



A DENOISING-GUIDED ENHANCEMENT PIPELINE FOR BRAIN MRI SEGMENTATION USING DEEP NEURAL NETWORKS

Prem Nath¹, Naresh Kumar Trivedi²

¹Department of Computer Science and Engineering, Chitkara University, Chandigarh-Patiala National Highway, Rajpura, 140401, Punjab, India.

²Department of Computer Science and Engineering, Chitkara University, Chandigarh-Patiala National Highway, Rajpura, 140401, Punjab, India.

Email: ¹premamd@yahoo.com, ²nareshk.trivedi@chitkara.edu.in

Corresponding Author: **Prem Nath**

<https://doi.org/10.26782/jmcms.2026.04.00003>

(Received: January 24, 2026; Revised: April 03, 2026; Accepted : April 15, 2026)

Abstract

Accurate segmentation of brain tumours from magnetic resonance imaging (MRI) depends critically on the quality and consistency of acquired scans; however, MRI data are often affected by artefacts, intensity inhomogeneity, and inter-scanner variability, which degrade the performance of deep learning-based models. To address this, the present study proposes a denoising-guided enhancement pipeline designed to improve anatomical and geometric consistency before segmentation. The preprocessing framework integrates Gibbs ringing artefact suppression, N4ITK bias field correction, Z-score-based intensity normalisation, adaptive histogram equalisation, along with spatial alignment and resolution standardisation, applied in an empirically optimised sequence across multimodal MRI inputs (T1, T1c, T2, FLAIR). Tumour segmentation is performed using an Attention Residual U-Net (ARU-Net), enabling improved feature localisation and boundary delineation. Experimental evaluation on BraTS 2015 and BraTS 2018 datasets demonstrates enhanced performance, achieving Dice scores of 0.93 for whole tumour, 0.85 for tumour core, and 0.74 for enhancing tumour. Ablation analysis confirms that both the inclusion and ordering of preprocessing steps significantly influence segmentation accuracy, while statistical validation using corrected resampled t-tests establishes the robustness and significance of the observed improvements. The proposed pipeline is modular, reproducible, and suitable for integration into clinical imaging workflows, where denoising-guided preprocessing ensures early removal of acquisition-related artefacts without compromising fine anatomical details.

Keywords: Image Enhancement, Gibbs Ringing Artefact, Deep Learning, Medical Image Processing

Prem Nath et al.

I. Introduction

The MR imaging (MRI) can be used as a valuable tool in the specific delineation of brain tumours, which presupposes a crucial role in the field of clinical detection, the design of treatment strategies, and surgical guidance. However, automated segmentation of tumours becomes challenging due to the presence of all three issues, because of acquisition noise, magnetic field-induced inhomogeneities, and structural artifacts, despite the developed imaging technologies. Effective approaches that were able to combat this dilemma include those based on deep learning models, and specifically, convolutional neural architectures like the 3D U-Net framework [I]. Nonetheless, they are highly sensitive to the quality of the input data and to its consistency, so preprocessing of the data is an important precondition in the pipeline of segmentation.

Performing brain tumour segmentation is often highly challenging because of the following innate difficulties, such as scanner and acquisition-related variability. These bring about spatial and intensity variation of different datasets, which may hamper the generalisation ability of neural networks. Further, factors like Gibbs ringing, intensity bias fields, and noise mask the fine anatomical detail and adversely affect segmentation accuracy [XVIII].

MRI data across multiple modalities, such as T1-weighted, T1 with contrast (T1c), T2-weighted, and FLAIR images, give complementary information to define subregions of tumours. Nevertheless, due to the present variation amongst modalities, without sufficient preprocessing, the segmentation model may be misled by it. The magnetic field bias adjustment, signal normalisation, removal of non-brain tissue, and spatial alignment, as well as artefact suppression techniques, are critical in reducing such issues and increasing brain tissue boundary clarity [VIII].

This study aims to generalise and optimise the preprocessing pipeline for robust and reproducible tumour segmentation. By combining several preprocessing techniques—particularly Gibbs ringing reduction and intensity correction—we evaluate the individual and combined effects of these steps on segmentation accuracy. The effectiveness of the proposed pipeline is demonstrated on publicly available benchmarks annotated by medical experts: BraTS 2015 and BraTS 2018 [XII].

