, F., Rosenberg, J., Lan, E., Zhang, H., and Zaharchuk, G. (2023). Predicting fdg-pet images from multi-contrast mri using deep learning in patients with brain neoplasms. *J Magn Reson Imaging*.

Pang, Y., Lin, J., Qin, T., and Chen, Z. (2021). Image-to-image translation: Methods and applications. *arXiv*. arXiv:2101.08629.

Pini, L., Pievani, M., Bocchetta, M., Altomare, D., Bosco, P., Cavedo, E., Galluzzi, S., Marizzoni, M., and Frisoni, G. B. (2016a). Brain atrophy in alzheimer's disease and aging. *Ageing Research Reviews*, 30:25–48. Brain Imaging and Aging.

Pini, L., Pievani, M., Bocchetta, M., Altomare, D., Bosco, P., Cavedo, E., Galluzzi, S., Marizzoni, M., and Frisoni, G. B. (2016b). Brain atrophy in alzheimer's disease and aging. *Ageing Research Reviews*, 30:25–48.

Plewes, D. and Kucharczyk, W. (2012). Physics of mri: A primer. *Journal of Magnetic Resonance Imaging*, 35:1038–1054.

Price, J., Klunk, W., Lopresti, B., and et al. (2005). Kinetic modeling of amyloid binding in humans using pet imaging and pittsburgh compound-b. *Journal of Cerebral Blood Flow & Metabolism*, 25(11):1528–1547.

Ran, M., Hu, J., Chen, Y., et al. (2019). Denoising of 3d magnetic resonance images using a residual encoder-decoder wasserstein generative adversarial network. *Med Image Anal*, 55:165–180.

Reimand, J., Groot, C., Teunissen, C. E., Windhorst, A. D., Boellaard, R., Barkhof, F., Nazarenko, S., van der Flier, W. M., van Berckel, B. N., Scheltens, P., Ossenkoppele, R., and Bouwman, F. (2020). Why is amyloid-beta pet requested after performing csf biomarkers? *Journal of Alzheimer's Disease*, 73(2).

Ritchie, C., Smailagic, N., Noel-Storr, A. H., Ukoumunne, O., Ladds, E. C., and Martin, S. (2017). Csf tau and the csf tau/abeta ratio for the diagnosis of alzheimer's disease dementia and other dementias in people with mild cognitive impairment (mci). *Cochrane Database of Systematic Reviews*, 2017(3).

Ronneberger, O., Fischer, P., and Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. *arXiv e-prints*, page arXiv:1505.04597.

Sadigh-Eteghad, S., Sabermarouf, B., Majdi, A., Talebi, M., Farhoudi, M., and Mahmoudi, J. (2015). Amyloid-beta: A crucial factor in alzheimer's disease. *Med Princ Pract*, 24(1):1–10.

Saharia, C., Chan, W., Chang, H., Lee, C. A., Ho, J., Salimans, T., Fleet, D. J., and Norouzi, M. (2021). Palette: Image-to-Image Diffusion Models. *arXiv e-prints*, page arXiv:2111.05826.

Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. (2016). Improved Techniques for Training GANs. *arXiv e-prints*, page arXiv:1606.03498.

Scheltens, P., De Strooper, B., Kivipelto, M., Holstege, H., Chételat, G., Teunissen, C. E., Cummings, J., and van der Flier, W. M. (2021a). Alzheimer's disease. *The Lancet*, 397(10284):1577–1590.

Scheltens, P. et al. (2021b). Alzheimer's disease. *The Lancet*, 397(10284):1577–1590.

Schäfer, A., Chaggar, P., Thompson, T. B., Goriely, A., and Kuhl, E. (2021). Predicting brain atrophy from tau pathology: a summary of clinical findings and their translation into personalized models. *Brain Multiphysics*, 2:100039.

Shah, J., Gao, F., Li, B., et al. (2022). Deep residual inception encoder-decoder network for amyloid pet harmonization. *Alzheimer's % Dementia*, 18(12):2448–2457.

Shaw, L. M., Arias, J., Blennow, K., Galasko, D., Molinuevo, J. L., Salloway, S., Schindler, S., Carrillo, M. C., Hendrix, J. A., Ross, A., Illes, J., Ramus, C., and Fifer, S. (2018). Appropriate use criteria for lumbar puncture and cerebrospinal fluid testing in the diagnosis of alzheimer's disease. *Alzheimer's & Dementia*, 14(11):1505–1521.

Shigemizu, D., Akiyama, S., Mitsumori, R., Niida, S., and Ozaki, K. (2022). Identification of potential blood biomarkers for early diagnosis of alzheimer's disease through immune landscape analysis. *npj Aging*, 8:15.

Sikka, A., Skand, Virk, J., and Bathula, D. (2021). Mri to pet cross-modality translation using globally and locally aware gan (gla-gan) for multi-modal diagnosis of alzheimer's disease. *arXiv*. arXiv:2108.02160.

Soria Lopez, J. A., González, H. M., and Léger, G. C. (2019). Chapter 13 - alzheimer's disease. In Dekosky, S. T. and Asthana, S., editors, *Geriatric Neurology*, volume 167 of *Handbook of Clinical Neurology*, pages 231–255. Elsevier.

Stelzmann, R., Schnitzlein, H. N., and Murtagh, F. R. (1995). An english translation of alzheimer's 1907 paper, "Über eine eigenartige erkankung der hirnrinde". *Clinical Anatomy*, 8:429–431.

Sun, H., Jiang, Y., Yuan, J., Wang, H., Liang, D., Fan, W., Hu, Z., and Zhang, N. (2022). High-quality pet image synthesis from ultra-low-dose pet/mri using bi-task deep learning. *Quantitative Imaging in Medicine and Surgery*, 12(12):5326–5342.

