Super-Resolution of Brain MRI via U-Net Architecture

Aryan Kalluvila
Athinoula Martinos Center, Harvard Medical School, Charleston, MA, Hartford Union High School, Hartford, WI

***Abstract*—This paper proposes a U-Net-based deep learning architecture for the task of super-resolution of lower resolution brain magnetic resonance images (MRI). The proposed system, called MRI-Net, is designed to learn the mapping between low-resolution and high-resolution MRI images. The system is trained using 50-800 2D MRI scans, depending on the architecture, and is evaluated using peak signal-to-noise ratio (PSNR) metrics on 10 randomly selected images. The proposed U-Net architecture outperforms current state-of-the-art networks in terms of PSNR when evaluated with a 3 x 3 resolution downsampling index. The system's ability to super-resolve MRI scans has the potential to enable physicians to detect pathologies better and perform a wider range of applications. The symmetrical downsampling pipeline used in this study allows for generically representing low-resolution MRI scans to highlight proof of concept for the U-Net-based approach. The system is implemented on PyTorch 1.9.0 with NVIDIA GPU processing to speed up training time. U-Net is a promising tool for medical applications in MRI, which can provide accurate and high-quality images for better diagnoses and treatment plans. The proposed approach has the potential to reduce the costs associated with high-resolution MRI scans by providing a solution for enhancing the image quality of low-resolution scans.**

Keywords—MRI; U-Net; Super-Resolution; PyTorch; SRCNN; SR-GAN; Deep Learning; GPU

I. INTRODUCTION

Medical imaging has been an essential tool in the diagnosis and treatment of various diseases. Among the different types of medical imaging, magnetic resonance imaging (MRI) has been widely used because it provides detailed images of the internal structures of the body. However, the resolution of MRI images is limited by several factors, such as the hardware used and the acquisition parameters, which can affect diagnostic accuracy and make it challenging to identify subtle anatomical features.

To address this limitation, researchers have been exploring the use of super-resolution techniques to enhance the resolution of MRI images [1]. Super-resolution is a computational method that enhances the resolution of an image beyond the physical limits of the imaging hardware. The primary objective of this approach is to generate high-resolution images that can provide more detailed information about the anatomy and pathology of the imaged area.

Recent advancements in deep learning-based super-resolution techniques for MRI have shown promising results in generating high-resolution MRI images with improved details and contrast [3]. These techniques use advanced machine learning algorithms to learn the mapping between low-resolution and high-resolution images. High-resolution MRI images are particularly important in the diagnosis and treatment of neurological disorders, where small changes in the brain's anatomy can have significant implications for patient outcomes. For example, in the diagnosis of brain tumors, high-resolution MRI images can help to accurately identify the location, size, and shape of the tumor, which is critical for surgical planning and treatment.

The use of super-resolution techniques in MRI has the potential to revolutionize the field of medical imaging, leading to significant improvements in medical diagnosis and treatment. With the increasing availability of large datasets and the progress of machine learning algorithms, there is enormous potential for further advancements in this field. However, the effectiveness of super-resolution techniques in enhancing MRI images depends on several factors, including the choice of algorithms and evaluation metrics. In this context, several deep learning-based super-resolution algorithms have been developed and tested to improve the resolution of MRI images. The choice of the appropriate algorithm depends on several factors, including the quality of the input data, the complexity of the target structures, and the available computational resources. Additionally, it is crucial to use appropriate evaluation metrics to assess the effectiveness of these algorithms in improving the resolution of MRI images.

In this paper, we aim to evaluate the effectiveness of super-resolution techniques for enhancing MRI images, with a specific focus on the U-Net algorithm. One of the strengths of our approach is the utilization of the Mean Squared Logarithmic Error (MSLE) as the main loss function, which enables accurate reconstruction of high-resolution MRI images. We compare the performance of the U-Net algorithm against four other commonly used networks in terms of generating high-quality images. To assess the quality of the enhanced images, we employ the widely accepted metric, Peak Signal-to-Noise Ratio (PSNR). Notably, our results demonstrate that the U-Net algorithm surpasses the other networks, yielding higher average PSNR values. These findings have significant implications for the diagnosis and treatment of neurological disorders, as enhanced MRI images obtained through the U-Net algorithm can provide improved insights and precision in medical imaging. The advancements showcased in this research have the potential to make a substantial impact on the field of medical imaging as a whole.

