

Low Field MRI Deep Learning Framework for Non-Hospitalizing Early Detection and Characterization of Alzheimer's Disease Pathology

Aarav Minocha

Great Neck South High School, Reteena
New York, United States
aminochal@student.gn.k12.ny.us

Seungyong Yang

North London Collegiate School Jeju, Reteena
Seoul, South Korea
syyang26@pupils.nlcsjeju.kr

Jainish H. Patel

Lake Minneola High School, Reteena
Florida, United States
jainish@reteena.org

Ivan Ma

Winston Churchill High School, Reteena
Maryland, United States
159606@mcpsmd.net

Abstract—Alzheimer's Disease (AD) is the most common type of neurodegenerative disease and it significantly disrupts brain function, leading to memory loss, cognitive decline, and more. Currently, there is no effective cure or treatment for AD, making early diagnosis critical so preventative actions can be taken before large-scale degradation occurs. The most common diagnostic techniques involve various neuroimaging modalities, which provide structural and functional information about the brain. However, many of these techniques are extremely expensive, have long scan times, and require access to hospitals equipped with high-grade equipment such as MRI, PET, and CT scanners.

To address these challenges, our team focused on the emerging use of Low-Field MRI scanners, which are cheaper, portable, and have much shorter scan times compared to traditional MRI scanners due to their lower magnetic field strength. However, these scanners provide significantly less information than expensive hospital-grade MRI scanners. Therefore, the goal of our study was to develop an image enhancement framework capable of effectively segmenting brain regions associated with AD (such as the amygdala, hippocampus, and ventricles) and automatically diagnosing AD using Low-Field MRI scans.

In our approach, we first created a framework to generate synthetic Low-Field (LF) MRI scans using a Fourier Transformation framework. We then developed a deep learning framework to enhance these LF MRI scans by utilizing a SRCNN based architecture and UNET++ models, which performed super-resolution and segmentation, respectively, to obtain volumetric information. Our SRCNN model achieved a Mean Squared Error of 214.54, a Peak Signal-to-Noise Ratio of 31.2 dB, and a Structural Similarity Index Measurement of 0.82. The UNET++ model achieved an accuracy of 96.3 percent, a precision of 89.3 percent, a recall of 85.6 percent, and a Dice score of 0.93.

Finally, we utilized the volumetric information to classify the scans as either AD or normal using a soft voting majority framework, where the final diagnosis was based on the majority consensus of models. This soft voting framework consisted of three machine learning (ML) models: Linear Regression, Support Vector Machine, and Multilayer Perceptron. The overall framework achieved an accuracy of 0.96, a precision of 0.96, and a recall of 0.98.

Index Terms—computational neuroscience, deep learning, neu-

roimaging, neurodegenerative disease, Alzheimer's Disease

I. INTRODUCTION

Alzheimer's Disease (AD) is the most common type of neurodegenerative disease and the fifth leading cause of death among individuals aged 65 and up. The disease significantly disrupts the functioning of the brain, leading to memory loss, cognitive decline, and more. AD is characterized by a gradual onset, with manifestations including memory loss, cognitive decline, and a variety of psychological and neurological issues [3]. Once patients are diagnosed with AD, there are no treatments that can effectively cure the disease, putting emphasis on early diagnosis during the prodromal stage of the disease, known as mild cognitive impairment (MCI), resulting in slow actions being taken. Traditional diagnostic techniques rely on the presence of significant biomarkers in cognitive mental assessments and/or neuroimaging modalities for diagnosis. Magnetic Resonance Imaging (MRI) is a neuroimaging technique that is used to obtain in-depth information about the anatomical structures of the brain. It provides significant information about the structural changes within the brain as a result of AD, including gray and white matter, ventricles, and important regions such as the amygdala and hippocampus. Segmentation of the various regions of the brain within the images provided by MRI brain scans is a common technique used by radiologists to analyze the volumetric changes within the structure of the brain [4]. However, many times once these apparent clinical symptoms become noticeable, patients are already in the severe late stages of AD, where no treatments can slow down or stop disease progression [8]. An early diagnosis usually requires meticulous analysis of patient data by radiologists and clinical experts. Although this technique is effective, it is time-consuming, tedious, and subject to external bias. The use of deep learning to improve the segmentation of brain MRI scans has been found to be faster, more effective,

