

Are Human Brains Sexually Dimorphic?

by

Stefan Ungurean

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
BACHELOR OF APPLIED SCIENCE

in the School

of

Engineering Science

© Stefan Ungurean 2022

SIMON FRASER UNIVERSITY

April 2022

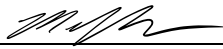
All rights reserved. This work may not be reproduced in whole or in part, by photocopy or other means, without the permission of the author. The Extraction Lens technology enclosed herein is the property of Simon Fraser University

APPROVAL

Name: Stefan Ungurean

Degree: BAsC, Biomedical Engineering

Title of Thesis: Are Human Brains Sexually Dimorphic?



Michael Sjoerdsma,
Acting Director, School of Engineering Science

Examining Committee:



Dr. Mirza Faisal Beg, P.Eng
Professor, School of Engineering Science



Dr. Da Ma, Assistant Professor,
Wake Forest School of Medicine



Dr. Hyunwoo Lee, Division of Neurology,
Department of Medicine, UBC

Date Approved:

April 28, 2022

Abstract

Machine learning and artificial intelligence tools are on the precipice of becoming a popular method for clinical diagnosis and disease prediction. Their ability to solve specific problems given sets of constraints are unmatched and they provide novel solutions to a plethora of different tasks. This thesis will examine a deep learning image classifiers capability of successfully classifying male and female MR images into their respective counterparts. The results of these models are visualized using a Grad-Cam to help gain insight on the sexual dimorphisms present within the human brain. The models showed sex dimorphisms exist in previously known areas like the frontal, temporal and precuneus gyri along with the cerebellum and thalamus. But other regions such as the cingulate, postcentral, and fusiform gyri illustrated differences not commonly mentioned in literature. This paper deviates away from the traditional statistical approaches of neuroimaging and analysis techniques and provides a new method to draw conclusions on individual volumes.

Acknowledgments

This project could not have been completed without the help of several important individuals. Their help in providing guidance and much-needed support was crucial in the completion of this work.

First, I would like to thank Dr. Mirza Faisal Beg, for accepting me as a part of his research team as an undergraduate student. His expertise in designing machine learning models and neurological imaging were of paramount importance to the success of my research. He was always patient and met my questions with sympathy and understanding.

Next, I would like to thank Dr Da Ma for guiding me every step of the way throughout these past few months and Dr Hyunwoo Lee for supporting and providing me with the time needed. Both Dr Ma and Dr Lee were tremendously helpful and understanding in answering my technical questions and clarifying key concepts I was unfamiliar with.

Lastly, I would like to sincerely thank my parents for their support over the last several months.

Table of Contents

Chapter 1: Introduction	1
1.1 Background	1
1.2 Related Works	1
1.3 Research Objectives	2
1.4 Thesis Significance and Organization	2
Chapter 2: Classification Problems for Brain MR Images, Network Designs and Background Information	4
2.1 Overview	4
2.2 Convolutional Neural Networks and Related Architectures	4
2.3 Network Construction	6
2.4 Grad-Cams and Visualization Methods	7
2.5 MRI Registration and AAL Atlas	8
2.6 Dataset Information	9
Chapter 3: Methodology, Results and Discussion.....	10
Part 1: Designing a deep-learning-based sex-classifier	10
3.1.2 Methods	10
3.1.3 Results	11
Raw Models	13
Affined Model	16
3.1.4 Discussion	18
Part 2: Explainable AI to explore brain-regions that determine sex-based diffeomorphism.....	20
3.2.2 Methods	20
3.2.3 Results	21
Raw Model Visualization	21
Affined Model Visualization	27
3.2.4 Discussion	33
Chapter 4: Future Works and Conclusion	35
References	37

List of Figures

Figure 2.1: Generic CNN Architecture

Figure 2.2: Res-Net Skip Connection

Figure 2.3: Overview of Res-Net Architectures

Figure 2.4 3-D Convolution and T1 Structural MRI Example

Figure 2.5 Classical CNN Architecture for MRI Sex Classification

Figure 2.6 Grad-Cam Process

Figure 2.7 Example of MRI Normalization

Figure 2.8 AAL Atlas Slices taken at Index 75 and 100 from axial, sagittal and coronal planes

Figure 3.1 UKBiobank Raw, Rigid, Affined and MNI 305 Template Image

Figure 3.2 OASIS Rigid, Affined and MNI 305 Template Images

Figure 3.3 Classical (Top) and ResNet18(Bottom) Training and Testing Loss/Accuracies

Figure 3.4 ROC Curves per Model and Testset. Blue represents the Classical Architecture and Red shows ResNet18 performance across thresholds. The green dot shows the arbitrary 0.5 classification threshold chosen.

Figure 3.5 Classical (Top) and ResNet18(Bottom) Training and Testing Loss/Accuracies

Figure 3.6 ROC Curves per Model and Testset. Blue represents the Classical Architecture and Red shows ResNet18 performance across thresholds. The green dot shows the arbitrary 0.5 classification threshold chosen.

Figure 3.7 Population-wide female saliency map generated from the raw OASIS testset using the Classical Model

Figure 3.8 Population-wide male saliency map generated from the raw OASIS testset using the Classical Model

Figure 3.9 Population-wide female saliency map generated from the raw UKB testset using the Classical Model

Figure 3.10 Population-wide male saliency map generated from the raw UKB testset using the Classical Model

Figure 3.11 Population-wide female saliency map generated from the raw OASIS testset using ResNet18

Figure 3.12 Population-wide male saliency map generated from the raw OASIS testset using ResNet18

Figure 3.13 Population-wide female saliency map generated from the raw UKB testset using ResNet18

Figure 3.14 Population-wide male saliency map generated from the raw UKB testset using ResNet18

Figure 3.15 Population-wide female saliency map generated from the affined OASIS testset using the Classical Model

Figure 3.16 Population-wide male saliency map generated from the affined OASIS testset using the Classical Model

Figure 3.17 Population-wide female saliency map generated from the affined UKB testset using the Classical Model

Figure 3.18 Population-wide male saliency map generated from the affined UKB testset using the Classical Model

Figure 3.19 Population-wide female saliency map generated from the affined OASIS testset using ResNet18

Figure 3.20 Population-wide male saliency map generated from the affined OASIS testset using ResNet18

Figure 3.21 Population-wide female saliency map generated from the affined UKB testset using ResNet18

Figure 3.22 Population-wide male saliency map generated from the affined UKB testset using ResNet18

List of Tables

Table 3.1 K-fold Cross Validation Results for Models Trained on Raw MR Images. Optimal parameters for best performing model highlighted in yellow

Table 3.2 Model Performance Statistics with Classification Threshold at 0.5 on OASIS Dataset

Table 3.3 Model Performance Statistics with Threshold at 0.5 on UKBiobank Testset

Table 3.4 K-fold Cross Validation Results for Models Trained on Affined MR Images. Optimal parameters for best performing model highlighted in yellow

Table 3.5 Model Performance Statistics with Classification Threshold at 0.5 on OASIS Dataset

Table 3.6 Model Performance Statistics with Threshold at 0.5 on UKBiobank Testset

Table 3.7 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and classical architecture for the raw OASIS test set.

Table 3.8 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and classical architecture for the raw UKB test set.

Table 3.9 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and ResNet18 for the raw OASIS test set.

Table 3.10 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and ResNet18 for the raw UKB test set.

Table 3.11 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and classical architecture for the affined OASIS test set.

Table 3.12 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and classical architecture for the affined UKB test set.

Table 3.13 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and ResNet18 for the affined OASIS test set.

Table 3.14 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and ResNet18 for the affined UKB test set

Table 3.15 List of reviewed studies and their results

Chapter 1: Introduction

1.1 Background

Sex differences within the brain have historically been a controversial and very debatable area of study; with many researchers claiming males and females share no anatomical differences other than size [1] [2]. As opposed to psychological disparities across sexes, variations between anatomical and physiological organs of the brain are argued to be non-existent or at least not significant [2]. However, there are academics that oppose this view and attempt to illustrate differences present in the underlying neurological structures [3]. Many go on to claim these discrepancies extend past structure influencing many physiological functions such as psychology, emotion, memory, vision, and hearing [3].

The sexual dimorphism of the human brain is a pillar in understanding the medicine and medical practices of the future [4]. It is important to gain a detailed picture of these differences to be able to assess neurological conditions that differ in their rates of manifestation between the sexes [5]. Diseases like Alzheimer's Dementia (AD) and Major Depressive Disorder (MDD) are more prevalent in women whereas males have a higher chance of being diagnosed with autism spectrum disorder or schizophrenia [6]. Currently, finding and administering treatment for a particular disease is reliant on the assumption that the human brain is identical to its sexual counterpart. Granted, this assumption has been accepted and employed in treatments due to the limitations in statistical power in generating substantial results [5]. Nevertheless, the ability to see the human brain as "two complex mosaics" [3] rather than one identical entity is a steppingstone to identifying and understanding complex diseases and their effect on neurological physiologies.

1.2 Related Works

To this date, some related works and studies have shown there are no existing structural and functional differences between male and female brains other than size. This is supported extensively in a paper published in 2020 in the *Neuroscience & Biobehavioral Reviews* journal [2]. However, there are others who firmly oppose this view claiming that sex dimorphisms within the brain have not been fully discovered [3].

Larry Cahill was the one of the first researchers to publish an in-depth article debunking the "myth" surrounding identical brains [3]. Cahill discusses the common misconceptions surrounding sexual dimorphisms and focuses on difference findings [3] in multiple brain organs. The article summarizes preliminary results on the disparities between male and female brains and urges future work be done by first analyzing areas that differ fundamentally from one another.

Another article published in *The Journal of Cerebral Cortex* in 2018 provides a larger study on the sex differences in the adult human brain. The goal of the study was to present a large and comprehensive analysis of sex difference in mean and variance that covered structural, diffusion and functional MRI measures [5]. The investigation found several structural and functional differences that were previously unnoticed, such as a stronger connectivity in the default mode network of females, and stronger connectivity in males between unimodal sensory and motor

cortices [5]. The paper also sees large overlaps between both sexes, but its main strength resides in the large sample size of the entire study.

Lastly, an important article was published in 2013 in the *Magnetic Resonance Imaging* journal by J. M. Goldstein *et al* [7]. The paper also investigates the structural and functional differences of the brain. More specifically, it finds sex differences in the white matter microstructure of differing organs and uncovers differing neural connectivity in the resting state functions of the corpus callosum, the anterior cingulate cortex and orbitofrontal cortex.

Although these publications have brought forth sexual dimorphisms in the human brain, they have not yet provided a methodology for accurate classification of the organs within. The development of such classification would not only be beneficial to understanding human cognition and behavior but would also aid in the detection and treatment of many neurodegenerative diseases [8]. According to a paper published in by Dr Raznahan in 2020 in the PNAS journal, categorizing and characterizing sex differences within the brain could provide vital information for how we understand biological sex and variations in disease risk and prevalence [8]. Knowing this, classification of male and female brains could hold the key, or at least the blueprint, to discovering neurological remedies that were previously believed to be non-existent.