II. LITERATURE REVIEW

Super-resolution of MRI is a critical task in the medical field that involves increasing the resolution of magnetic resonance images to improve their quality and accuracy. Many researchers have proposed various techniques to achieve super-resolution of MRI, including deep learning, image registration, and image fusion. In this literature review, we will summarize some of the recent research works on super-resolution of MRI.

Li et al. [1] proposed a critic-guided framework for super-resolution of low-resolution MRI scans. In clinical practice, vast quantities of MRI scans are routinely acquired but are of sub-optimal quality for precision medicine, computational diagnostics, and neuroimaging research. To address this limitation, the authors utilized feature-importance and self-attention methods in their model to improve interpretability. Their framework was evaluated on paired low- and high-resolution MRI scans from various manufacturers and was shown to produce qualitatively faithful results to ground-truth scans with high accuracy (PSNR = 35.39; MAE = 3.78E−3; NMSE = 4.32E−10; SSIM = 0.9852).

Ottesen et al. [2] proposed a densely connected cascading deep learning reconstruction framework to improve accelerated MRI reconstruction. The authors modified a cascading deep learning reconstruction framework by incorporating three architectural modifications, namely input-level dense connections, an improved deep learning sub-network, and long-range skip-connections. The proposed framework, called the Densely Interconnected Residual Cascading Network (DIRCN), was evaluated on the NYU fastMRI neuro dataset with an end-to-end scheme at four- and eightfold acceleration. The authors performed an ablation study where they trained five model configurations and evaluated them based on the structural similarity index measure (SSIM), normalized mean square error (NMSE), and peak signal to noise ratio (PSNR). The results showed that the proposed DIRCN with all three modifications achieved an SSIM improvement of 8% and 11%, a NMSE improvement of 14% and 23%, and a PSNR improvement of 2% and 3% for four- and eightfold acceleration, respectively.

Similarly, Qiu, Wang, and Guo [3] proposed a novel deep learning-based approach for super-resolution of MRI, which utilizes a generative adversarial network (GAN) to generate high-resolution MRI images from low-resolution images. Due to limitations in hardware, scan time, and throughput, obtaining high-quality MR images can be a challenging task in clinical settings. Therefore, the authors aimed to use a super-resolution approach to enhance the image quality without requiring any hardware upgrades. In that paper, they proposed an ensemble learning and deep learning framework for MR image super-resolution. To create their framework, the authors first enlarged low resolution images using five commonly used super-resolution algorithms, resulting in differentially enlarged image datasets with complementary priors. Then, they trained a generative adversarial network (GAN) with each dataset to generate super-resolution MR images. Finally, they used a convolutional neural network for ensemble learning, which synergized the outputs of the GANs to produce the final MR super-resolution images.

De Leeuw den Bouter et al. [4] highlighted the potential of low-field MRI scanners to make MRI technology more accessible globally due to their significantly lower cost compared to high-field counterparts. However, images acquired using low-field MRI scanners tend to be of relatively low resolution, which limits their clinical utility. To address this limitation, the authors presented a deep learning-based approach to transform low-resolution low-field MR images into high-resolution ones. They trained a convolutional neural network to carry out single image super-resolution reconstruction using pairs of noisy low-resolution images and their noise-free high-resolution counterparts obtained from the NYU fastMRI database. The trained network was subsequently applied to noisy images acquired using a low-field MRI scanner, producing sharp super-resolution images with most of the high-frequency components recovered. The authors demonstrated the potential of a deep learning-based approach to increase the resolution of low-field MR images.