Suppiah, S., Didier, M. A., and Vinjamuri, S. (2019). The who, when, why, and how of pet amyloid imaging in management of alzheimer's disease-review of literature and interesting images. *Diagnostics (Basel, Switzerland)*, 9(2):65.

Tanzi, R. E. and Bertram, L. (2005). Twenty years of the alzheimer's disease amyloid hypothesis: A genetic perspective. *Cell*, 120(4):545–555.

Thompson, P. M., Hayashi, K. M., de Zubicaray, G., Janke, A. L., Rose, S. E., Semple, J., Herman, D., Hong, M. S., Dittmer, S. S., Doddrell, D. M., and Toga, A. W. (2003). Dynamics of gray matter loss in alzheimer's disease. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 23(3):994–1005.

Tikkakoski, T. (2011). Advanced magnetic resonance imaging techniques in the diagnostics of alzheimer's disease: how much imaging is enough? *Acta Radiologica*, 52(2):127.

Tosun, D., Schuff, N., Mathis, C., Jagust, W., Weiner, M., and (ADNI), A. D. N. I. (2011). Spatial patterns of brain amyloid-beta burden and atrophy rate associations in mild cognitive impairment. *Brain*, 134(Pt 4):1077–1088.

Ulyanov, D., Vedaldi, A., and Lempitsky, V. (2016). Instance Normalization: The Missing Ingredient for Fast Stylization. *arXiv e-prints*, page arXiv:1607.08022.

Valentin, J. (2005). Protecting people against radiation exposure in the event of a radiological attack. *Annals of the ICRP*, 35(1):iii–iv.

Vega, F., Addeh, A., Ganesh, A., Smith, E., and MacDonald, M. (2023). Image translation for estimating two-dimensional axial amyloid-beta pet from structural mri. *Journal of Magnetic Resonance Imaging*.

Vega, F., Addeh, A., and MacDonald, M. E. (2023a). Amyloid-Beta Axial Plane PET Synthesis from Structural MRI: An Image Translation Approach for Screening Alzheimer's Disease. page arXiv:2309.00569.

Vega, F., Addeh, A., and MacDonald, M. E. (2023b). Amyloid-beta pet synthesis from structural mri: A potential alternative method for alzheimer's disease screening.

Vega, F., Addeh, A., and MacDonald, M. E. (2023c). Denoising Simulated Low-Field MRI (70mT) using Denoising Autoencoders (DAE) and Cycle-Consistent Generative Adversarial Networks (Cycle-GAN). *arXiv e-prints*, page arXiv:2307.06338.

Vemuri, P. and Jack, C. R. (2010). Role of structural mri in alzheimer's disease. *Alzheimer's Research & Therapy*, 2:23.

Wang, J., He, L., Zheng, H., and Lu, Z.-L. (2014). Optimizing the magnetization-prepared rapid gradient-echo (mp-rage) sequence. *PLOS ONE*, 9(5):e96899.

Wang, Y.-J., Zhou, H.-D., and Zhou, X.-F. (2006). Clearance of amyloid-beta in alzheimer's disease: progress, problems and perspectives. *Drug Discovery Today*, 11(19):931–938.

Whitwell, J. L. (2010). Progression of atrophy in alzheimer's disease and related disorders. *Neurotoxicity research*, 18(3-4):339–346.

Wimo, A., Seeher, K., Cataldi, R., Cyhlarova, E., Dielemann, J. L., Frisell, O., Guerchet, M., Jonsson, L., Malaha, A. K., and Nichols, Emma, e. a. (2023). The worldwide costs of dementia in 2019. *Alzheimer's & Dementia*.

Wittenberg, R., Knapp, M., Karagiannidou, M., Dickson, J., and Schott, J. (2019). Economic impacts of introducing diagnostics for mild cognitive impairment alzheimer's disease patients. *Alzheimer's % Dementia: New York*, 5:382–387.

Wong, W. (2020). Economic burden of alzheimer disease and managed care considerations. *American Journal of Managed Care*, 26(8 Suppl):S177–S183.

Yamashita, R., Nishio, M., Do, R., and Togashi, K. (2018). Convolutional neural networks: an overview and application in radiology. *Insights into Imaging*, 9(4):611–629.

Yamin, G. and Teplow, D. (2017). Pittsburgh compound-b (pib) binds amyloid beta-protein protofibrils. *Journal of Neurochemistry*, 140(2):210–215.

Yang, G., Yu, S., Dong, H., et al. (2018). Dagan: Deep de-aliasing generative adversarial networks for fast compressed sensing mri reconstruction. *IEEE Trans Med Imaging*, 37(6):1310–1321.

Yasuno, F., Kazui, H., Morita, N., Kajimoto, K., Ihara, M., Taguchi, A., Yamamoto, A., Matsuoka, K., Takahashi, M., Nakagawara, J., Iida, H., Kishimoto, T., and Nagatsuka, K. (2017). Use of t1-weighted/t2-weighted magnetic resonance ratio to elucidate changes due to amyloid-beta accumulation in cognitively normal subjects. *NeuroImage: Clinical*, 13:209–214.

Zhang, J., He, X., Qing, L., Gao, F., and Wang, B. (2022). Bpgan: Brain pet synthesis from mri using generative adversarial network for multi-modal alzheimer's disease diagnosis. *Comput Methods Programs Biomed*, 217.

Zhou, W., Bovik, A., Sheikh, H., and Simoncelli, E. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612.

Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. (2017). Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. *arXiv e-prints*, page arXiv:1703.10593.