1.3 Research Objectives

The main objective of this thesis will be to use explainable AI methods to investigate sex differences within the structures of the human brain, while ruling out individual differences in overall size, volume, and shape. The goal is to present a new method for finding structural differences between the sexes and visualizing their locations in the cranium.

To accomplish the objective, two specific aims with detailed hypothesis and questions are created to help guide the research. First, design a deep-learning-based sex classifier that can predict sex based on individual brain MR volumes. This aim hypothesizes that convolutional neural networks can correctly classify sex with high precision and accuracy. Given the plethora of classification architectures available, the thesis will analyze which models are optimal for the binary classification task and if they can predict sex with an accuracy over 80% on the test sets provided.

If the first aim is successful, the second uses explainable AI to explore brain-regions that determine sex-based diffeomorphisms. Ideally, the aim hypothesizes that Grad-Cam along with saliency maps and deep-learning visualization tools can provide novel information of region-wise sex-based differences. Additionally, these visualizations will attempt to show that brain classification is not only driven by overall size and shape. By looking at raw and affinely transformed MR scans, results can be drawn about the location of differences present when size and shape are removed as features of interest. From there, the top-ranked brain regions responsible for classification can be extracted and compared against previous literature.

1.4 Thesis Significance and Organization

The overall purpose of this research is to observe whether machine learning classification models can highlight neurological sex differences within healthy human brains. These dimorphisms can lead to a better understanding of treatments, diagnosis, and progression of many neurological and

psychological diseases. More specifically, sex differences for psychiatric disorders could be rooted in sex-specific brain behavior and organization [6]. For example, from an evolutionary standpoint, females have an advantage in risk aversion [9] [10] which could explain the activation patterns of stress systems like anxiety, major depressive disorder and post-traumatic stress disorder [6]. These activations could have sex-specific structural counterparts that manifest themselves in certain brain regions. Conversely, males tend to show better spatial abilities than females [11] and a greater affinity for aggressive behaviors and patrolling [12] which could be linked to risk-taking and social behavior disorders like autism, schizophrenia and impulsive aggression [6]. As with females, these behaviors and disorders could have structural manifestations within the brain that have not yet been uncovered. Furthermore, recent literature found dissimilarities in brain regions where genes on the sex chromosomes are disproportionately expressed [8]. This can help understanding surrounding the manifestation of neurological disease symptoms and associated risk levels. By discovering sex-based brain differences, we can learn more about our species and transform medical practices to suit every patient's needs rather than using a one-size-fits all approach.

The project is split into two parts. First, it will analyze a deep-learning classifier's ability to perform sex-based predictions on individual volumes for affined and non-affined MR images. Second, it will use explainable AI to visualize, explore and compare the classification decisions of these models. This paper will begin in Chapter 2 with a technical overview of the important aspects of the research. This section covers classification problems, machine learning models, grad-cam visualizations, and MRI registration along with details on the dataset used and architectures designed. Chapter 3 provides the methodology, results, and detailed discussion on the findings. Chapter 4 concludes with future works and challenges faced.

Chapter 2: Classification Problems for Brain MR Images, Network Designs and Background Information

2.1 Overview

Machine learning and AI is at the precipice of becoming the dominant tool in the software and technology industry due to its ever-increasing ability to solve and visualize complex problems. More specifically, over the last several years, the integration of machine learning algorithms with medicine has resulted in new methods for diagnosing and predicting disease progression. Deep learning, a popular sub section of machine learning, is constantly evolving in its predictive capabilities due to increases in computational power and access to large datasets.

In the context of medicine, images are one of the primary forms of confirming diagnosis and understanding human anatomy and physiology. For brain imaging, T1 structural MRI's are used to visualize the entire cranium and its contents. These MRI's can be visualized as a cube of voxels where a slice in the sagittal, coronal, or axial plane forms a different image. T1 MRIs provide very high-resolution images that can hold a plethora of information invisible to the naked eye, but noticeable to machine learning networks.

2.2 Convolutional Neural Networks and Related Architectures

Convolutional neural networks are a specific subsection of deep learning models where an image is passed through a set of consecutive convolutional layers to extract important features that best describe the data. These feature maps are then concatenated and fed to a traditional fully connected neural network with some activation applied to the output layer. Figure 2.1 below shows a generic convolutional neural network model.




Figure 2.1 Generic CNN Architecture [13]

These types of architectures have become extremely popular for image recognition, segmentation, and classification because their powerful generalization and feature extraction capabilities. However, as these models become larger and more complex, they become prone to suffer from the vanishing gradient problem; an issue where the partial derivatives of the error functions become increasingly small, effectively zero, as the network undergoes backpropagation hence preventing the subsequent weights from updating correctly.

As a result, many CNN architectures have been designed over the years to overcome this issue and improve performance, such as LeNet-5, AlexNet, and ResNet-18/30/50/150 [14] [15] [16] all with great classification and segmentation capabilities. However, one the best is the Residual

Neural Network developed in 2016 [16]. The ResNet architecture overcomes the classic vanishing gradient problem suffered by classical deep-layered models by implementing skip-connections shown in figure 2.2 below.




Figure 2.2 ResNet Skip Connection [16]

These skip-connections allow an alternate shortcut path for the gradient to flow through [16] hence reducing the length the gradient must travel to update the necessary weights.

The Res-Net models comes in varying depths and forms with Figure 2.3 below showing all the generic architectures of the multilayered models.

layer name	output size	18-layer	34-layer	50-layer	101-layer	152-layer
conv1	112×112	7×7, 64, stride 2				
		3×3 max pool, stride 2				
conv2_x	56×56	$\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 2$	$\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \end{bmatrix} \times 3$
conv3_x	28×28	$\begin{bmatrix} 3 \times 3, 128 \\ 3 \times 3, 128 \end{bmatrix} \times 2$	$\begin{bmatrix} 3 \times 3, 128 \\ 3 \times 3, 128 \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128 \\ 1 \times 1, 512 \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128 \\ 1 \times 1, 512 \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128 \\ 1 \times 1, 512 \end{bmatrix} \times 8$
conv4_x	14×14	$\begin{bmatrix} 3 \times 3, 256 \\ 3 \times 3, 256 \end{bmatrix} \times 2$	$\begin{bmatrix} 3 \times 3, 256 \\ 3 \times 3, 256 \end{bmatrix} \times 6$	$\begin{bmatrix} 1 \times 1, 256 \\ 3 \times 3, 256 \\ 1 \times 1, 1024 \end{bmatrix} \times 6$	$\begin{bmatrix} 1 \times 1, 256 \\ 3 \times 3, 256 \\ 1 \times 1, 1024 \end{bmatrix} \times 23$	$\begin{bmatrix} 1 \times 1, 256 \\ 3 \times 3, 256 \\ 1 \times 1, 1024 \end{bmatrix} \times 36$
conv5_x	7×7	$\begin{bmatrix} 3 \times 3, 512 \\ 3 \times 3, 512 \end{bmatrix} \times 2$	$\begin{bmatrix} 3 \times 3, 512 \\ 3 \times 3, 512 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1, 512 \\ 3 \times 3, 512 \\ 1 \times 1, 2048 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1, 512 \\ 3 \times 3, 512 \\ 1 \times 1, 2048 \end{bmatrix} \times 3$	$\begin{bmatrix} 1 \times 1, 512 \\ 3 \times 3, 512 \\ 1 \times 1, 2048 \end{bmatrix} \times 3$
	1×1	average pool, 1000-d fc, softmax				
FLOPs		1.8×10^9	3.6×10^9	3.8×10^9	7.6×10^9	11.3×10^9

Figure 2.3 Overview of Res-Net Architectures [16]

2.3 Network Construction

To properly understand and visualize the performance of machine learning models on sex-based classification, it is paramount to construct and compare multiple architectures to ensure optimal performance and generalizability. As a result, two different CNN architectures are designed using a series of three-dimensional convolutions for feature extraction. Three-dimensional image convolution is identical to that in two-dimensions, with the kernel sliding over each voxel adding the weighted values of all its neighbors over the area it covers. Figure 2.4 below presents an example.




Figure 2.4 3-D Convolution and T1 Structural MRI Example

Of the two models the first constructed is a classical sequential CNN model. The details are illustrated in Figure 3.1 below.

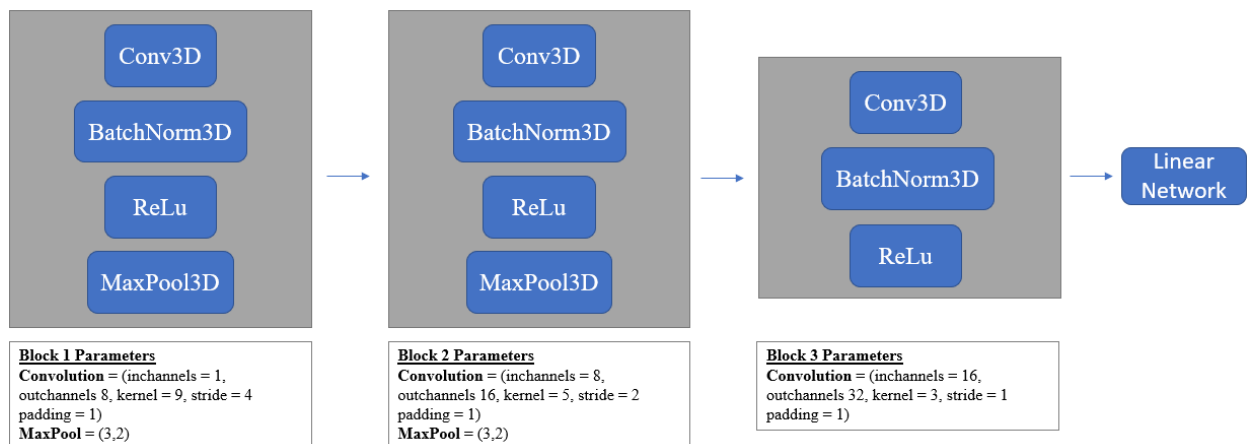


Figure 2.5 Classical CNN Architecture for MRI Sex Classification

This architecture is designed in three convolutional blocks to reduce the dimensionality of each image to a set of 32 feature maps each with dimensions 5x5x5. These feature maps are then fed

to a three layered fully connected network each with a rectified linear unit activation. The network is a lightweight and quick solution for feature extraction and model training.

The second architecture constructed was a ResNet18 model modified using three-dimensional convolutions. It is more complex than its sequential CNN counterpart and requires more computational time to properly train. The first column in Figure 2.3 above illustrates the design.

The output of both architectures are fed to a logistic function to produce a probability for prediction. Given the binary nature of the problem, a sigmoid with a classification threshold of 0.5 is reasonable; with anything above the threshold being classified as male and anything underneath being female.

2.4 Grad-Cams and Visualization Methods

Grad-Cam is a tool developed in 2017 that produces visual explanations for decisions from a large-class of CNN based models [17]. As an image is passed through the model, the last convolutional layer retains spatial information [17] and is the best for extracting the visualizations needed for determining the features which contribute most to the decision of the network. These feature maps are combined and backpropagated through the network to create a coarse heatmap of features of interest [17]. Figure 2.4 shows the process for a multi-class CNN, however the concept is the same for binary classification.