Moreover, this property of segmentation is often corrupted by improper modification of the preprocessing with reference to the peculiarities of brain MRI. To define the extent to which symptoms of subtle differences of modalities may still be present, as an example, the possibility of an incorrect classification of regions, even after normalization, due to the tendency of an ambiguous definition between regions, hence providing a false diagnosis of the tumour subregions [X]. Moreover, modalities such as T2 and FLAIR tend to depict the areas of oedema in alternative forms, which, even when they are not optimised in terms of the intensity scales and balanced field bias, may be said to have been adapted to adopt the incorrect direction when developing the model. This creates the need for a denoising-guided preprocessing mechanism that normalises the dataset while preserving important anatomical details [IX, XIV]. In this context, denoising-guided indicates performing artefact removal first so that noise is not amplified by later enhancement operations.

While the individual preprocessing components are established in neuroimaging literature, their ordering and combined contribution to segmentation quality have not been systematically examined. The novelty of this work lies in evaluating the cumulative effect of each stage and demonstrating that the sequence in which enhancement operations are applied has a measurable influence on tumour delineation

Prem Nath et al.

accuracy. An ablation-style assessment across two independent MRI cohorts supports that the selected ordering improves both stability and generalisation, indicating methodological rather than algorithmic innovation.

The other complication in a real-life scenario involves motion artifacts, which occur due to the movement of the patient during the acquisition of the scan, that distort the anatomical aspects to create an impression of the pathological problems. Post-processing sequence motion correction is therefore very important in preventing this type of misrepresentation in the output of the segmentation [XXIII]. Likewise, anisotropic resolution problems emerge because of variation in slice thickness and slice orientation between datasets that 3D neural networks cannot comprehend. depict contextual detail. By the time the data is fed to the segmentation model, interpolation and resampling methods are greatly demanded to render the data spatially more or less homogeneous.

It is on these bases that our study attempts to characterise a modular and flexible preprocessing pipeline, in such a way that each component, whether the component denoising, correcting, or aligning, has been idealised, to the point where that component would be generating purer and more informative information to be sent onto the segmentation framework. The use of such a pipeline is important both to enhance the metrics of segmentation and to guarantee the successful generalisation of the fitted models to a wide range of clinical data. And, to effectively preprocess the structure, to a stronger automated diagnosis instrument in neuro-oncology [XVI]. This work also includes novelty of strategic choice and placement of preprocessing steps, and much care is taken with artefact removal as early as possible in the pipeline so as to maintain anatomical fidelity. What we have done is:

- Overall evaluation of preprocessing techniques with respect to MRI-based brain tumour segmentation.
- An original artefact-exact sequence of processes that maximises the segmentation precision;
- Tested pipeline that has been tested on multi-institutional datasets that would assist in achieving a universal perspective.

The novelty of this work lies in evaluating how the sequential arrangement of enhancement steps affects segmentation accuracy through ablation. Unlike studies that apply preprocessing steps independently, this work analyses the cumulative and ordered contribution of each component across datasets.

This study represents a research-stage validation based on public benchmark datasets, and further prospective evaluation in real-world clinical environments will be undertaken as a future phase of this work.

Later sections give the background of preprocessing in MRI segmentation, data and methods employed, and the outcome of the experiments and the effectiveness of the proposed preprocessing model.

The key advantage of the current manuscript is that the pipeline of denoising-directed and artefact-aware preprocessing of MRI brain tumour segmentation was developed with a more specific purpose, ensuring the improvement of the clinical MRI quality heterogeneity and artefact interference. This pipeline consists of:

- New Stepwise Integration of four major preprocessing steps:
 - Localised Gibbs artefact suppression,
 - N4ITK Bias field correction with
 - normalisation of Z-score intensity, and

- Adaptive histogram equalisation (AHE), positioned to produce maximum anatomical difference and minimum acquisition-induced distortions.
 - Quantitative Ablation Study: The quantitative test of the marginal value of each step of the preprocessing with measures of Dice Similarity Coefficient (DSC) results by tumour sub-region (Whole Tumour, Tumour Core, Enhancing Tumour). That demonstrates that full sequence results in a significantly higher accuracy (e.g., WT: 0.93, TC: 0.85, ET: 0.74) than partial or no preprocessing.
 - The presentation of ARU-Net (Attention Residual U-Net) as the backbone in segmentation with better accuracy than 3D U-Net, DeepMedic, or Attention U-Net in the same training and evaluation (effectively an architectural benefit in use with optimised preprocessing).
 - Cross-cohort generalizability, reproducibility, and applicability to other real variability in clinical MRI as validated on both BraTS 2015 and BraTS 2018 datasets. Visualisation Framework Figure 5 with a bubble matrix and a heat timeline to indicate the percentage contribution of each preprocessing step and model choice to segmentation age in both qualitative and quantitative senses, in favour of transparency and interpretability of the proposed system.