Wang et al. [5] proposed a CNN-based multi-scale attention network (MAN) to improve the performance of convolutional super-resolution (SR) networks. While convolutional neural networks can compete with transformer-based methods in many high-level computer vision tasks, transformers with self-attention still dominate the low-level vision, including the super-resolution task. The authors exploit large kernel decomposition and attention mechanisms in their design. The proposed MAN consists of multi-scale large kernel attention (MLKA) and a gated spatial attention unit (GSAU). Within the MLKA, the authors rectify LKA with multi-scale and gate schemes to obtain the abundant attention map at various granularity levels. This approach jointly aggregates global and local information and avoids potential blocking artifacts. In GSAU, a gate mechanism and spatial attention are integrated to remove the unnecessary linear layer and aggregate informative spatial context. The authors evaluate MAN with multiple complexities by simply stacking different numbers of MLKA and GSAU. Experimental results illustrate that their MAN can achieve varied trade-offs between state-of-the-art performance and computations.

Bahrami et al. [6] proposed a novel method for predicting high-resolution 7T-like MR images from low-resolution 3T MR images. The predicted 7T-like MR images demonstrate higher spatial resolution compared to 3T MR images, as well as prediction results obtained using other comparison methods. The authors suggest that such high-quality 7T-like MR images could better facilitate disease diagnosis and intervention. This paper demonstrates proof of concept for reconstruction in even high-resolution MRI dynamics.

Koonjoo et al's paper introduces AUTOMAP, a deep learning method for improving image quality in low-field MRI systems [10]. AUTOMAP outperforms traditional Fourier reconstruction and two contemporary denoising algorithms, reducing noise and artifacts in the reconstructed images. It achieves substantial signal-to-noise ratio gains for both human brain and plant root data, demonstrating the potential of deep learning in enhancing image quality in low-field MRI. This approach contributes to advancing resolution and image quality in low-field MRI applications.

The U-Net architecture outperforms other techniques [9] in medical image super-resolution of brain MRI due to its dedicated design for image segmentation tasks and effective feature extraction. Its skip connections enable the preservation and utilization of both high-level and low-level features, resulting in enhanced resolution. Unlike other architectures, such as the SRCNN, GAN-based approaches, or multi-scale attention networks, the U-Net consistently achieves superior resolution improvement. Its ability to capture fine details and preserve structural information makes it the preferred choice for medical image super-resolution tasks.

III. Methodology

A. Downsampling Pipeline

In order to accurately simulate the low-resolution scans typically obtained from lower field strength MRI scanners, we employed a symmetrical down sampling pipeline, as depicted in Fig. 1. This pipeline involved reducing each dimension of the scanner by a factor of three, replicating the effects of decreased resolution. By implementing this downsampling technique, we were able to mimic the conditions of low field strength and lower-quality scanners commonly associated with compromised image resolution and reduced overall image quality. To further ensure the authenticity of the simulated low-resolution scans, we applied bilinear interpolation, a widely adopted interpolation method, to generate the corrupted scans. This approach effectively captures the characteristic imperfections and limitations of lower field strength MRI scanners, providing a reliable basis for evaluating the performance and effectiveness of our super-resolution techniques in enhancing the quality and resolution of these low-resolution MRI images.

In order to enhance the computational efficiency of the U-Net architecture, we employed a technique known as residual learning. Instead of directly generating the complete high-resolution scan, our U-Net model was trained to focus on learning the difference between the high-resolution scan and the bilinear interpolated output. By utilizing this residual learning approach, the model became more adept at capturing the fine details and nuances present in the high-resolution image that may be lost during the bilinear interpolation process.

During the inference stage, the U-Net model would generate the residual, which represented the additional information needed to transform the interpolated scan into a super-resolved MRI scan. This residual was then added to the bilinear interpolated scan, resulting in the creation of a high-resolution image with enhanced details and improved quality. This approach not only improved the computational efficiency of the U-Net architecture but also ensured that the generated super-resolved MRI scan closely resembled the original high-resolution scan by effectively compensating for the limitations of the bilinear interpolation.