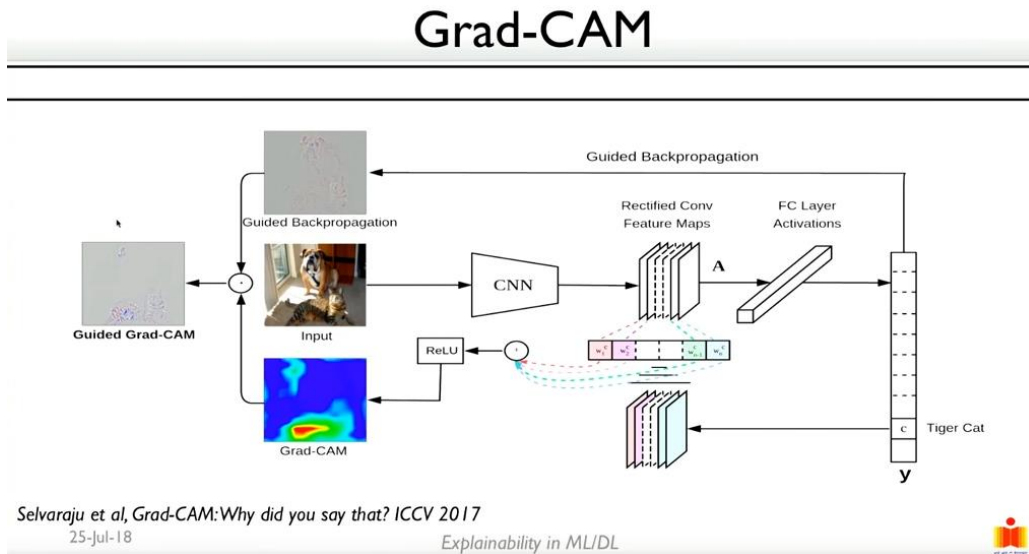


Figure 2.6 Grad-Cam Process [17]

More specifically for this project, the visualization will be performed via a sensitivity analysis [18] introduced by Simonyan et al in 2014. The sensitivity of a pixel, or voxel in this case, measures how much small changes in its value contribute to the classification function f of a network [19]. This is done by taking the norm over partial derivatives, acquired via backpropagation through the network, of a specific color channel c , shown in Formula 2.1 below.

$$h_p = \left\| \frac{\partial}{\partial x_p} f(x) \right\|_{\infty}$$

Formula 2.1 Sensitivity Analysis Formula [19]

The sensitivity analysis used in this project is provided by Rieke et al. [20] from their paper analyzing four visualization methods for MRI-based convolutional networks.

2.5 MRI Registration and AAL Atlas

T1 structural MRIs come in many different shapes and sizes meaning two images taken on the same machine a few months apart may not look identical. This is a problem when it comes to diagnosing conditions like tumor growth or neurological disease progression [21]. MRIs must be registered, or standardized, to a normal space so that images taken months apart or from separate machines can be properly compared and analyzed. MRI registration involves performing affine and rigid transformations on the original image to fit a standard space. A classic template used is the MNI 305 Average Brain, which is an average of 305 T1-weighted MRI scans linearly transformed to Talairach space [22] with a dimensionality of 172x220x156. This template is used for registration in the project. Figure 2.5 below shows an example of MRI normalization on a set of individual images. As seen, the normalization provides a clear way to visualize MR images from which more conclusive results can be drawn.

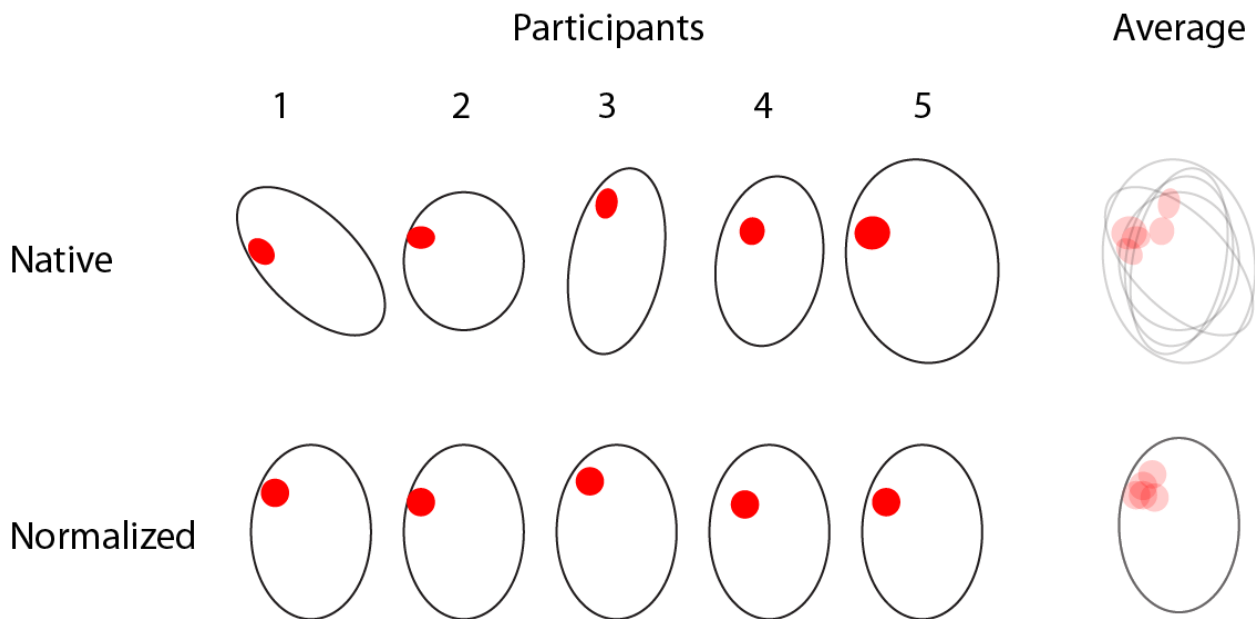


Figure 2.7 Example of MRI Normalization [23]

If MRI registration provides a simple method for normalizing brains to a standard space for visualization, then it is also useful to have a tool that can annotate and display which areas of interest are being analyzed. The AAL atlas is an automated anatomical parcellation of the spatially normalized average from 27 T1 volumes [24]. The parcellation results in 54 3D areas within each brain hemisphere where each region is manually drawn every 2mm on the axial

slices. The 108 regions are then reconstructed and assigned a label [24]. Each brain region is assigned a specific pixel value with all distinct areas sharing a different value. For example, 108 regions each have a pixel value ranging from 0 to 108 with 0 being the background. Figure 2.8 shows several slices taken from the atlas.

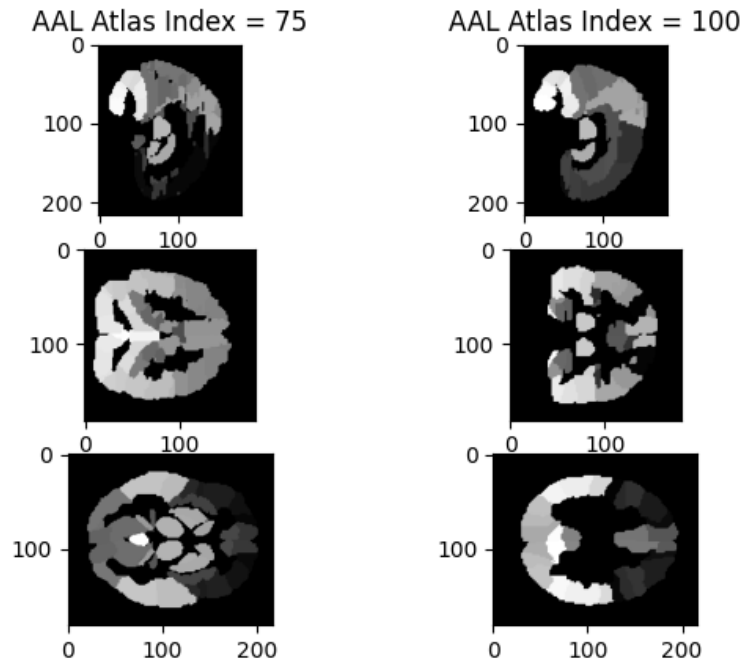


Figure 2.8 AAL Atlas Slices taken at Index 75 and 100 from axial, sagittal and coronal planes

2.6 Dataset Information

To properly construct the image classifiers, the patient image data must subscribe to a series of parameters to remove unwanted features during model training. First, a healthy subject without any underlying neurological conditions or diseases is required. Second, subjects must be between twenty-five and fifty years old. This ensures all patients had fully developed brains without any signs of deterioration.

The following research used images from two datasets, UKBiobank [25] and OASIS [16]. From the UKBiobank set, one thousand 3T brain MR images are chosen to train the image classifiers. This training stack consisted of 585 females and 411 males. Additionally, a test set of four hundred 3T images are acquired from the UKBiobank set, consisting of 225 female and 175 male scans.

To further test the model performance and generalizability, 81 1.5T MR images are chosen from the OASIS dataset. This test set contains 44 female, and 37 male scans that fit the data constraints.

Chapter 3: Methodology, Results and Discussion

The following chapter presents the methodology, procedures, results, and discussion of the project. Each section begins by stating its hypothesis and specific questions for clarity and understanding.

Part 1: Designing a deep-learning-based sex-classifier

It is important to restate the hypothesis mentioned in Chapter 1 to better understand the methodology of the section. Earlier, the first aim hypothesized that CNNs could correctly classify sex with high accuracy, over 80%. This section analyzes and selects the optimal model architecture for the binary classification task that gives the best generalizable results on raw and affinely transformed MR scans.

3.1.2 Methods

The methodology for the chapter is broken into the following three steps. First, acquiring and creating the necessary training and test sets for both the raw and affined image stacks. Second, model training and k-fold cross validation for optimal parameter selection. Lastly, model validation on the corresponding test stacks.

1. Data Acquisition and Organization

The training data is acquired from a subset of the UKBiobank dataset. It consists of one-thousand MR images of varying shapes and sizes, of which, 411 are male and 585 are female. Using Free Surfer pipelines, these images undergo a rigid transformation to the MNI 305 space for centering via the regalading pipeline from the nifty reg package. No affined registration is performed. The stack is then split into a training and testing set of 950 and 50 images respectively. To test the model, two separate datasets are acquired from the OASIS and the UKBiobank databank. The OASIS dataset consists of 81 MR images, 44 females and 37 males, which are also registered via a rigid transformation to the MNI 305 space. The UKBiobank test set consists of 400 MR images, 225 females and 175 males, rigidly transformed to the MNI305 space.

Now that each training and validation group has been created. A separate training and validation stack is formed by performing an affined registration on each image in the UKBiobank and OASIS stacks to the MNI 305 space. The regaladin pipeline from the niftyreg package is used to perform the affine transformations for each image. The split of training and testing images remains the same as before.

Two stacks of training and testing data are created: one for training and predicting raw images and another for training and predicting on affined images.