and more accurate. However, one major issue with common imaging techniques such as MRI is that they have very long scan times and require patients to have access to a hospital with expensive equipment. An emerging technique within current research to address these problems is the use of low-field (LF) MRI scanners, which are composed of 50-100mT as opposed to common scanners that are 1-3T. This allows these scanners to be significantly cheaper, smaller, and have much shorter scan times compared to traditional MRI scanners [2]. These scanners can be transported to patients' homes and allow for an initial detection of AD without requiring patients to have access to a hospital with expensive equipment. However, these scanners provide less information compared to expensive hospital-grade MRI scanners. This can lead to less detailed images and the potential to miss subtle abnormalities, even though they are more accessible. Therefore, it is necessary to improve the structural information these scans provide for a diagnosis.

A. Deep Learning techniques for data extraction

In recent years, deep learning neural networks have been used effectively in many computer vision tasks, because of its superior feature learning ability [7]. One network which has been researched extensively for super resolution and transformation tasks is the SRCNN [8]. SRCNN excels in super-resolution by effectively reconstructing high-resolution details from low-resolution images through its expansive feature mapping pathway where the spatial resolution of the image is increased by utilizing learned complex patterns from the original image to hypothesize on every individual pixel. The UNET architecture uses a combination of encoder-decoder blocks and convolutional layers for medical segmentation related tasks [9]. The encoder block captures the semantic landscape and extracts the features while the decoder block reconstructs the image from the extracted features and upsamples to enhance spatial resolution. For segmentation, its encoder-decoder structure, complemented by skip connections, ensures precise delineation of complex structures within images, where pixels are separated into regions and labeled based on the region. The SRCNN performs super resolution on a 3D LF brain MRI and the UNET segments the super-resolved scan for volumetric data extraction for regions of the brain.

B. ML Algorithms for diagnosis

The first model used is Support Vector Machine (SVM), a linear algorithm that is particularly attractive for AD diagnosis because it can easily separate data into multiple classes and is resistant to overfitting [5]. The second model used is Linear regression, a linear algorithm that performs classification tasks by predicting the probability that a data instance belongs to a particular class. The final model used is Multilayer Perceptrons (MLPs), a type of artificial neural network that is typically used for classification and pattern recognition and is suited for large sets of data in various modalities [11]. Soft Voting is an ML technique that weights the predictions of multiple models to produce a final output. This final output makes a

more accurate prediction than any single base model. The final prediction is made by taking a weighted sum of the probabilities obtained from each base model, with the weights assigned according to the confidence of each model. This allows for a more accurate prediction of the final classes (normal, MCI, and AD)

II. METHODOLOGY

A. Scan Acquisition

600 de-identified expertly segmented 3D High Field (HF) MRI scans were obtained from more than 200 patients ranging from the ages of 50-80 from the OASIS 3 Brains Dataset. Scans were taken on T1w and T2w scanners with Tesla strengths of either 1.5T and 3T from various hospitals and diagnosis centers around the United States. These scans were automatically segmented for various brain regions using FreeSurfer software before being validated and modified by a radiologist to increase the accuracy of the segmentations. Each patient within the dataset was also diagnosed for AD by a specialist and underwent multiple MR sessions.

B. Low Field MRI Dataset Generation

To improve the resolution of LF scans using deep learning, a dataset must be created containing corresponding HF and LF scans for the same patient. No current dataset exists, so synthetic LF MRI scans were created by removing fine-grained details and blurring the image. A Fourier transformation converts the image from the spatial domain into the frequency domain, with higher frequencies corresponding to finer details. Truncation selectively removes higher frequency components, leaving only lower frequency/detail components. An inverse Fourier transformation reconstructs the lower-resolution image back to the spatial domain, resulting in a blurred image with lower detail. These new transformed images have a lower peak-signal-to-noise ratio (PSNR) and contrast-to-noise ratio compared to HF MRI scans. This transformation was performed on each individual scan to create a dataset consisting of HF scans and corresponding LF scans for each MR session of a patient. The new dataset was split into 80 percent training and 20 percent testing for the deep learning image enhancement model framework.