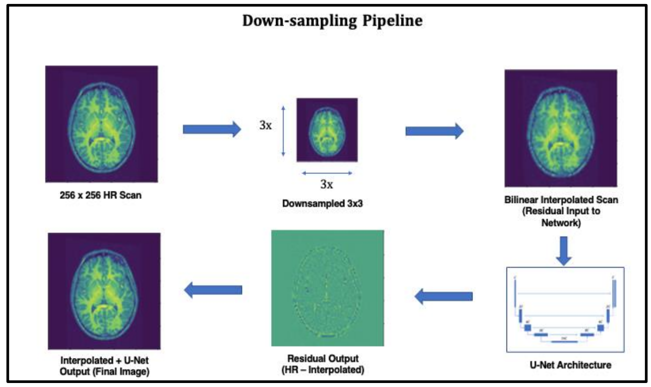

Fig. 1. Downsampling flow chart to create corrupted MRI scans.

B. U-Net Architecture

We used a U-Net architecture as denoted by Fig. 2, as the main super-resolution algorithm to improve the resolution of brain MRI. The U-Net architecture is a type of deep learning neural network that is particularly well-suited for image segmentation tasks, which involve dividing an image into multiple segments to identify specific structures or features within the image.

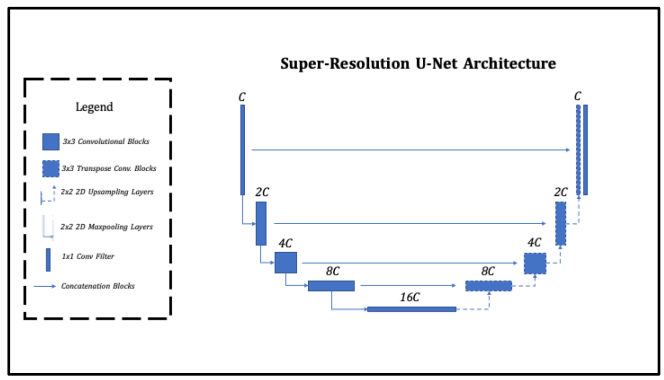

Fig. 2. U-Net architecture for super-resolution task [7].

The U-Net architecture is specifically designed for biomedical image analysis, making it an effective choice for super-resolution of brain MRI. The architecture consists of an encoder, which gradually reduces the resolution of the input image, and a decoder, which gradually increases the resolution of the image to produce the final high-resolution output. The encoder and decoder are connected by a bottleneck layer that contains information about the original image, allowing for precise reconstruction of the high-resolution output.

Compared to other deep learning architectures, such as fully convolutional networks (FCNs) or residual networks (ResNets), U-Net has several advantages for super-resolution of brain MRI. First, the U-Net architecture allows for the preservation of fine details, which is important for identifying subtle anatomical features in MRI images. Second, U-Net is less prone to overfitting, a common problem in deep learning models, as it contains skip connections that enable the model to learn from features at multiple scales. Finally, U-Net is computationally efficient, allowing for faster training and inference times compared to other architectures.

C. Evaluation Metrics

When evaluating the performance of image super-resolution techniques for MRI, several metrics are commonly used, including Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM). While both metrics are useful for evaluating image quality, PSNR is generally considered to be more important than SSIM in the context of super-resolution for MRI. PSNR is a widely used metric that measures the quality of an image by calculating the ratio of the peak signal power to the mean squared error of the image. Higher PSNR values indicate better image quality, with a perfect image having a PSNR of infinity. In the context of MRI super-resolution, PSNR is important because it reflects the ability of the super-resolution technique to accurately reconstruct high-frequency details in the image. This is particularly important in MRI, where small details can have significant clinical implications.

SSIM, on the other hand, is a metric that measures the structural similarity between two images. Specifically, SSIM calculates the similarity of the luminance, contrast, and structure of the images being compared. While SSIM is useful for evaluating overall image similarity, it is less sensitive to high-frequency details, which are important in the context of super-resolution for MRI. Therefore, while SSIM can provide useful information about the overall similarity of two images, it may not be as effective at evaluating the ability of a super-resolution technique to accurately reconstruct high-frequency details.