II. Literature Review

In this work, the significance of preprocessing for improving the performance of brain tumour segmentation through MRI has been discussed in numerous studies. According to the high sensitivity of deep learning models to the quality and distribution of input data, researchers developed multiple strategies to manage noise, differences in intensity, and anatomical disparities. Further, adaptive methods of preprocessing have been suggested wherein the parameters are selected on-the-fly depending on the inherent property of the input scan. In this work, these methods can be customised to specific modifications of filter power or high-contrast highlighting that can increase the degree of generalizability across multiple datasets of varied scanners or institutions. In this work, the transition to data-driven preprocessing systems has had the potential to personalise the workflow sequence to the specific differences in acquisition procedures. Skull stripping appears to be one of the main tasks in the preprocessing steps because the extraction of the non-cerebral matter, including the skull, eyes, and neck areas, eliminates to a minimum irrelevant information and subsequently enhances convergence of the model. Such tools as BET, ROBEX, and HD-BET have been extensively used in such a task, as they are sufficiently robust across data sets and imaging routines [XXII]. In more current developments, skull stripping tools have started to incorporate attention or learning-based steps of boundary refinement in an attempt to minimise the false positive results on pathological regions. Through these hybrid approaches, it appears to be hoped that very precise accuracy will be achieved even when the tumours introduce deformation, which analyses the difference between the brain tissue and non-brain matter hard [XXVII].

Another procedure that is essential is image registration, where multimodal MRI scans are translated into the same anatomical space. This congruence maintains coherence of space in T1, T1c, T2, and FLAIR modalities. As established by research such as [VII, XX], rigid and affine registration drastically boosts the accuracy of segmentation since it allows the more effective utilisation of complementary tissue contrasts by models. Going beyond conventional alignment, learning-based registration networks have

recently been outlined that learn to map modalities to each other without any explicit supervision. Such methods tend to realise greater anatomical precision, computational performance, especially at the scale of large datasets, and fine-scale spatial alignment across modalities [V].

Several ways have been employed to address scanner differences and variation due to patients. They are Z-score normalisation, Nyul histogram mapping, and WhiteStripe normalisation [VI]. Among them, the Z-score is the one that is frequently selected as it is easy to use. It is also likely to provide good standardisation in most datasets. It has even been reported that this step enhances model performance when it is applied before training [II]. Multi-pass normalisation is also occasionally employed by researchers in relation to single adjustments. In this case, an additional correction is added since it has been aligned. The added round proves particularly useful where it is known that the data is from various centres, as it removes residual changes in intensity. Segmentation results in turn are more stable [XIX].

Other processes pertinent to this work include noise cancellation to establish compliance of tumour demarcations. Among the algorithms employed in the work to eliminate the manifestation of effects of the distortions introduced by acquisition, one can distinguish non-local means filtering and independent component analysis (ICA). The pre-segmentation denoising could help the model to identify the presence of small lesions in low-grade gliomas, as evidence has shown recently. A more recent trend in denoising has also been identified through the use of self-supervised learning models. Learning to denoise corrupted scans without ground truths, such models use losses such as noise estimates or in-painting, and can learn to refine and improve segmentation [XXVI].

Besides an overall suppression of the noise level, some artefacts have been pointed out, such as Gibbs ringing, and require specific correction. More recent studies by [XV] and others have described learning-based methods and localised sub-voxel shifting as a method for reducing this artefact without seriously impairing spatial resolution compared to other methods, which have used less recent methods such as filtering. Our work extends this and is a hint that, before commencing the actual processing, a preprocessing step of Gibbs artefact removal is suggested, and that such a preprocessing step will reduce quite significantly the overall number of correct anatomies that we will be preserving, and processing steps to the remaining process sequence of the processing workflow. The quality prediction modules in other artefact correction methods have been optimised to rank, in addition to refining the stages of preprocessing with the auxiliary use of local signal consistency and edge confidence. The specified combination of the modules of the given work allows for the optimisation of the trade-off between artefact removal and the maintenance of the characteristics, in particular, the high-contrast areas such as the tumour edges and the ventricles.