2. Model Training

Two architectures are employed to achieve the objectives listed above: a classical CNN and a ResNet18 model. These are described in further detail in Chapter 2. Each model is trained for 50 epochs and is evaluated using three-fold cross validation, where each fold is trained using either the PyTorch SGD or Adam optimizers with a learning rate ranging from $1e-2$ to $1e-3$. The best parameters and hyperparameters are chosen by comparing the average loss and accuracy of all

the folds. The optimal model is then chosen with the hyperparameters that result in the lowest of these averages. Once the appropriate values are found, the model is trained on the entire training stack for 100 epochs and validated on the OASIS and UKB set for proper results. The process is repeated for the raw and affined image stack resulting in two models for sex-based prediction.

3. Model Validation

With every model being trained on the UKB training set, each is then validated on the OASIS and UKBiobank test images respectively, and their performance are analyzed via a confusion matrix. Models could not be trained on the OASIS dataset due to a small number of images. The accuracy, precision, recall and AUC/ROC metrics are used to gain insight on the model. Finally, these metrics are used to compare the classical and ResNet18 architectures to determine the optimal model for each image stack.

3.1.3 Results

To begin, it is best to visualize the rigid and affined transformation of each training and validation set to better understand the data. Figures 3.1 and 3.2 below illustrate the rigid and affined transformations of the UKBiobank and OASIS datasets. To help understand information about the data used, Table 3.1 displays the breakdown of datasets.

	Number of Males	Number of Females	Mean Age	Standard Deviation of Age	Scanner Strength
UKB Train	411	585	42	8.4	3T
UKB Test	175	225	37	8.8	3T
OASIS Test	37	44	36	9.3	1.5T

Table 3.1 Breakdown of training and testing dataset information




Figure 3.1 UKBiobank Raw, Rigid, Affined and MNI 305 Template Image




Figure 3.2 OASIS Rigid, Affined and MNI 305 Template Images

Raw Models

The first several models are trained on the rigid registrations and evaluated using 3-fold cross validation to ensure optimal parameters for training. Table 3.1 below illustrates the average training and testing losses/accuracies based on varying parameters.

Model Architecture	Optimizer	Learning Rate	Avg Training Loss	Avg Training Acc	Avg Validation Loss	Avg Validation Acc
Classical	SGD	0.01	2.838	98.18	3.262	96.32
Classical	SGD	0.001	10.305	692.43	6.001	90.74
Classical	Adam	0.01	6.840	96.38	9.848	93.87
Classical	Adam	0.001	2.775	98.32	5.178	95.46
ResNet18	SGD	0.01	3.045	98.23	2.974	96.30
ResNet18	SGD	0.001	7.854	95.33	4.66	94.27
ResNet18	Adam	0.01	6.146	96.19	5.227	94.12
ResNet18	Adam	0.001	3.782	97.72	3.859	95.72

Table 3.1 K-fold Cross Validation Results for Models Trained on Raw MR Images. Optimal parameters for best performing model highlighted in yellow

The best parameters and models are highlighted in yellow. The corresponding loss and accuracy plots are shown below. From the loss plots, there does not seem to be any signs of overfitting for either model architecture but there is a greater instability in accuracies for the Resnet18 model compared to the sequential CNN.

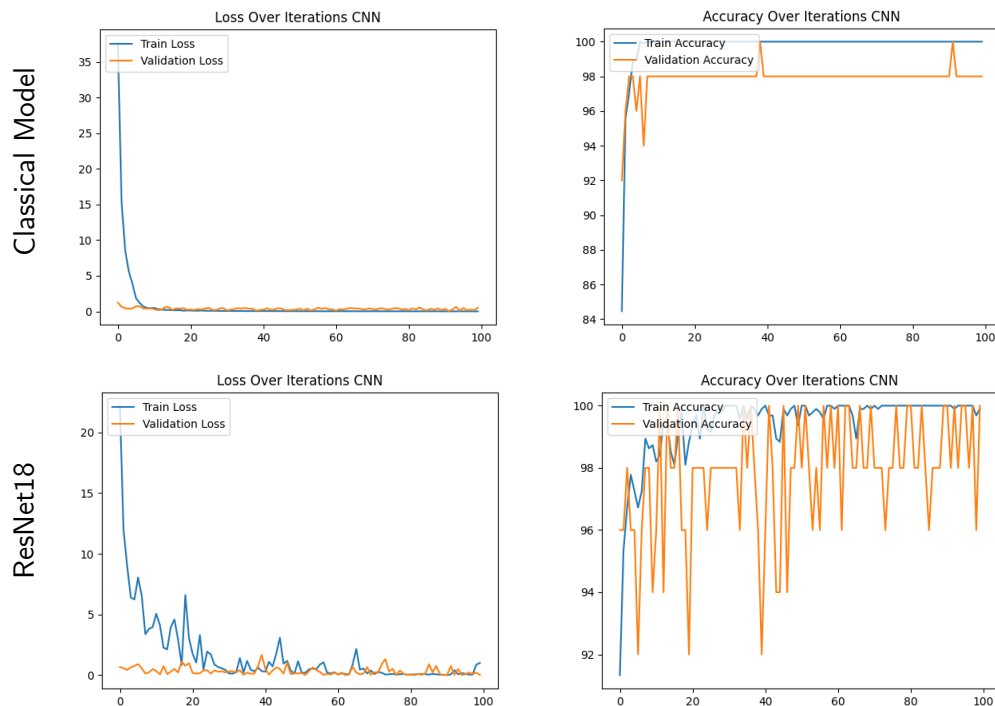


Figure 3.3 Classical (Top) and ResNet18(Bottom) Training and Testing Loss/Accuracies

Following the training, the models are tested on the OASIS and UKBiobank images. With a classification threshold of 0.5. The following result are displayed in Tables 3.2 and 3.3.

Classical Model, Thresh = 0.5		ResNet18 Model, Thresh = 0.5	
OASIS Test Dataset	Evaluation Metrics	OASIS Test Dataset	Evaluation Metrics
Accuracy	90.1%	Accuracy	54.32%
Precision	96.77%	Precision	Nan
Recall	81.08%	Recall	0%
Male Accuracy (Sensitivity)	81.08%	Male Accuracy	0%
Female Accuracy (Specificity)	97.72%	Female Accuracy	100.0%

Confusion Matrix Classical		
	0	1
0	43	1
1	7	30

Confusion Matrix ResNet18		
	0	1
0	44	0
1	37	0

*Note 0 is Female, 1 is Male

Table 3.2 Model Performance Statistics with Classification Threshold at 0.5 on OASIS Dataset

Classical Model, Thresh = 0.5		ResNet18 Model, Thresh = 0.5	
UKB Test Dataset	Evaluation Metrics	UKB Test Dataset	Evaluation Metrics
Accuracy	97.0%	Accuracy	56.25%
Precision	96.05%	Precision	Nan
Recall	97.14%	Recall	0%
Male Accuracy (Sensitivity)	97.14%	Male Accuracy	0%
Female Accuracy (Specificity)	96.88%	Female Accuracy	100.0%

Confusion Matrix Classical		
	0	1
0	218	7
1	5	170

Confusion Matrix ResNet18		
	0	1
0	225	0
1	175	0

Table 3.3 Model Performance Statistics with Threshold at 0.5 on UKBiobank Testset

Tables 3.2 and 3.3 really show the success of the classical CNN compared to its Resnet counterpart. On both test sets, the sequential model averaged extremely high-performance scores. On the other hand, Resnet18 suffered with male classification, predicting all volumes to be female under the 0.5 threshold. With the performance of both architectures on the test sets, the following ROC/AUC curves were generated.

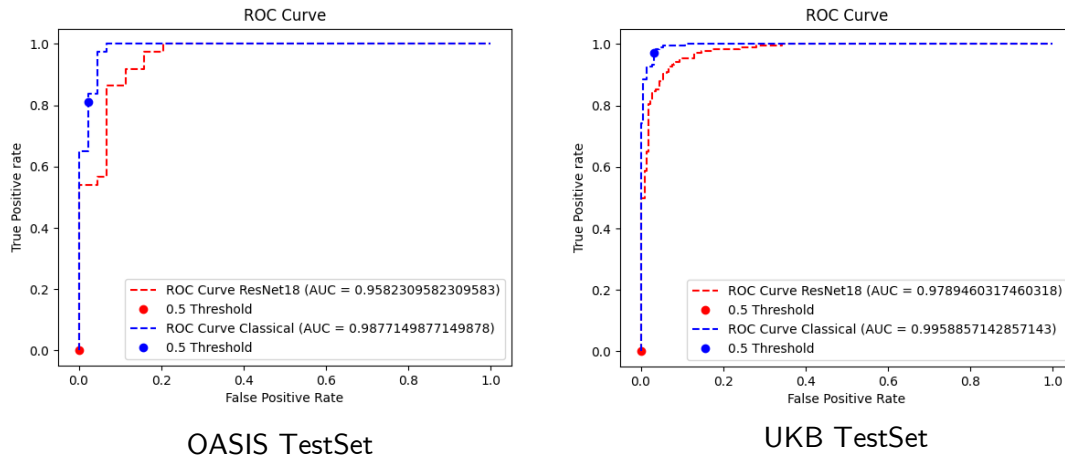


Figure 3.4 ROC Curves per Model and Testset. Blue represents the Classical Architecture and Red shows ResNet18 performance across thresholds. The green dot shows the arbitrary 0.5 classification threshold chosen.

The resulting ROC curves for both architectures show very good performance across classification thresholds. However, in terms of model generalizability, ResNet18 requires careful selection of its threshold to provide accurate results whereas the sequential CNN classifies well with the arbitrary 0.5 threshold. In addition, the roc plot of the classical model is more desirable for this classification problem.

The ResNet18 threshold is found to be 0.008 for UKB and 0.0075 for OASIS. These optimal thresholds are found by taking a geometric average of the sensitivity and specificity on the test set over a range of values. Despite having a good ROC curve, the model's generalizability is lacking.

Affined Model

After completing the training and evaluation on the raw, rigidly transformed images, the next step involved training and validating models on the affined image stack. Table 3.6 below shows the 3-fold cross validation of the models trained on affined MR images.

Model Architecture	Optimizer	Learning Rate	Avg Training Loss	Avg Training Acc	Avg Validation Loss	Avg Validation Acc
Classical	SGD	0.01	1.997	99.06	3.223	96.72
Classical	SGD	0.001	10.718	96.01	6.792	93.97
Classical	Adam	0.01	2.887	98.40	4.392	96.47
Classical	Adam	0.001	1.699	99.08	3.139	97.32
ResNet18	SGD	0.01	2.980	98.44	4.277	96.11
ResNet18	SGD	0.001	6.598	96.85	4.658	95.52
ResNet18	Adam	0.01	5.737	96.89	5.143	95.61
ResNet18	Adam	0.001	3.875	97.97	4.161	96.27

Table 3.4 K-fold Cross Validation Results for Models Trained on Affined MR Images. Optimal parameters for best performing model highlighted in yellow

The respective losses and accuracies of each model architecture are shown below in Figure 3.6.