In our paper, we include only PSNR as a metric for evaluating the performance of our super-resolution technique for brain MRI. However, we primarily focused on the PSNR results as this metric provides a more accurate reflection of the ability of the technique to accurately reconstruct high-frequency details. Our results showed that the use of our super-resolution technique significantly improved the PSNR values compared to the baseline low-resolution images, indicating that our technique was effective at accurately reconstructing high-frequency details.

$$PSNR = 10\log\frac{255^2}{MSE}$$

D. Loss Function Determination

In our paper, we sought to explore the use of different loss functions to optimize the performance of our super-resolution technique for MRI. While Mean Squared Error (MSE) is a commonly used loss function for image super-resolution, we found that it underperformed during the training process for our specific application. As a result, we decided to experiment with different loss functions to identify the most effective option.

After extensive experimentation, we found that Mean Squared Logarithmic Error (MSLE) performed significantly better than MSE in terms of producing higher resolution metrics. MSLE is particularly useful for image super-resolution applications as it is less sensitive to outliers, which can be a common issue in medical imaging data. MSLE also places a higher weight on errors for lower pixel values, which is important in the context of MRI super-resolution as lower pixel values typically correspond to high-frequency details.

As a result of our experimentation, we established MSLE as the baseline loss function for our super-resolution technique and used it to test all of the networks. This allowed us to accurately compare the performance of different network architectures and identify the most effective approach for our specific application. By using MSLE as our main loss function, we were able to achieve significant improvements in image resolution and quality, ultimately leading to a more effective super-resolution technique for MRI [8].

$$MSLE = \frac{1}{n}\log(\sum Y_i - Y_j)^2$$

E. Technology and Datasets

In order to ensure that our network is well-equipped to handle a wide range of imaging scenarios, we utilized a large dataset of 50-800 3D MRI scans (depending on the architecture). All of these images came from the ABIDE dataset. These scans were carefully selected to include a variety of imaging parameters, such as field strength, contrast, and resolution, in order to ensure that our network is trained to handle the full range of imaging scenarios that it may encounter in clinical practice.

Each of the 3D scans in our dataset contained 256 slices, from which we selected the 128th slice to train our model. This approach was chosen to ensure that we have a sufficient number of training examples while also avoiding any potential bias that might arise from using only a subset of the available slices. By selecting the middle slice of each 3D scan, we can be confident that our training data is representative of the full range of imaging parameters present in each scan.

To evaluate the performance of our network, we tested it on a set of 10 randomly selected images. We used the peak signal-to-noise ratio (PSNR) as our metric for evaluation. These metrics are widely used in the image processing community and are commonly used to assess the quality of reconstructed images.

All of the experimentation for this study was completed on PyTorch 1.9.0 with NVIDIA GPU processing to speed up training time. The use of GPU processing allowed us to train our network more efficiently, enabling us to complete the necessary experimentation in a timely manner. Additionally, the use of PyTorch provided a powerful and flexible framework for implementing and testing our network, allowing us to easily modify and iterate on our approach as needed.

IV. Results

In our study, we utilized Fig. 3 to present the qualitative observations of our super-resolution technique based on the U-Net architecture. The U-Net algorithm is a sophisticated deep learning model that has gained widespread popularity for its high efficacy in image super-resolution tasks, particularly in the medical imaging domain. One of the key strengths of the U-Net algorithm is its ability to reconstruct fine details from low-resolution inputs, particularly in the deeper regions such as the hippocampal areas. These regions are particularly critical in detecting neurological conditions such as Alzheimer's and Parkinson's disease.

By effectively improving the resolution of MRI images in these regions, the U-Net algorithm can enhance the accuracy and reliability of disease detection, ultimately leading to improved patient outcomes. Furthermore, our study also compared the U-Net algorithm against four other commonly used networks in terms of average Peak Signal-to-Noise Ratio (PSNR) as denoted by Table I. The results of our study demonstrate that the U-Net algorithm outperformed the other networks in terms of average PSNR. This highlights the superiority of the U-Net algorithm in producing high-quality, super-resolved MRI images, which is of paramount importance in medical diagnosis and treatment.