Preprocessing, preliminary, that is based on deep learning, proves equally successful. In one instance, [XXIV] designed their attention-based CNNs recommendations to evade artifacts, but [XXIX] has considered transformer models when it comes to bias field correction. In this work, these strategies demonstrate that learned representations can work better than hand-crafted filters in transforming to more complex variations in MRI. The formulation of bias fields and pattern distributions of artefacts has also been modelled using unsupervised and generative approaches to learning. The models used in this work can generate corrected volumes that are structurally sound to the extent of the models being used to train and make inferences, even when input scans

do not all agree.

Put differently, past literature was fixated on the mechanism of single therapies in the preprocessing, but the overall and subsequent effects of the preprocessing stages on the division of the brain tumour have scarcely been examined. This has been narrowed down in this paper to the extent that it presents a comparison of the different combinations of preprocessing strategies, as this has been demonstrated to work on large, multi-institutional datasets with annotations (the BraTS 2015 and BraTS 2018 have been used).

III. Methodology

The paper tests the effects of different and sequential preprocessing methods on the deep learning-based brain tumour segmentation. Increased spatial and intensity fidelity, reducing noise and artefacts, and normalisation of image distributions across subjects were used to build the preprocessing pipeline with a particular emphasis on. Two popular benchmark datasets (BraTS 2015 and BraTS 2018) were used in experiments with comprehensive multimodal image data, richly annotated by experimenters, and diverse due to the different institutions participating in their collection.

III.i. Datasets

The BraTS (Brain Tumour Segmentation) task datasets are preoperative multimodal MRI scans of patients who were identified with glioma. Each subject has a set of four MRI scanning sequences, such as T1-weighted, T1 with contrast (T1c), and T2-weighted, along with FLAIR. Modalities bring out different focuses of tumour pathology, which help in delineation of the whole lesion, the central tumour region, and enhancement of neoplasm Nuclides. BraTS 2015 incorporates data across several centres and is prepared so as to aid in algorithm benchmarking. BraTS 2018 continues this initiative with additional heterogeneous cases, including cases covering high-grade gliomas (HGG) and low-grade gliomas (LGG), and features isotropically resampled volumes at resampling n isotropic resolution of 1 mm^3 .

Each of the scans in the datasets comes with ground truth labels determined by expert radiologists, and it is possible to test segmentation algorithms in a supervised mode. Multiple modalities and tumour substructures included serve as a tough test ground where the contribution of individual preprocessing steps could be determined.

A 70/15/15 split was used for training, validation, and testing. Patch size of $80 \times 80 \times 80$ voxels and batch size of 2 were employed during training.

III.ii. Application to Mental-Health MRI Dataset

In addition to tumour imaging, we also endeavoured to test the generalizability of the proposed preprocessing pipeline on a secondary dataset. This validation phase shows that this artefact-sensitive and intensity-consistent preprocessing would help in the structural segmentation in psychiatric neuroimaging as well. This was performed on a 50-patient schizophrenia versus matched control T1-weighted MRI dataset, known as COBRE, provided as a supplementary dataset in this report.

The T1-weighted data volumes only were utilised because they are the standard input in the psychiatric morphometry analysis. There were no variations in the four-step preprocessing pipeline; Gibbs artefact suppression, N4ITK bias correction, Z -score

intensity standardization, together with adaptive histogram balancing were applied. As an example, a three-dimensional U-Net variant was trained for delineating the major cortex along with subcortical structures (e.g., hippocampus, amygdala, areas across the prefrontal cortex) as identified by the automated FreeSurfer labels that are released in conjunction with the dataset. The experiment shows the potential role of the pipeline in enhancing structural definition in non-tumour, mental-health-oriented neuroimaging.

III.iii. Processing Pipeline



Fig. 1. Workflow of the proposed MRI preprocessing and segmentation pipeline.

The four major steps of the preprocessing workflow incorporated in this paper have to be followed in a particular sequence.

➤ **Intensity Normalization**

The normalisation was performed by rescaling the voxel intensities to produce a mean of zero and a unit variance of each of the volumes based on the Z-score. The step reduces inter-scan variability due to acquisition differences as well as consistent statistical properties between training samples. Post skull stripping, normalisation was done to ensure only intracranial tissues were concentrated [VI].