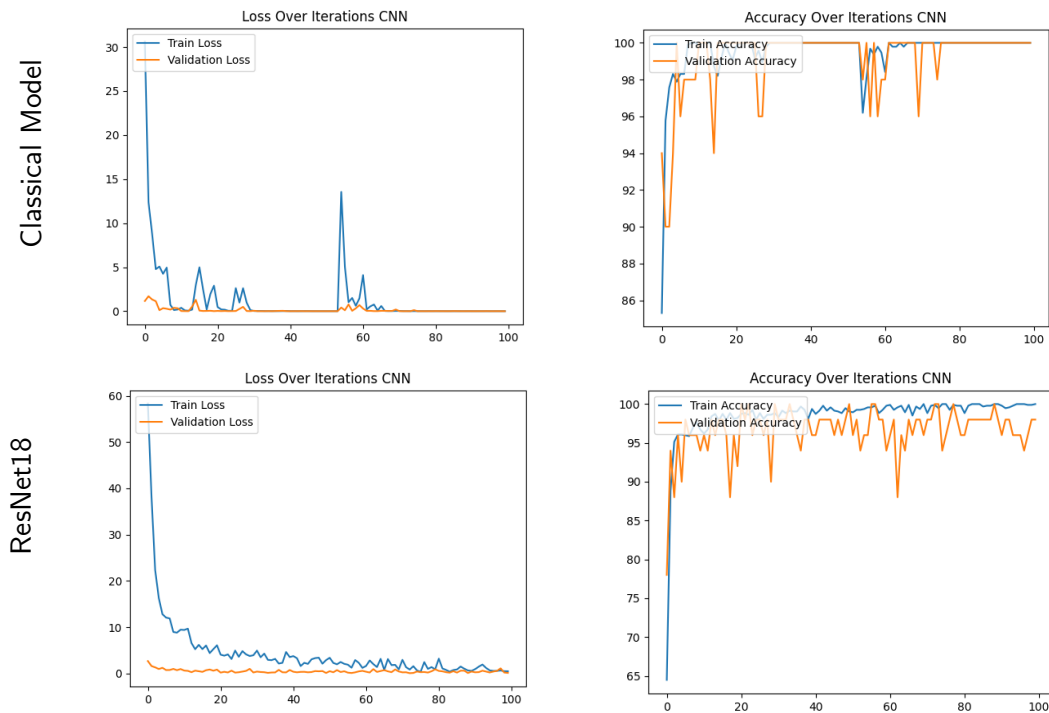


Figure 3.5 Classical (Top) and ResNet18(Bottom) Training and Testing Loss/Accuracies

Like the models trained on rigid images, there does not seem to be any signs of overfitting for either model architecture but there is a greater instability in testing accuracies for the Resnet18 model compared to the sequential CNN.

With the models fully trained, their respective validation scores on the OASIS and UKB test sets are shown in Table 3.4 and 3.5.

Classical Model, Thresh = 0.5		ResNet18 Model, Thresh = 0.5	
OASIS Test Dataset	Evaluation Metrics	OASIS Test Dataset	Evaluation Metrics
Accuracy	91.4%	Accuracy	54.32%
Precision	100.0%	Precision	Nan
Recall	81.1%	Recall	0%
Male Accuracy (Sensitivity)	81.1%	Male Accuracy	0%
Female Accuracy (Specificity)	100.0%	Female Accuracy	100.0%

Confusion Matrix Classical		
	0	1
0	44	0
1	7	30

Confusion Matrix ResNet18		
	0	1
0	44	0
1	37	0

*Note 0 is Female, 1 is Male

Table 3.5 Model Performance Statistics with Threshold at 0.5 on OASIS Dataset

Classical Model, Thresh = 0.5		ResNet18 Model, Thresh = 0.5	
UKB Test Dataset	Evaluation Metrics	UKB Test Dataset	Evaluation Metrics
Accuracy	98.0%	Accuracy	56.25%
Precision	98.26%	Precision	Nan
Recall	97.14%	Recall	0%
Male Accuracy (Sensitivity)	97.14%	Male Accuracy	0%
Female Accuracy (Specificity)	98.67%	Female Accuracy	100.0%

Confusion Matrix Classical		
	0	1
0	222	3
1	5	170

Confusion Matrix ResNet18		
	0	1
0	225	0
1	175	0

Table 3.6 Model Performance Statistics with Threshold at 0.5 on UKBiobank Testset

The same pattern as before follows, with the classical network heavily outperforming Resnet with the 0.5 threshold. Resnet suffers from the same issue, classifying all scans as female. Furthermore, the sequential CNN outperforms itself on the affined image stack showing higher accuracy, precision, recall along with better sensitivity and specificity than its rigid image counterpart.

Using an ROC/AUC curve the threshold for the OASIS dataset classification can be modified. Figure 3.6 illustrates the varying threshold levels and the corresponding optimal performance statistics.

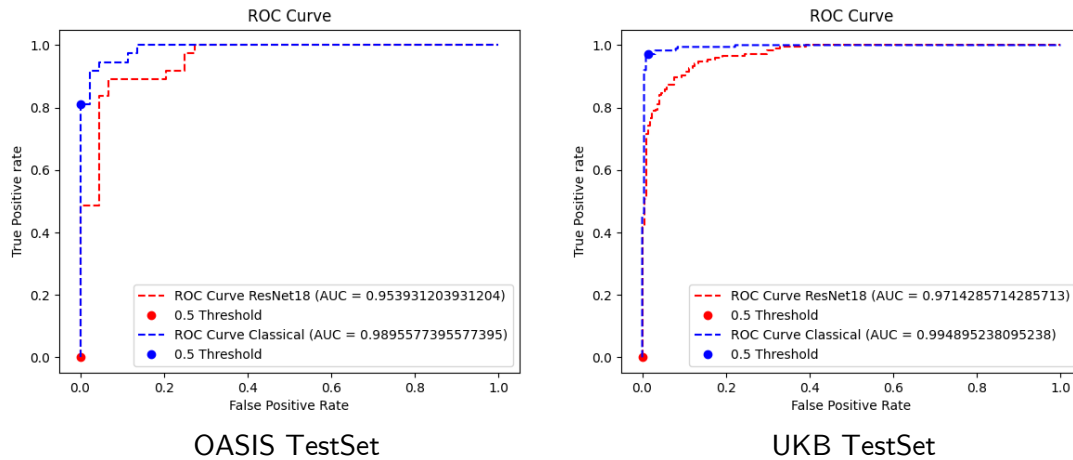


Figure 3.6: ROC Curves per Model and Testset. Blue represents the Classical Architecture and Red shows ResNet18 performance across thresholds. The green dot shows the arbitrary 0.5 classification threshold chosen.

Both architectures show promise in their classification abilities however the Resnet model again needs careful selection of its threshold to properly classify each test set. As a result, the classical model outperforms Resnet in terms of generalizability across test sets using the arbitrary 0.5 threshold.

The ResNet18 threshold is found to be 0.055 for UKB and 0.059 for OASIS. These optimal thresholds are found by taking a geometric average of the sensitivity and specificity on the test set over a range of values. Despite having a good ROC curve, the model’s generalizability is lacking.

3.1.4 Discussion

By comparing the 3-fold cross validation performance of different model parameters, it became easy to recognize which architecture was optimal for the binary classification task. A 3-fold cross validation was used in both cases due to the number of training images. 5-fold would have drastically reduced the number of testing images and may not have given a good view on how the model performs with each fold.

For the raw, rigidly transformed images, the SGD optimizer with a learning rate of 0.01 outperformed the other parameters by having the highest training and testing accuracies while

maintaining the lowest training and testing loss. The loss and accuracy plots for the sequential and ResNet18 model shown in Figure 3.3 demonstrate no overfitting and very high accuracies, between 96% and 100% on the testing and training sets. Interestingly, the sequential model appears to be more stable in accuracy and loss compared to ResNet18. This is supported further by the performance metrics where each model is tested on the OASIS and UKB test sets. Tables 3.2 and 3.3 illustrate the classification performance with a threshold of 0.5. It becomes clear that the sequential architecture outperforms ResNet18 in terms of generalizability given the high accuracies for both test sets. In contrast, ResNet18 struggles to classify males in both test sets, predicting all images to be female. Each test set has its own ROC/AUC score shown in Figure 3.4. The sequential model still outperforms ResNet18 with both test sets averaging higher AUC scores and better ROC curves. In addition, the classical architecture is more generalizable with a classification threshold of 0.5 as opposed to ResNet18 which requires a different threshold based on the test set used. The classical model is much more lightweight than its ResNet18 counterpart, requiring less computational power and time to train.

The success of the sequential architecture carries over to the models trained on affined images. Table 3.4 displays the optimal parameters for each design highlighted in yellow. The corresponding loss and accuracy plots are illustrated in Figure 3.6. As with the raw images, the classical model seems more stable in terms of accuracy and loss than ResNet18. The accuracy of both models is quite high, again ranging from 96% to 100% and the loss plots show no sign of overfitting. When analyzing the performance metrics on the OASIS and UKB affined test sets with a classification threshold of 0.5, it becomes clear the classical CNN generalizes much better than ResNet18 averaging an accuracy over 90% for both test sets. ResNet18 suffers from the same issue as before, in where the classification fails for all male images. The corresponding ROC/AUC curves demonstrate good classification capability from both architectures. Like the raw image models, the classical CNN outperforms ResNet18 by being more generalizable across test sets with a threshold of 0.5.

Furthermore, the ResNet18 model might not have performed as well since the training images were not intensity normalized. This may be why the model acted sporadically and did not have such high classification accuracies. With further data preprocessing, the Resnet18 model could have very different results.

Nevertheless, the classical CNN is the optimal performer for both raw and affined images. Furthermore, it can extend its classification predictions across MRI field strengths ranging from 1.5T to 3T without sacrificing performance. Given the complex design of ResNet18, the model may be looking for specific features that are not generalizable or relevant to the research problem. The simple architecture employed was capable of classifying male and female brains with an accuracy of over 90% across test sets and transformations; hence answering the research hypothesis and questions proposed at the beginning of the section.

Part 2: Explainable AI to explore brain-regions that determine sex-based diffeomorphism

As stated in Chapter 1, given the success of the model classification, the hypothesis is that grad-cams and saliency maps can provide novel information on region-wise sex-based diffeomorphisms while showing that brain classification is not fully driven by size and shape. This section explores the visualization of the classification decisions of both Resnet18 and classical models. It also shows if the generated saliency maps are different between the raw and affined image stacks and lists the features which contribute most to the model classification decision.

3.2.2 Methods

The methodology for the chapter is broken into the following two steps. First, the construction of the grad-cam and related saliency maps for visualization. And second, the explanation of model classification decisions based on differentiating features uncovered by the deep-learning networks.

1. Grad-Cam Construction

To properly visualize the results of the models, a grad-cam is used via a sensitivity analysis [20] [18]. This tool highlights relevant areas that contribute to a model's classification decision as a heatmap overlaid on the input image. The sensitivity analysis developed by Rieke [20] for MRI-based convolutional networks is used to perform these visualizations. The saliency maps generated are then utilized in creating a population-wise grad-cam.

2. Explanation of model classification decisions

A single sample grad-cam is not suitable for visualizing sex differences as it is prone to noise as the model backpropagates to generate the saliency map [20]. A population wise grad-cam is more appropriate for examining differences to determine conclusive results for large subject groups. A grad-cam heatmap is created for all images in the OASIS and UKBiobank test set and an average is taken among the resulting male and female saliency maps. This averaging removes any outliers or noise otherwise present in individual grad-cams. The result is two population-wise heatmaps, one for males and one for females, which are then overlaid onto a subject for visualization.