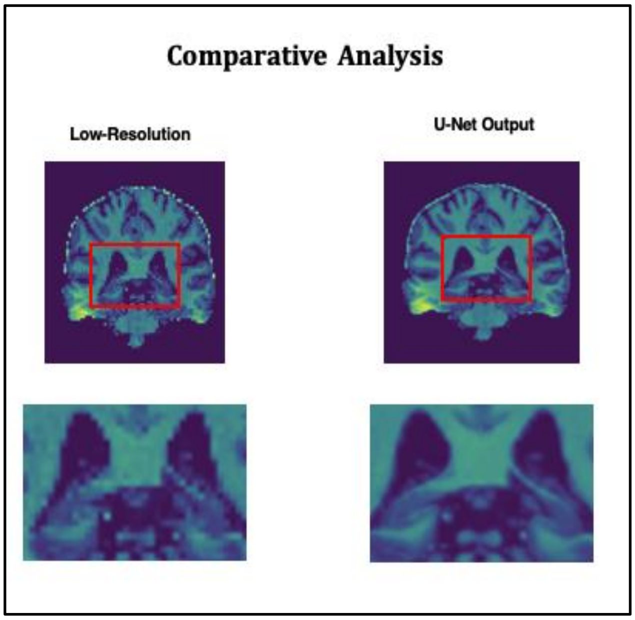

Fig. 3. Qualitative observations from low-resolution (left) to u-net output (right).

TABLE I. PSNR COMPARISON

Dataset I	PSNR
SR DenseNet	29.62458398
VDSR	29.9758636
U-Net	**30.40278329**
U-Net++	30.18895619
SR CNN	30.13047352

V. DISCUSSION AND CONCLUSION

Our study has delved into the potential of deep learning-based super-resolution techniques to improve the resolution of MRI images in the medical imaging domain. The results of our research are highly promising, demonstrating the effectiveness of the U-Net architecture in reconstructing fine details from low-resolution inputs, specifically in the hippocampal regions that play a crucial role in detecting neurological conditions such as Alzheimer's and Parkinson's disease. Our comparison with four other commonly used networks has highlighted the superiority of the U-Net algorithm in producing high-quality, super-resolved MRI images. The U-Net performed with an average PSNR of 30.40, outperforming all other algorithms in terms of PSNR.

By improving the accuracy and reliability of medical diagnosis and treatment in the field of neurological disorders, our research could have a significant impact on the lives of patients and their families. However, our study is just the beginning, and there are numerous potential avenues for future research.

For example, we could explore the integration of multiple deep learning models for improved accuracy and efficiency. This could potentially lead to even more precise and comprehensive diagnosis and treatment for neurological disorders.

Another possible direction for future research is to expand the dataset used for training to encompass a broader range of neurological conditions and patient populations. This could help to improve the generalizability of our results and make our super-resolution techniques even more widely applicable in the medical imaging field.

While the U-Net architecture has demonstrated remarkable effectiveness in MRI super-resolution tasks, it is not without limitations. One notable limitation is the potential for overfitting, especially when dealing with limited or imbalanced training datasets. Due to the large number of parameters in the U-Net model, there is a risk of the model memorizing specific features from the training data rather than learning generalizable patterns. This can result in reduced performance when faced with unseen or diverse data during the testing phase. Another limitation is that our U-Net struggled to perform well with regards to the SSIM metric, underperforming current state of the art. Modifying loss function optimized to SSIM could potentially fix this issue.

Additionally, the U-Net architecture may struggle to capture long-range dependencies and complex spatial relationships within the MRI images, which could impact the accurate reconstruction of fine details. Moreover, the U-Net's performance may vary depending on the specific MRI imaging modality or imaging protocols, making it less universally applicable across different types of MRI scans. Addressing these limitations through appropriate regularization techniques, larger and more diverse training datasets, and exploration of alternative architectures could further enhance the performance and generalizability of the U-Net in MRI super-resolution tasks.

Moreover, we could investigate the potential of combining super-resolution techniques with other image processing techniques such as image segmentation and registration. By integrating these techniques, we could potentially achieve even more precise and comprehensive diagnosis and treatment for neurological disorders.