C. Scan Preprocessing

Raw scans were preprocessed before being passed into the deep learning framework for image enhancement in order to increase computational efficiency. This is because the raw scans have high dimensionality (up to 256 x 256) and a large array of varying pixel intensities. A preprocessing pipeline is necessary to reduce computational strain and remove unnecessary information within scans.

First, the image intensities for all pixels in the image were normalized between values of 0 and 1 [1]. Then, the skull and neck were stripped from the scan to reduce dimensionality using FSL software [1]. Scans were then normalized to the same size (200 x 160), due to the varying dimensionality of the scans after skull stripping, using bilinear interpolation,

which hypothesized on all missing pixels by averaging the four nearest pixels.

D. Deep Learning Image Enhancement

The first step of the image enhancement pipeline was to improve LF scans through image super resolution to improve the quality and detail of LF MRI scans, making LF scans comparable to their HF counterparts. A Super Resolution Convolutional Neural Network (SRCNN) was used to achieve image super resolution [13]. SRCNN operates through a three-layered architecture designed for enhancing the resolution of images. The first layer, known as the patch extraction and representation layer, processes the input low-resolution image by extracting overlapping patches. These patches are then analyzed to identify and represent inherent features such as color, shape, and patterns, effectively converting spatial details into high-dimensional feature vectors. The second layer performs a crucial non-linear mapping, transforming these extracted feature vectors into a refined feature space that embodies the characteristics of a higher resolution. This transformation is key to preparing the data for the final step of super-resolution. In the last stage, the reconstruction layer takes over, synthesizing these enhanced features to reconstruct the final super-resolved image. This layer meticulously combines the processed features to produce a detailed, high-resolution output from the initially low-resolution input, showcasing the model's ability to infer and amplify the intricate details necessary for high-quality image super-resolution. By training SRCNN with pairs of low and high-resolution images, it learns to predict the missing high-frequency details in low-resolution images, significantly enhancing the quality and utility of images produced by LF MRI scans. This technology ensures that despite the lower signal-to-noise ratio and contrast-to-noise ratio inherent in LF MRI images, the output closely approximates that of higher field strength MRI images in terms of detail and clarity, making it viable for accurate diagnosis. The SRCNN was trained on 80 epochs with a batch size of 8 scans per iteration. An Adam's optimizer was used to minimize the loss in the loss function, being mean squared error (MSE). Calculates the average of the squared differences between the actual pixel values y_i and the predicted pixel values \hat{y}_i .

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (1)$$

The second step of the image enhancement pipeline was to segment out the three regions of interest (ROIs) linked to AD (Amygdala, Hippocampus, and Ventricles) from the super resolved scans. UNET++ was used to segment out the ROIs because of its ability to delineate different objects from an image. The UNET architecture specializes in the accurate segmentation of medical images into various labels [14]. The UNET++ architecture is a convolutional network that uses an encoding path and a symmetric decoding path to understand complex features within an image. The encoding pathway involves repeated application of two 3x3 convolutional layers, rectified linear unit activation, and a 2x2 max-

pooling operation to downsample the image. This reduces dimensions and increases feature depth to understand complex patterns. The decoder block reconstructs the image using learned features, upsampling convolutional layers to enhance spatial dimensions. Skip connections connect the encoder and decoder, ensuring precise delineation of complex features onto the image. This structure allows the UNET++ to excel in differentiating between tissue types and segmenting ROIs onto super-resolved MRI scans. The model was trained on 100 epochs with a 16-scan batch size each iteration, using the Adam optimizer to minimize loss in the loss function, specifically the Dice loss.

$$\text{Dice Loss} = 1 - \frac{2 \sum_{i=1}^n y_i \hat{y}_i}{\sum_{i=1}^n y_i + \sum_{i=1}^n \hat{y}_i} \quad (2)$$