These two grad-cams contain information on the existence and location of potential dimorphisms within the brain. A rigorous visual inspection is conducted across both heatmaps to pinpoint specific areas of interest. This process is then repeated for the models trained on the raw and affined image stacks, and the resulting population-wide saliency maps are compared. A relevance table is then constructed based on the aal Atlas [24] to list the most important brain areas responsible for the sex classification. This provides a numerical method to complement the visual inspection for rigor and verification.

This relevance table is constructed by generating a series of area masks each associated with a specific region in the aal atlas. Then the saliency maps generated are multiplied and summed with each area mask to create a relevance score. The more overlapping pixels in the area mask and saliency map will result in a higher score. These scores are then merged with the

corresponding area names and sorted in descending order. The top scoring regions are the most relevant to the model classification being the areas where the saliency map overlaps the individual masks the most.

3.2.3 Results

Raw Model Visualization

Classical Architecture

Given the success of the sequential model in correctly predicting sex, it is important to visualize where and how the CNN makes its decision. Performing a population grad-cam on the OASIS and UKB test image stack, the saliency maps generated for the classical architecture are shown in Figures 3.7 to 3.10 below.




Figure 3.7 Population-wide average female saliency map generated from the rigidly transformed OASIS testset using the Classical Model




Figure 3.8 Population-wide average male saliency map generated from the rigidly transformed OASIS testset using the Classical Model

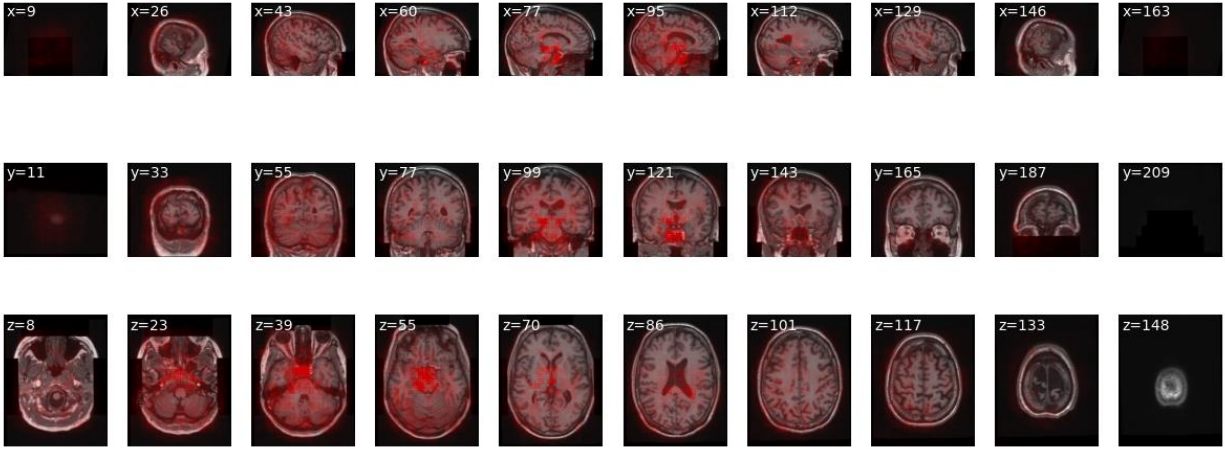


Figure 3.9 Population-wide average female saliency map generated from the rigidly transformed UKB testset using the Classical Model




Figure 3.10 Population-wide average male saliency map generated from the rigidly transformed UKB testset using the Classical Model

Both sets of images show similar highlighted regions across test sets. The model seems to look for differences towards the center of the brains with some scattering towards the lateral and posterior sides. Despite being trained on 3T volumes, the classifier sees similar differences in lower resolution images like the 1.5T Oasis MR images. Overlaying the heatmaps with aal atlas and creating a sorted table of brain regions most important to the model classification, a relevance table is created for both test sets and displayed in Tables 3.7 and 3.8. Out of the 110 available brain regions in the atlas, the top ten areas that contribute to the classification are shown.

	Male Relevance	Female Relevance
Relevance	Brain Area	Brain Area
1	Middle Temporal Gyrus	Middle Temporal Gyrus
2	Precuneus	Precuneus
3	Cerebellum Crus 1	Cerebelum_Crus1
4	Inferior Temporal Gyrus	Inferior Temporal Gyrus
5	Postcentral	Postcentral
6	Lingual	Lingual
7	Middle Occipital Gyrus	Middle Occipital Gyrus
8	Fusiform	Middle Frontal Gyrus
9	Middle Frontal Gyrus	Fusiform
10	Precentral	Precentral

Table 3.7 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and classical architecture for the rigidly transformed OASIS test set.

	Male Relevance	Female Relevance
Relevance	Brain Area	Brain Area
1	Middle Temporal Gyrus	Middle Temporal Gyrus
2	Precuneus	Precuneus
3	Cerebellum Crus 1	Postcentral
4	Postcentral	Cerebellum Crus 1
5	Inferior Temporal Gyrus	Inferior Temporal Gyrus
6	Mid Occipital Gyrus	Middle Frontal Gyrus
7	Middle Frontal Gyrus	Middle Occipital Gyrus
8	Precentral	Precentral
9	Fusiform	Fusiform
10	Lingual	Lingual

Table 3.8 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and classical architecture for the rigidly transformed UKB test set.

As listed above, the top regions of interest with the highest contribution are the temporal gyri, cerebellum, and precuneus and postcentral gyri. Despite generating saliency maps for MR

images of differing field strengths, both test sets list similar ROI's that contribute to model classification.

ResNet18 Architecture

Despite not performing as well as the classical architecture, it is important to visualize the decisions of the ResNet18 model to better understand where the model searches for discrepancies. The saliency maps illustrating the classification decisions for the ResNet18 architecture for the raw OASIS and UKB test sets are displayed in Figures 3.11-3.14 below.




Figure 3.11 Population-wide average female saliency map generated from the rigidly transformed OASIS testset using ResNet18




Figure 3.12 Population-wide average male saliency map generated from the rigidly transformed OASIS testset using ResNet18




Figure 3.13 Population-wide average female saliency map generated from the rigidly transformed UKB testset using ResNet18




Figure 3.14 Population-wide average male saliency map generated from the rigidly transformed UKB testset using ResNet18

The figures illustrate a wide-ranging grad-cam, highlighting the entire volume as an ROI rather than a specific feature. Concentrations of pixels are found around the skull hinting that the model is looking at the overall size and shape of the brain itself rather than the brain structures.

However, the relevance table constructed using the overlay from the aal atlas shows similar areas to that of the sequential CNN. The relevance of each area for the classification is shown below in Tables 3.9 and 3.10

	Male Relevance	Female Relevance
Relevance	Brain Area	Brain Area
1	Middle Frontal Gyrus	Middle Frontal Gyrus
2	Middle Temporal Gyrus	Middle Temporal Gyrus
3	Inferior Temporal Gyrus	Inferior Temporal Gyrus
4	Superior Frontal Gyrus	Superior Frontal Gyrus
5	Cerebellum Crus 1	Middle Occipital Gyrus
6	Inferior-Frontal Gyrus, Triangular	Inferior-Frontal Gyrus, Triangular
7	Postcentral	Precuneus
8	Fusiform	Postcentral
9	Precuneus	Superior Frontal Gyrus, Medial
10	Precentral	Cerebellum Crus 1

Table 3.9 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and ResNet18 for the rigidly transformed OASIS test set

	Male Relevance	Female Relevance
Relevance	Brain Area	Brain Area
1	Middle-Temporal Gyrus	Middle-Temporal Gyrus
2	Precuneus	Precuneus
3	Cerebellum Crus 1	Postcentral
4	Postcentral	Cerebellum Crus 1
5	Inferior-Temporal Gyrus	Inferior-Temporal Gyrus
6	Middle-Occipital Gyrus	Middle Frontal Gyrus
7	Middle Frontal Gyrus	Middle-Occipital Gyrus
8	Precentral	Precentral
9	Fusiform	Fusiform
10	Lingual	Lingual

Table 3.10 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and ResNet18 for the rigidly transformed UKB test set

The temporal, precuneus and postcentral gyri are at the forefront of the classification, like in the previous model. However, this is overshadowed by the conglomeration of red pixels around the skull of the volume as shown in the figures above. This leads to question if ResNet18 looks for sex specific features or rather just size and shape of the individual volumes.

Affined Model Visualization

Classical Architecture

Performing a population grad-cam on the affined OASIS and UKB test stack, the saliency maps generated for the classical architecture are shown in Figures 3.15 to 3.18 below.




Figure 3.15 Population-wide average female saliency map generated from the affined OASIS testset using the Classical Architecture




Figure 3.16 Population-wide average male saliency map generated from the affined OASIS testset using the Classical Architecture

The resulting population saliency maps at the 1.5T field strength show highlights similar to those of the raw image grad-cams from above. With highlighted regions around the center, posterior and lateral sides of the brain. The 3T images from the UKB test set below show similar concentrations of red pixels.




Figure 3.17 Population-wide average female saliency map generated from the affined UKB testset using the Classical Architecture




Figure 3.18 Population-wide average male saliency map generated from the affined UKB testset using the Classical Architecture

By inspection, similar regions are highlighted in red between the OASIS and UKB test sets, around the center of the brain with some spots toward the frontal and parietal lobe. To further understand which regions contribute the most, the relevance of each brain area for sex classification is constructed and shown below in Table 3.6 and 3.7.

	Male Relevance	Female Relevance
Relevance	Brain Area	Brain Area
1	Middle Temporal Gyrus	Inferior Temporal Gyrus
2	Middle Frontal Gyrus	Middle Temporal Gyrus
3	Inferior Temporal Gyrus	Middle Frontal Gyrus
4	Cerebellum Crus 1	Median Cingulate Gyrus
5	Middle Occipital Gyrus	Cerebellum Crus 1
6	Superior Frontal Gyrus	Fusiform
7	Postcentral	Postcentral
8	Median Cingulate Gyrus	Middle Occipital Gyrus
9	Precentral	Precuneus
10	Superior Frontal Gyrus, Medial	Calcarine

Table 3.11 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and classical architecture for the affined OASIS test set

	Male Relevance	Female Relevance
Relevance	Brain Area	Brain Area
1	Middle Temporal Gyrus	Middle Temporal Gyrus
2	Middle Frontal Gyrus	Middle Frontal Gyrus
3	Inferior Temporal Gyrus	Inferior Temporal Gyrus
4	Superior Frontal Gyrus	Superior Frontal Gyrus
5	Cerebellum Crus 1	Cerebellum Crus 1
6	Postcentral	Postcentral
7	Middle Occipital Gyrus	Superior Frontal Gyrus, Medial
8	Median Cingulate Gyrus	Middle Occipital Gyrus
9	Precentral	Fusiform
10	Fusiform	Median Cingulate Gyrus

Table 3.12 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and classical architecture for the affined UKB test set

At a quick glance, the top areas of interest are similar amongst both test sets, with the temporal and frontal gyri holding the most significance. Like the raw images in the previous section, the affined grad-cam picks up the same ROI's; hinting that both models are searching for sex-specific features similar across affined and rigid transformations.