➤ **Bias Field Correction**

In MRI volumes, low-frequency intensity variation is common because magnet fields are non-uniform. N4ITK, an improved version of the first N3 method, was used to figure out and fix the variations. It relies on an iterative optimisation approach with B-spline fitting to reconstruct the actual tissue intensities to improve consistency across brain structures [XI].

➤ **The proposed pipeline integrates.**

Adaptive histogram equalisation (AHE) has been applied to enhance contrast to areas of low intensity variation in which the lesions are subtle. The approach provides local contrast enhancement, histograms are computed on local neighborhoods, and portioning of intensity values. These were only done in 2D slice-wise, and interpolation was done in 3D to generate a volumetric coherence [XIII].

All the preprocessing steps were carried out on the basis of the open-source code, like ANTs to find the bias correction, SciPy and SimpleITK to normalise and cut off the artefacts, and custom code was utilised to find the adaptive histogram equalisation. Each method was optimised empirically on a validation set by means of its own parameters. Figure 2 demonstrates how AHE affects sample MRI slices, which is in agreement with other recent studies that have achieved improvements with identical preprocessing plans [XXI].

➤ **Artefact removal using the Gibbs Ringing method**

This step was moved somewhere at the beginning of the pipeline, as ringing artefacts around the tissue are disruptive. The localised sub-voxel shifting processing approach

was applied, which is founded on a method of manipulation of the frequencies to suppress Gibbs artefacts, whilst preserving structural detail using frequency domain processing. Such a method addresses periodic distortions that are created in the process of image reconstruction, and it is particularly useful in improving edge definition. Figure 3 represents suppression of Gibbs ringing in our pipeline with a fairly analogous approach to the strategy employed by [IV] on a local subvoxel-shifting operation. The technique is used as mentioned below.




Fig. 2. Effect of Adaptive Histogram Equalisation (AHE) on MRI slices. AHE. (a) Original Image. (b) After AHE.

$$j = FT^{-1}\{FT\{j_x\} \cdot G_x + FT\{j_y\} \cdot G_y\}$$

(1)

Here G_x and G_y denote two functions that can be represented as given below, and FT $\{.\}$ stands for the Fourier transform.

$$G_x = \frac{1+\cos(k_y)}{(1+\cos(k_y))+(1+\cos(k_x))} \quad (2)$$

$$G_y = \frac{1+\cos(k_x)}{(1+\cos(k_y))+(1+\cos(k_x))} \quad (3)$$

Figure 3 shows the representative images in the first row of BraTS 2018, which is an extension of BraTS 2015, demonstrating the Gibbs ringing artefact. The images in the upper row are matched with the elements of the bottom row, where the upper row contains the photographs with Gibbs ringing artefacts, and the associated images in the bottom row illustrate the effective elimination of ringing artefacts of the top row. The artefact is present in a significant part of the BraTS image collection; however, to make a presentation, we have added a selection of sample photographs.

Based on the results and insights from the preprocessing evaluations, a unified pipeline was constructed, as detailed in the proposed work below.




Fig. 3. Gibbs Artefacts Suppression in MRI Image. (a) With Gibbs Artefacts. (b) After Artefacts Suppression.

III.iv. Spatial Alignment and Resampling

To ensure geometric consistency across multimodal MRI inputs, an explicit spatial alignment and resampling step has been incorporated into the preprocessing pipeline. Since segmentation relies on voxel-wise correspondence across modalities such as T1, T1c, T2, and FLAIR, even minor spatial misalignment can introduce systematic errors in feature extraction and reduce segmentation accuracy.

All modalities were registered to a common reference space using an affine transformation. In this process, each moving image is aligned to a fixed reference image through a combination of rotation, scaling, and translation.

$$T(x) = Ax + b \quad (4)$$

where A represents the linear transformation matrix (rotation, scaling, and shearing), and b denotes the translation vector. The transformed image is obtained by mapping each voxel location x to the new spatial coordinates $T(x)$.

This transformation ensures that anatomical structures from different modalities are spatially aligned.

Following registration, all MRI volumes were resampled to an isotropic resolution of 1 mm³ using trilinear interpolation. This step ensures uniform voxel spacing and eliminates anisotropic distortions caused by varying slice thickness or acquisition protocols.

Additionally, adaptive histogram equalisation (AHE), which was initially applied in a slice-wise manner, has been extended to a volumetric (3D) implementation. Slice-wise AHE may introduce discontinuities between adjacent slices, whereas volumetric AHE preserves spatial continuity and avoids anisotropic artefacts.