E. ML Diagnosis

The three ROIs (Amygdala, Hippocampus, and Ventricles) are analyzed using volumetric data in a soft voting machine learning framework for automatic AD diagnosis, with the final diagnosis determined by consensus. The first model was a linear regression which plotted the data in a 3D feature space according to the volume of its three ROIs and created the optimal linear hyperplane to separate the data into the two classes, AD and non-AD [10]. This allowed for a clear separation between the data according to the 3D plane of best fit through its plotted features. The second model utilized was a support vector machine (SVM), which operated similarly to a linear regression by plotting the data in a 3D plane, and creating a hyperplane separating the data into two classes [11]. The difference with a SVM is that it is not confined to a linear separation and can more accurately separate the data into the two classes by creating a non-linear hyperplane between the data points. The final model utilized was a multilayer perceptron (MLP), which was a type of artificial neural network capable of learning how to separate the data into classes through complex learned patterns and relationships within the data [12]. This was the most effective model utilized because it could learn from the nuanced volumetric information of the ROIs more effectively, allowing it to be more accurate in its diagnosis.

A majority voting consensus between the models was utilized for the final diagnosis. This took into account each model's individual diagnosis for a patient and outputted the diagnosis which the majority of the models agreed upon. This provided the most accurate final diagnosis as it took into account the output from all three models.

III. RESULTS

A. High Field Dataset

The dataset was constructed from the OASIS3 Brains dataset consisting of more than 600 High Field (1.5T-5T) scans from 200 distinct patients. An example 3D scan is shown in Figure 1 for all three scan views (axial, coronal, and sagittal). The dataset also contained expert validated segmentations which we utilized to train the image segmentation model.




Fig. 1: High Field Scan



Fig. 3: Synthetic LF to HF

B. Synthetic Low Field Scan Generation

Utilizing a blurring process we were able to synthetically generate LF Scans to address the inability to access real ones. This process consisted of a fourier transformation, where the image was converted to the frequency domain, and then truncation where the high frequency signals, corresponding to fine details in the image were removed. Finally, an inverse fourier transformation was performed to convert the scan back into the image domain where the image was blurred, resembling a simulated LF scan. To validate our synthetic scan generation methodology we measured the PSNR of real LF scans from the M4Raw dataset [16]. We compared across 50 MRI scans and achieved 27.6 dB indicating an almost similar amount of noise found between the scans.

C. Preprocessing

After generating synthetic LF MRI scans, we preprocessed the scans to normalize them to reduce computational strain when training the models. The preprocessing method aimed to reduce computational strain by utilizing intensity normalization between 0 to 1, skull stripping using FSL software, clipping of edges and 3D patching into 64 x 64 x 64 segments. This allowed for increased computational efficiency and more accurate image enhancement. The output of the preprocessing transformation for a LF scan is shown in Figure 2, where patches of the images would be created for image enhancement.



Fig. 2: Preprocessed Synthetic LF Scan

D. Super Resolution Model

Super Resolution is the first part of the image enhancement process. Super Resolution enhanced the detail and information within the scan while also reducing blurring. Super resolution was performed by a SRCNN with 3 2D convolutional layers. The output of the SRCNN was a super resolve scan, with less blurring and greater specific details present in the image. Figure 3 displays a super resolved scan for a patient's LF scan.

E. Super Resolution Model Metrics

To evaluate efficiency of the SRCNN model the Mean Squared Error (MSE), Peak Signal to Noise Ratio (PSNR), and Structural Similarity Index Measurement (SSIM) were calculated.

	MSE	PSNR (dB)	SSIM
SRCNN	214.54	31.2 dB	0.82

TABLE I: Metrics for SRCNN Model

F. Scan Segmentation Model

The UNET++, a newer version of the UNET architecture, uses convolutional layers and max pooling layers to learn segment patterns in 3D scans. It reduces dimensionality and displays segments back onto the image, increasing dimensionality using skip connections. This improves segmentation accuracy for 3D scans by overlaying learned features more frequently. UNET++ also uses the volume of three ROIs as features in ML diagnosis models to obtain a final diagnosis of normal control or Alzheimer's disease. The output of a segmentation of the UNET++ is shown in Figure 4.