ResNet18 Architecture

Like with the raw image model, it is important to visualize the classification of ResNet18 despite its poor performance at the general 0.5 threshold. The resulting saliency maps for the ResNet18 architecture are displayed in Figures 3.19-3.22 below.




Figure 3.19 Population-wide average female saliency map generated from the affined OASIS testset using ResNet18




Figure 3.20 Population-wide average male saliency map generated from the affined OASIS testset using ResNet18

ResNet18 highlights the entire volume with concentrations of pixels around the skull indicating size and shape as differentiating features. This is further illustrated in the UKB saliency maps below.




Figure 3.21 Population-wide average female saliency map generated from the affined UKB testset using ResNet18




Figure 3.22 Population-wide average male saliency map generated from the affined OASIS testset using ResNet18

Like ResNet18 for raw images, the population wide grad-cam does not emphasize any particular set of features as different between male and female brains in any test set. However, the list of relevant areas to brain classification shown in Tables 3.13 and 3.14 are quite similar to those of the sequential CNN, illustrating similar regions of interest.

	Male Relevance	Female Relevance
Relevance	Brain Area	Brain Area
1	Middle Frontal Gyrus	Middle Frontal Gyrus
2	Middle Temporal Gyrus	Middle Temporal Gyrus
3	Inferior Temporal Gyrus	Superior Frontal Gyrus
4	Superior Frontal Gyrus	Inferior Frontal Gyrus
5	Postcentral	Postcentral
6	Precuneus	Precuneus
7	Precentral	Precentral
8	Inferior Frontal Gyrus, Triangular	Superior Frontal Gyrus, Medial
9	Superior Temporal Gyrus	Middle Occipital Gyrus
10	Middle Occipital Gyrus	Inferior Frontal Gyrus, Triangular

Table 3.13 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and ResNet18 for the affined OASIS test set

	Male Relevance	Female Relevance
Relevance	Brain Area	Brain Area
1	Middle Frontal Gyrus	Middle Frontal Gyrus
2	Middle Temporal Gyrus	Middle Temporal Gyrus
3	Inferior Temporal Gyrus	Inferior Temporal Gyrus
4	Postcentral	Superior Frontal Gyrus
5	Superior Frontal Gyrus	Postcentral
6	Precentral	Precentral
7	Precuneus	Precuneus
8	Cerebellum Crus 1	Middle Occipital Gyrus
9	Middle Occipital Gyrus	Inferior Frontal Gyrus, Triangular
10	Superior Frontal Gyrus, Medial	Superior Frontal Gyrus, Medial

Table 3.14 Top 10 brain regions responsible for model classification sorted in order from most to least significant based on the aal brain atlas and ResNet18 for the affined UKB test set

As listed in the tables above, the temporal and front gyri hold the biggest significance for the model classification along with the precuneus and, precentral and postcentral gyri. Interestingly, the ResNet18 model list similar areas to those of the classical CNN, however, due to the poor visualization and prediction capabilities, it is difficult to be certain of these results since the grad-cam hints at size being a key-differentiating feature between the sexes.

3.2.4 Discussion

Figures 3.7 to 3.10 illustrate the population wide saliency maps generated by a sensitivity analysis for males and females for the OASIS and UKB test sets. The areas highlighted in red are regions that contribute most to the model sex classification. These figures are complemented by Tables 3.7 and 3.8 where the regions of interest are shown in order of importance to the classification task. The top ten most influential areas to the decision are listed for each sex. Among the OASIS and UKB test sets, the classical model uncovers similar regions of importance to classification. Among the areas listed, the middle temporal gyrus, precuneus, cerebellum, postcentral and inferior temporal gyri contribute most to the model's decision. Figures 3.7-3.10 show similar concentrations of pixels in the heatmap between test sets.

On the other hand, figure 3.11 to 3.14 are more difficult to interpret. By inspection, it seems the ResNet18 saliency maps highlight the entire brain volume for both sexes across test sets. The area relevance listed in Tables 3.9 and 3.10 list the middle frontal, inferior temporal, and middle temporal gyri as most influential to the classification. However, the saliency maps tend to show a spread across the entire image. This could be due to the model searching for hyper-specific features across the entire volume, rather than focusing on those that contribute to sex differences.

Comparing the population-wide raw saliency results to those of the affined images, the classical model (Figures 3.15 to 3.18) shows similar brain regions highlighted in red. This is further supported by the brain area relevance in Tables 3.11 and 3.12. For both test sets, the top overlapping areas for model classification are the middle and inferior temporal gyri along with the cerebellum, middle frontal gyrus and superior frontal gyrus.

In contrast to the sequential CNN, ResNet18 highlights the entire volume in both test sets in Figures 3.19 to 3.22. As before, due to the complexity of the classifier, it searches for specific features that are not general to sex classification but rather to the volume itself. In addition, along the axial plane (3rd row of each image), a concentration of red pixels can be seen around the skull towards the corners of the image. This indicates, the model is viewing the entire MRI as different rather than features within the volume. Nevertheless, tables 3.13 and 3.14 show similar areas as those in tables 3.11 and 3.12. Despite highlighting the entire brain, it still sees some specific regions as more important than others, indicating a higher concentration of red pixels in the same areas as the classical model.

Table 3.15 below shows a list of related literature and its results on structural sex differences within the brain. Comparing the results presented with previous findings helps gauge the validity and novelty of the research.

Authors	Purpose	Sample Size and Characteristics	Summary of Results
Julia Sacher et. Al. [26] (2011)	A comprehensive summary exploring sex dimorphisms in the human brain at the structural and functional level	Review Article	Increases in brain volume in temporal and occiput gyri in men. Women showed more gray matter in cingulate cortices and parietal lobules. White matter differences in fusiform, hippocampus along with frontal and temporal lobes.
Xi Zhang et. Al. [27] (2020)	To show that gender differences are encoded differently within the structure and function of the human brain	290 T1 MRI's. Voxel-wise comparison performed using two t-tests	Gender differences detected structurally in frontal, parietal, occipital lobes and cerebellum with greater GMV in females
A. N. Ruigrok et. Al. [28] (2014)	A meta-analysis of typical sex differences on global brain volume	Combination of multiple studies using a new Gaussian-process regression coordinate-based meta-analysis for examination	Regional sex differences in volume and tissue density include the amygdala, hippocampus and insula, areas known to be implicated in sex-biased neuropsychiatric conditions.
Jill M Goldstein et. Al. [7] (2001)	A comprehensive evaluation of normal sexual dimorphisms of cortical and subcortical brain regions, using in vivo magnetic resonance imaging	48 T1 Weighted 1.5T MRI's. Scans analyzed and segmented via MR technicians	Women had larger volumes in particularly in the frontal and medial paralimbic cortices with men showing larger volumes in the frontal medial cortex, amygdala and hypothalamus
Stuart J. Ritchie et. Al. [5] (2018)	Large single sample study of structural and functional differences in the human brain.	5216 T1 weighted 3T MRI's. Analyzed via various statistical methods	GMV differences present in thalamus, amygdala and hippocampus with females having smaller volumes but more cortical thickness across entire brain.
Xiaohua Chen et. Al. [29] (2007)	The study examined sex-related differences in regional gray matter (GM) in 44-48 year old healthy individuals	411 T1 weighted MRI's. Analyzed via voxel based morphometry	Regional sex dimorphism was present, with men having more GM volume in midbrain, left inferior temporal gyrus, right occipital lingual gyrus, right middle temporal gyrus, and both cerebellar hemispheres. Women showed more GM in dorsal anterior, posterior and ventral cingulate cortices, and right inferior parietal lobule

Table 3.15 List of reviewed studies and their results

The differences observed using the classical CNN are of great interest given the model's generalizability and high accuracy predictions. Literature surrounding brain sex differences shows differences in volume and tissue densities in several areas including the cerebellum, precuneus, inferior and middle frontal gyri and the occipital cortex [28]. Other publications also show differences in GMV in the frontal regions like the middle and superior frontal gyri along with the cerebellum, thalamus, and inferior parietal lobule [27]. Additionally, some research present differences in cortical thickness in the frontal, parietal, and temporal lobes and within the structural volume of the cerebellum and amygdala [30]. These findings align partially with the regions responsible for classification, seeing the similarities across the temporal gyri, cerebellum, and frontal regions. These regions are responsible for processing sound, language comprehension and processing along with self-awareness and spatial processing [31] [32] [33]. However, the classifier sees new areas not frequently mentioned in research like the postcentral and fusiform gyri (Tables 3.11 and 3.12). These areas are responsible for object recognition, visual processing along with touch sensory receptive areas [34] [35] [36]. Furthermore, the saliency maps show intense highlighting in regions surrounding the hippocampus and amygdala; areas known to be structurally and functionally dimorphic among brains [3].

Comparing the saliency maps to known areas of interest for neurological diseases we see overlaps in the results. For example, major depressive disorder manifests itself within cortical areas such as the dorsal and medial prefrontal cortex, the dorsal and ventral anterior cingulate cortex, the orbital frontal cortex, and the insula predominantly [37]. In subcortical-limbic regions, structural changes are observed within the amygdala, hippocampus, and dorsomedial thalamus [37]. On the other hand, disorders such as autism spectrum disorder tend to manifest themselves within the amygdala, orbitofrontal cortex (OFC), temporoparietal cortex (TPC), and insula [38]. Overlapping these areas with tables 3.11 and 3.12, we see a connection within areas such as the frontal cortex, temporal lobules, and cingulate cortices. These areas are shown to be significant in neurological diseases and sex differences hence opening the door for potential further research into how and why they are important.

Despite showing similar regions of interest, the models trained on affined MR images hold more information on sex differences than their rigid counterpart. By performing, the affine transformations, variations within the individual organ volumes are removed and standardized. The results presented differ from standard literature since the ROI's are not found via a voxel-based morphometry, covariate regression or other traditional statistical methods, but rather by visualizing classification decisions of a deep learning model on an individual volume. Pulling all the findings together, we see that size and shape are not the only differentiating features between male and female brains but rather a fraction of the dimorphisms present.

Chapter 4: Future Works and Conclusion

The research presented in this paper opens the door to various new avenues of research. Looking into the future, this research could be furthered in several ways. First, more T1 MR images can be used for training. The UKBiobank has access to over 41,000 MR images which can be used for analysis. Second, the criteria on data selection can be widened in scope. Constricting age between 25- and 50-years old accounts for only a small percentage of the population. Increasing

the range to be between 25- and 75-years old could show a significant improvement in classification accuracy and generalizability. Third, more preprocessing should be done to the individual volumes before training. Rigid and affine transformations are good to start but it would be beneficial if the images were also skull stripped before training. In this way, the skull is removed as a feature and the model only looks for differences within the gray matter volume. And lastly, increasing the number of testing datasets and images could provide more profound insight on model generalizability and sex differences amongst a large population sample.

This thesis project examined the use of deep learning models in classifying sex from individual MR volumes in hopes to better understand sex dimorphisms present in the human brain. With a better understanding of the differences present, treatments for many neurological conditions and diseases can be tailored to a patient's specific needs rather than using a "one size fits all" approach. Many researchers oppose the view that these sex differences exist, let alone cause a substantial change to the progression and diagnosis of disease [2]. However, there are a group of researchers and studies that show these dimorphisms may be crucial in our understanding of neuroscience, neuroanatomy, and neurodegenerative disorders [3].