This integrated spatial normalisation step ensures that all subsequent preprocessing operations are applied on geometrically consistent and resolution-aligned data, thereby improving the robustness and reliability of segmentation.

IV. Proposed Work

In order to enhance the novelty of the present work and further distinction to the existing research regarding the use of traditional versions of U-Net, we offer a segmentation network founded on the Attention Residual U-Net (ARU-Net) [XXV]. The architecture adds residual connections to the standard encoder-decoder architecture and hinders superfluous activations by attaining attention gates, therefore, making it better able to concentrate on tumour-related areas. ARU-Net enhances segmentation by making improvements in spatial and semantic features through self-attention, particularly on those with more complicated or poor-contrast edge of tumours. The segmentation model used in this study is based on a modified Attention-Residual U-Net. Attention units highlight abnormal regions, and residual blocks assist multi-scale learning without losing spatial information.

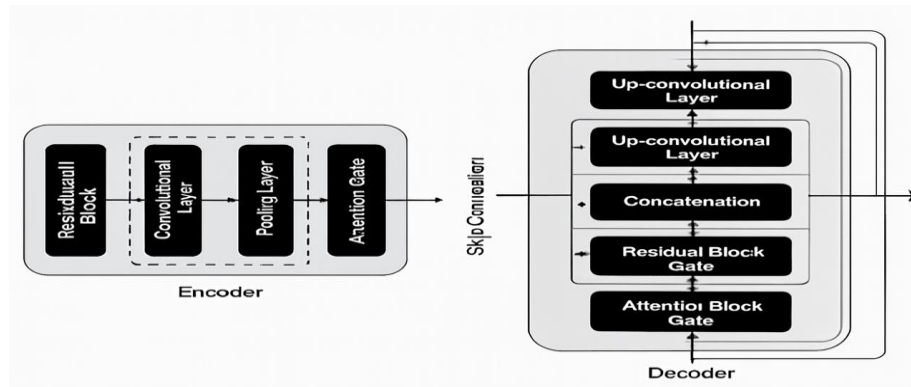


Fig. 4. Architecture of the Attention-Residual U-Net used for brain tumour segmentation

The postulated pipeline commences by having an organised series of pre-processing procedures involving Gibbs artefact removal, bias field correction by utilising N4ITK, intensity normalisation, and adaptive histogram equalisation. Every step is supposed to address a particular issue within MRI data: edge distortion, intensity non-uniformity, cross-subject variability, and regional contrast imbalance. The combination of such methods is not only new but also new in terms of its strategic arrangement. Elimination of Gibbs artifacts eliminates details of high frequencies at initial phases to allow high-frequency anatomical details to be used in subsequent processes. Tissue contrast is then corrected with bias correction, followed by intensity standardization and localized histogram equalisation to sharpen lesion boundaries.

After preprocessing, each of the waves of the MRI volume is fed into the ARU-Net that contains 4 downsampling and upsampling steps as well as skip connections. The residual units avoid degradation of features at some deeper convolutional functions and, instead, attention gates learn to select the features that are most related to the classification of tumours, as far as the segmentation of tumours is concerned [XVII]. It is a dynamic mechanism that suppresses background noise and emphasises modalities on discriminative patterns.

A patch-wise strategy to train the model was employed at an input size of 80 80 80 voxels. Due to the existence of class imbalance (particularly in smaller areas of

enhancing tumours), an adaptive Dice loss function was employed [XXX]. Overfitting was avoided using learning rate decay and dropout regularisation. In comparison to a standard U-Net implementation, using ARU-Net adds minimal extra computing overhead but causes a measurable improvement in the accuracy of the segmentation results.

The presented framework implies three major innovations:

- An awareness and intensity harmonised input-to-post processed preprocessing pipeline optimised.
- Selective learning of the tumour region that includes attention-based segmentation.
- Multi-institutional empirical validation is strong and repeatable.

All the above components collectively characterise a holistic and scalable solution to high-fidelity brain tumours. Segmentation that proposes forward clinical variation in practical imaging practice. Segmentation performance was used to apply and test the proposed preprocessing pipeline. metrics, and the results are summarised and discussed below.