Fig. 4: Ventricle Segmentations

G. Scan Segmentation Metrics

To evaluate the efficacy of the UNET++ at performing segmentation of ROIs from super resolved scans, the accuracy, precision, recall, and dice score were calculated. The UNET++ model is compared to another study segmenting LF scans as shown in Table II.

H. Soft Voting Metrics

The Soft Voting classifier, consists of Support Vector Machine, Linear Regression, and Multilayer Perceptron achieved an accuracy of 96 percent, precision of 96 percent, and a recall of 98 percent as shown in Figure 5.

	Accuracy	Precision	Recall	Dice Score
U-Net (Kumar et al., 2023) [6]	91.5	88.4	87.1	0.90
	Accuracy	Precision	Recall	Dice Score
Proposed U-Net++	96.3	89.3	85.6	0.93

TABLE II: Comparison of U-Net and Proposed U-Net++

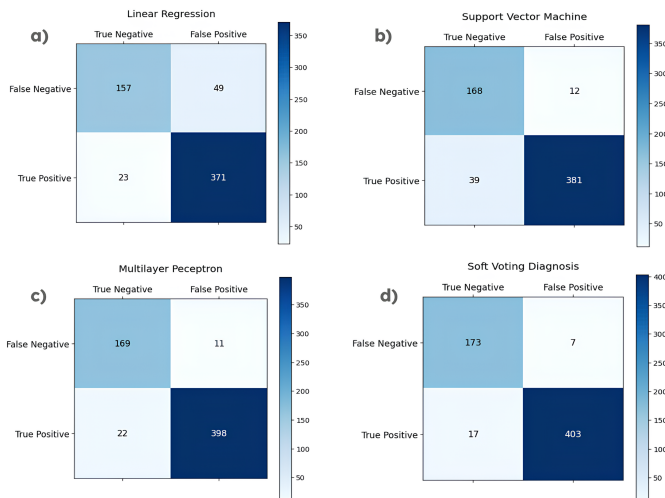


Fig. 5: Soft Voting Results

IV. DISCUSSION

A. Pipeline

Prior studies have not conducted super resolution on LF MRI scans; therefore, we utilized three benchmark metrics: MSE, SSIM, and PSNR. MSE and SSIM were calculated to measure the difference between the LF and HF scans, while PSNR was used to measure the noise in the image. Our adoption of the SRCNN resulted in a decrease in MSE and an increase in SSIM, along with PSNR. After these images were super resolved the UNET++ created volumetric segmentations of comparable accuracies to a previous study segmenting HF images with UNET (Kumar et al., 2023) [6]. The framework enhanced LF image quality through super resolution. Following this, UNET++ produced volumetric segmentations comparable to a prior study on HF images, which reported a Dice score of 0.87 for AD-related regions (Kumar et al., 2023) [6]. While this score may be inflated as it excluded ventricles, another study also reported a Dice score of 0.87 for ventricle segmentation using UNET. These findings suggest that segmenting with UNET after super resolution improves accuracy, and additional skip connections further enhance segmentation quality. A final diagnosis is obtained using a soft voting framework utilizing three different ML algorithms. A previous study diagnosing AD through ML with MRI scans obtained an accuracy of 95 percent, precision of 96 percent, and recall of 96 percent when creating binary classification diagnosis of either AD or non-AD with SVM [15]. Utilizing linear regression, SVM, and a MLP in a soft-voting framework resulted in a higher accuracy and recall rate.

B. Limitations

The study presents a deep learning framework for diagnosing neurodegenerative diseases like Alzheimer's Disease (AD) using LF MRIs. However, it has limitations. The framework was trained on synthetically generated LF scans, which may not be suitable for real-world scenarios. Additionally, the dataset used for training was mainly from affluent areas of the United States, leaving many regions out. The dataset also only included data for patients aged 50-81, which may not be suitable for patients seeking an AD diagnosis earlier. A more generalized dataset could make the framework more accessible to a broader population.