The research had two core aims. First, to find the optimal deep learning architecture capable of correctly predicting sex from T1 structural volumes and second, to visualize the model's classification decision to localize differentiating features of interest. Two architectures were compared for performance and predictive capabilities, a Resnet18 model and a simpler sequential CNN. Since size and shape are known differentiating features of human brains [2], the models were trained on raw and affinely transformed images. The affined transformations to the MNI305 space removed size and shape as features within the training set. Having two training sets resulted in two models for each architecture. These models were evaluated against 500 test images where the simple CNN outperformed the Resnet18 models in both scenarios averaging an accuracy over 90%. More interestingly, the classical model was able to generalize across MR field strengths, correctly classifying 1.5T and 3T images.

Visualizing the results of both classical and Resnet18 models, both architectures showed differences were present and aligned with previous research. The models showed sex dimorphisms exist in previously known areas like the frontal, temporal and precuneus gyri along with the cerebellum and thalamus. But other regions such as the cingulate, postcentral, and fusiform gyri illustrated differences not commonly mentioned in literature. These areas of the brain are responsible for proprioceptive touch and emotion processing indicating that known psychological differences manifest themselves in the structures of the brain.

In summary, this study provides a glimpse into the sexual dimorphisms present in the human brain and overall areas where these differences manifest. It analyzed and visualized classification results from differing CNN architectures trained on raw and affined MR images and chose the optimal model in terms of prediction capabilities and generalizability across datasets. This paper set the groundwork for many future studies and has veered away from the traditional statistical methods of analysis to a volumetric approach for brain MR exploration.

References

- [1] D. Joel, Z. Berman, I. Tavor, N. Wexler, O. Gaber, Y. Stein, N. Shefi, J. Pool, S. Urchs, D. S. Margulies, F. Liem, J. Hanggi, L. Jancke and Y. Assaf, "Sex beyond the genitalia: The human brain mosaic," *PNAS*, vol. 112, no. 50, pp. 15468-15473, 2015.
- [2] L. Eliot, A. Ahmed, H. Khan and J. Patel, "Dump the "dimorphism": Comprehensive synthesis of human," *Neuroscience and Biobehavioral Reviews*, 2021.
- [3] L. Cahill, "Why sex matters for neuroscience," *Nature Reviews Neuroscience*, 2006.
- [4] M. M. McCarthy, A. P. Arnold, G. F. Ball and J. D. Blaustein, "Sex Differences in the Brain: The Not So Inconvenient Truth," *Journal of Neuroscience*, vol. 32, no. 7, pp. 2241-2247, 2012.
- [5] S. J. Ritchie, "Sex Differences in the Adult Human Brain: Evidence from 5216 UK Biobank Participants," *Cerebral Cortex*, vol. 28, no. 8, pp. 2959-2975, 2018.
- [6] E. Choleris, L. A. M. Galea, F. Sohrabji and K. M. Frick, "Sex differences in the brain: Implications for behavioral and biomedical research," *Neurosci Biobehav Rev*, vol. 85, pp. 126-145, 2018.
- [7] J. M. Goldstein, L. J. Seidman, N. J. Horton, N. Makris, D. K. Kennedy, V. S. Caviness and S. V. Faraone, "Normal Sexual Dimorphism of the Adult Human Brain Assessed by In Vivo Magnetic Resonance Imaging," *Cerebral Cortex*, vol. 11, no. 6, pp. 490-497, 2001.
- [8] S. Liu, J. Seidlitz, J. D. Blumenthal, L. S. Clasen and A. Raznahan, "Integrative structural, functional, and transcriptomic analyses of sex-biased brain organization in humans," *PNAS*, vol. 117, no. 31, pp. 18788-18798, 2020.
- [9] J. Maner, J. Richey, K. Cromer, M. Mallot, C. Lejuez, T. Joiner and N. Schmidt, "Dispositional anxiety and risk-avoidant decision-making," *Personal Individ Diff*, vol. 42, pp. 665-675, 2007.
- [10] L. Betzig, "Means variances, and ranges in reproductive success: comparative evidence," *Evolution Human Behaviour*, vol. 33, pp. 309-317, 2012.
- [11] M. Spritzer, N. Solomon and D. Meikle, "Influence of scramble competition for mates upon the spatial ability of male meadow voles," *Anim Behav*, vol. 69, pp. 425-436, 2005.
- [12] L. Jacobs, S. Gaulin, D. Sherry and G. Hoffman, "Evolution of spatial cognition: sex-specific patterns of spatial behavior predict hippocampal size.," *Proc Natl Acad Sci USA*, vol. 87, pp. 6349-6352, 1990.
- [13] J. Bernal, K. Jushibar, D. S. Asfaw, S. Valverde, A. Oliver, R. Marti and X. Llado, "Deep convolutional neural networks for brain image analysis on magnetic resonance imaging: a review," *Artificial Intelligence in Medicine*, vol. 95, pp. 64-81, 2019.

- [14] Y. LeCun, L. Bottou, Y. Bengio and Patrick Haffner, "Gradient-Based Learning APplied to Document Recognition," *Proceedings of IEEE*, p. 46, 1998.
- [15] A. Krizhevsky, I. Sutskever and G. Hinton, "ImageNet Classification with Deep Learning Convolutiona Neural Networks," *NeurIPS*, p. 9, 2012.
- [16] K. He, X. Zhang, S. Ren and J. Sun, "Deep Learning for Image Recognition," *CVPR*, p. 12, 2016.
- [17] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh and D. Batra, "Grad-Cam: Visual Explanations from Deep networks vis Gradient-based Localization," *ICCV*, p. 23, 2017.
- [18] K. Simonyan, A. Vedaldi and A. Zisserman, "Deep Inside Convolutional Networks: Visualising," *Workshop at International Conference on Learning Representations*, 2013.
- [19] W. Samek, A. Binder, G. Montavon, S. Bach and K. Robert, "Evaluating the visualization of what a," *arXiv.org*, p. 13, 2015.
- [20] J. Rieke, F. Eitel, M. Weygandt, J. D. Haynes and K. Ritter, "Visualizing Convolutional Networks for," *Alzheimer's Disease Neuroimaging Initiative*, p. 9, 2018.
- [21] P. Kostelec and S. Periaswamy, "Image Registration for MRI," *Modern Signal Processing*, vol. 46, p. 24, 2003.
- [22] D. L. Collins, P. Neelin , T. M. Peters and A. C. Evans, "Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space," *J. Comput. Assist. Tomogr*, vol. 18, pp. 192-205, 1994.
- [23] J. Peelle, "jpeelle.net," Creative Commons BY-NC-SA 3.0 License, [Online]. Available: http://jpeelle.net/mri/image_processing/registration.html. [Accessed 15 January 2022].
- [24] T. E. Rolls, C. C. Huang, C. P. Lin, J. Feng and M. Joliot, "Automated anatomical labelling atlas," *Neuroimage*, p. 206, 2020.
- [25] UK Biobank Limited 2022, "ukbiobank," UK Biobank Limited 2022, [Online]. Available: ukbiobank.ac.uk.
- [26] J. Sacher, J. Neumann, H. Okon-Singer, S. Gotowiec and A. Villringer, "Sexual dimorphism in the human brain: evidence from neuroimaging," *Magnetic Resonance Imaging*, vol. 31, no. 3, pp. 366-375, 2013.
- [27] X. Zhang, M. Liang, W. Qin, B. Wan, C. Yu and D. Ming, "Gender Differences Are Encoded Differently in the Structure and Function of the Human Brain Revealed by Multimodal MRI," *Frontiers in Human Neuroscience*, vol. 14, pp. 1662-5161, 2020.
- [28] A. N. Ruigrok, G. S. Khorshidi, M.-C. Lai, S. Baron-Cohen, M. V. Lombardo, R. J. Tait and J. Suckling, "A meta-analysis of sex differences in human brain structure," *Elsevier*, vol. 39, no. 100, pp. 34-50, 2014.

- [29] X. Chen, P. S. Sachdev, W. Wen and K. J. Anstey, "Sex differences in regional gray matter in healthy individuals aged 44-48 years: a voxel-based morphometric study," *Neuroimage*, vol. 36, pp. 691-699, 2007.
- [30] O. Sangha, D. Ma, K. Popuri, J. Stocks, L. Wang and M. F. Beg, "Structural volume and cortical thickness differences between males and females in cognitively normal, cognitively impaired and Alzheimer's dementia population," *Elsevier*, vol. 106, pp. 1-11, 2021.
- [31] F. Boisgueheneuc, R. Levy, E. Volle, M. Seassau, H. Duffau, S. Kinkingnehun, Y. Samsun, S. Zhang and B. Dubois, "Functions of the left superior frontal gyrus in humans: a lesion study," *Oxford Academic*, vol. 129, no. 12, pp. 3315-3328, 2006.
- [32] A. Patel, G. M. N. R. Biso and J. B. Fowler, *Neuroanatomy, Temporal Lobe*, Florida: StatPearls Publishing, 2022.
- [33] S. Jimshelishvili and M. Dididze, *Neuroanatomy, Cerebellum*, Florida: StatPearls, 2022.
- [34] F. R. Jumah and R. H. Dossani, *Neuroanatomy, Cingulate Cortex*, Florida: StatPearls Publishing, 2021.
- [35] J. DiGuiseppi and P. Tadi, *Neuroanatomy, Postcentral Gyrus*, Florida: StatPearls Publishing, 2021.
- [36] K. S. Weiner and K. Zilles, "The anatomical and functional specialization of the fusiform gyrus," *Neuropsychologia*, vol. 83, pp. 48-62, 2016.
- [37] M. Pandya, M. Altinay, D. A. Malone and A. Anand, "Where in the brain is depression?," *Curr Psychiatry Rep*, vol. 14, pp. 634-642, 2012.
- [38] C. S. Weston, "Four Social Brain Regions, Their Dysfunctions, and Sequelae, Extensively Explain Autism Spectrum Disorder Symptomatology," *Brain Sci*, vol. 9, p. 130, 2019.
- [39] D. Ma, K. Popuri, M. Bhalla, O. Sangha, D. Lu, J. Cao, C. Jacova, L. Wang and M. F. Beg, "Quantitative assessment of field strength, total intracranial," *Human Brain Mapp*, vol. 40, pp. 1507-1527, 2018.
- [40] N. J. T. Brian B. Avants, "NeuroBattery, Smashing through the boundaries - of current neuroimage analysis limitations," [Online]. Available: <http://jeffduda.github.io/NeuroBattery/>. [Accessed 21 7 2021].
- [41] M. Notter, "NiPype Beginners Guide," 10 June 2021. [Online]. Available: <https://miykael.github.io/nipype-beginner-s-guide/neuroimaging.html>. [Accessed 31 December 2021].
- [42] M. Ingalhalikar, A. Smith, D. Parker, T. D. Satterthwaite, M. A. Elliott, K. Ruparel, H. Hakonarson, R. E. Gur and R. Verma, "Sex differences in the structural connectome of the human brain," *PNAS*, vol. 111, pp. 823-828, 2013.