In summary, our study highlights the significant potential of deep learning-based super-resolution techniques for medical imaging, particularly in the detection and treatment of neurological disorders. By enhancing the resolution of MRI images, our research can contribute towards improving patient outcomes and ultimately lead to a better quality of life for individuals suffering from these conditions. The outcomes of our research could pave the way for further advancements in the field, leading to even more accurate and efficient diagnosis and treatment in the future.

VI. Acknowledgment

I am grateful for the tremendous support and guidance provided by Dr. Matthew Rosen, an associate professor and low-field MRI media director at Harvard Medical School, who introduced me to this fascinating topic and generously shared his computational resources with me. His mentorship and expertise were invaluable to the success of my research. Additionally, I would like to extend my appreciation to Dr. Sean Young, whose support and encouragement were instrumental in my research journey. Their contributions have been instrumental in shaping my research and have inspired me to pursue my academic and professional goals with passion and dedication.

References

[1] Li BM, Castorina LV, Valdés Hernández MdC, Clancy U, Wiseman SJ, Sakka E, Storkey AJ, Jaime Garcia D, Cheng Y, Doubal F, Thrippleton MT, Stringer M and Wardlaw JM (2022) Deep attention super-resolution of brain magnetic resonance images acquired under clinical protocols. Front. Comput. Neurosci. 16:887633. doi: 10.3389/fncom.2022.887633.

[2] Ottesen, J.A., Caan, M.W.A., Groote, I.R. et al. A densely interconnected network for deep learning accelerated MRI. Magn Reson Mater Phy (2022). https://doi.org/10.1007/s10334-022-01041-3Zhou Z, Ma A, Feng Q, Wang R, Cheng L, Chen X, Yang X, Liao K, Miao Y, Qiu Y. Super-resolution of brain tumor MRI images based on deep learning. J Appl Clin Med Phys. 2022 Nov;23(11):e13758. doi: 10.1002/acm2.13758. Epub 2022 Sep 15. PMID: 36107021; PMCID: PMC9680577.

[3] Zhou Z, Ma A, Feng Q, Wang R, Cheng L, Chen X, Yang X, Liao K, Miao Y, Qiu Y. Super-resolution of brain tumor MRI images based on deep learning. J Appl Clin Med Phys. 2022 Nov;23(11):e13758. doi: 10.1002/acm2.13758. Epub 2022 Sep 15. PMID: 36107021; PMCID: PMC9680577. R. Nicole, "Title of paper with only first word capitalized," J. Name Stand. Abbrev., in press.

[4] de Leeuw den Bouter, M.L., Ippolito, G., O'Reilly, T.P.A. et al. Deep learning-based single image super-resolution for low-field MR brain images. Sci Rep 12, 6362 (2022). https://doi.org/10.1038/s41598-022-10298-6 M. Young, The Technical Writer's Handbook. Mill Valley, CA: University Science, 1989.

[5] Wang, Li, et al. "Multi-Scale Attention Network for Image Super-Resolution." Journal of Visual Communication and Image Representation, vol. 80, 2021, p. 103300., doi:10.1016/j.jvcir.2021.103300.

[6] Bahrami K, Shi F, Rekik I, Gao Y, Shen D. 7T-guided super-resolution of 3T MRI. Med Phys. 2017 May;44(5):1661-1677. doi: 10.1002/mp.12132. Epub 2017 Apr 22. PMID: 28177548; PMCID: PMC5686784.

[7] "Image Segmentation with Monte Carlo Dropout UNET and Keras¶." 42: A Blog on A.I., nchlis.github.io/2019_10_30/page.html.

[8] "Mean Squared Logarithmic Error Loss." *InsideAIML*, insideaiml.com/blog/MeanSquared-Logarithmic-Error-Loss-1035.

[9] M. Sharma, S. Chaudhury and B. Lall, "Deep learning based frameworks for image super-resolution and noise-resilient super-resolution," 2017 International Joint Conference on Neural Networks (IJCNN), 2017, pp. 744-751, doi: 10.1109/IJCNN.2017.7965926.

[10] Koonjoo, N., Zhu, B., Bagnall, G.C. et al. Boosting the signal-to-noise of low-field MRI with deep learning image reconstruction. Sci Rep 11, 8248 (2021).