REFERENCES

- [1] Z. Fan, et al., "U-Net Based Analysis of MRI for Alzheimer's Disease Diagnosis," Apr. 2021, pp. 13587–13599. Available: <https://doi.org/10.1007/s00521-021-05983-y>
- [2] Kimberly, W. T., Sorby-Adams, A. J., Webb, A. G., Wu, E. X., Beekman, R., Bowry, R., Schiff, S. J., de Havenon, A., Shen, F. X., Sze, G., Schaefer, P., Iglesias, J. E., Rosen, M. S., & Sheth, K. N. (2023). Brain imaging with portable low-field MRI. *Nature reviews bioengineering*, 1(9), 617–630. <https://doi.org/10.1038/s44222-023-00086-w>
- [3] A. Y. Yao, P. J. Halloran, G. Y. Ge, N. Singh, J. Zhou, J. Galske, W. He, R. Yan, and X. Hu, "Bace1 deletion in the adult reverses epileptiform activity and sleep-wake disturbances in ad mice," *The Journal of Neuroscience*. [Online]. Available: <https://www.jneurosci.org/>
- [4] M. Smits, "MRI biomarkers in neuro-oncology," *Nature Reviews Neurology*, vol. 17, 2021, pp. 486–500.
- [5] Z. Zhao, J. H. Chuah, K. W. Lai, C. O. Chow, M. Gochoo, S. Dhanalakshmi, N. Wang, W. Bao, and X. Wu, "Conventional ML and deep learning in Alzheimer's disease diagnosis using neuroimaging: A review," *Frontiers in Computational Neuroscience*, Feb. 2023.
- [6] Kumar, G. M., Parthasarathy, E. (2022). Development of an enhanced U-Net model for brain tumor segmentation with optimized architecture. *Biomedical Signal Processing and Control*, 81, 104427. <https://doi.org/10.1016/j.bspc.2022.104427>
- [7] R. A. Hazarika, A. K. Maji, R. Syiem, S. N. Sur, and D. Kandar, "Hippocampus Segmentation Using U-Net Convolutional Network from Brain Magnetic Resonance Imaging (MRI)," *Journal of Digital Imaging*, vol. 35, no. 4, Aug. 2022, pp. 893-909.
- [8] Z. Iqbal, D. Nguyen, G. Hangel, S. Motyka, W. Bogner, and S. Jiang, "Super-Resolution 1H Magnetic Resonance Spectroscopic Imaging Utilizing Deep Learning," *Frontiers in Oncology*, vol. 9, 2019, p. 1010.
- [9] Lu, H., She, Y., Tie, J., Xu, S. (2022). Half-UNet: a simplified U-Net architecture for medical image segmentation. *Frontiers in Neuroinformatics*, 16. <https://doi.org/10.3389/fninf.2022.911679>
- [10] A. Schneider, G. Hommel, and M. Blettner, "Linear regression analysis: part 14 of a series on evaluation of scientific publications," *Dtsch Arztebl Int*, vol. 107, no. 44, 2010, pp. 776-82.
- [11] Y. Zhang, "Support Vector Machine Classification Algorithm and Its Application," 2012, pp. 179-186.
- [12] T. Bikku, "Multi-layered deep learning perceptron approach for health risk prediction," *Journal of Big Data*, vol. 7, 2020, p. 50.
- [13] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *arXiv.org*, July 2015.
- [14] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," *arXiv.org*, May 2015.
- [15] Alatrany, A. S., Khan, W., Hussain, A., Kolivand, H., Al-Jumeily, D. (2024). An explainable machine learning approach for Alzheimer's disease classification. *Scientific reports*, 14(1), 2637. <https://doi.org/10.1038/s41598-024-51985-w>
- [16] Lyu, M., Mei, L., Huang, S. et al. M4Raw: A multi-contrast, multi-repetition, multi-channel MRI k-space dataset for low-field MRI research. *Sci Data* 10, 264 (2023). <https://doi.org/10.1038/s41597-023-02181-4>