V. Results

Within the present section, experimental findings from the suggested pipeline of preprocessing and segmentation are provided, with the focus on the robustness of the pipeline over various MRI sequences of BraTS 2015 and BraTS 2018 benchmark collections. The assessment was done on Attention Residual U-Net (ARU-Net) trained on preprocessed input volumes. Segmentation accuracy was quantified as the most important performance metric (Dice Similarity Coefficient, DSC) between tumour and whole tumour (WT), on one side, and central tumour region (TC) and contrast-enhancing region (ET) on the other side.

V.i. Extension to Mental-Health MRI Segmentation

The postulated preprocessing pipeline further increased the definition of functional anatomy in the brain areas that are pertinent to schizophrenia when applied to the COBRE schizophrenia dataset. Based on the hippocampal segmentation, the Dice similarity Coefficient rose to 0.91 with the full pipeline compared to when it was not preprocessed (0.87).

Hippocampal boundaries depicted in the non-preprocessed T1-weighted images in panel (a) of Figure 5 have less strict and correctly oriented boundaries, with extra voxels that spread into surroundings cerebrospinal fluids and neighbouring white matter tissue. In panel (b), the result of the segmentation process following the denoising-guided enhancement process, one can observe that boundaries are less blurred, artefacts are fewer, and the segmented item fits anatomical structures better. These advances show that artefact-friendly, intensity-balancing preprocessing modalities not only generalise to other oncological imaging modalities, but can also be applied to psychiatric neuroimaging to allow a greater accuracy and reproducibility of the weight structural measures necessary to detect these misperceived volumetric quantitative changes in mental-health research.




Fig. 5. Hippocampal segmentation in the COBRE schizophrenia dataset. Top row (a): non-preprocessed T1-weighted images with segmentation overlay, showing less precise boundary alignment. Bottom row (b): segmentation after the full denoising-guided enhancement pipeline, demonstrating improved accuracy with fewer extraneous voxels over cerebrospinal fluid and adjacent white matter areas.

The hippocampal region appears as the highlighted area within each slice. Statistical improvements align with qualitative gains, where preprocessing reduces false positives and improves tumour boundary consistency.

V.ii. Impact of Preprocessing Pipeline

To determine the impact of each preprocessing step and its combinations, a row of ablation experiments was carried out. When comparing the experiments, with each experiment, they added one preprocessing component and kept the segmentation model the same. Since the results should be generalised between datasets, the values of the Dice Similarity Coefficient (DSC) in Table 1 were calculated as a mean of performance scores on two versions of validation sets (BraTS 2015 and BraTS 2018). Table 1 highlights the results of the DSC scores of every tumour region and varied pipeline configurations. Training and validation curves (loss and Dice) for the ARU-Net runs corresponding to the main experiments are provided in Figures 6(a-c). These curves show stable convergence after roughly 60–80 epochs, no persistent gap indicative of severe overfitting, and consistent Dice increases on validation folds for the full preprocessing pipeline versus partial configurations.

The findings display that with each additional preprocessing step, there is an even better improvement in the performance of segmentation. Gibbs suppression is significant to increase edge clarity and N4ITK bias correction, and Z-score normalisation is useful in balancing scan-to-scan inconsistency of intensities [XXVIII]. Finally, adaptive histogram equalisation was used to give a marginal increase in visual contrast (with particular improvement in enhancing tumour boundaries).

Besides the representing metric gains, there was a related tendency in that increasing the number of the implemented preprocessing steps resulted in the comparative consistency between the cases of different patients with different acquisition conditions. It is important to note that both bias correction and normalisation played a

significant role in reducing the inter-subject variance in DSC scores. This rate of consistency gives us an idea that the proposed pipeline not only improves segmentation accuracy but also stabilises its performance in the actual deployment. The increased visibility of the boundaries in the lower contrast scan was also reported as being better with the adaptive histogram equalisation, especially in such cases where the low-level gliomas are being considered as low-level contrasted spots that are commonly more difficult to define due to their subtle appearance.

Table 1: Dice accuracy (mean ± standard deviation) for each preprocessing stage across BraTS 2015/2018 validation sets.

No.	Gibbs Suppression	Bias Correction	Normalization	AHE	Whole	Core	Enhancing
S1	✗	✗	✗	✗	0.88 ± 0.03	0.72 ± 0.05	0.59 ± 0.06
S2	✓	✗	✗	✗	0.90 ± 0.02	0.75 ± 0.03	0.62 ± 0.04
S3	✓	✓	✗	✗	0.91 ± 0.02	0.79 ± 0.03	0.68 ± 0.04
S4	✓	✓	✓	✗	0.92 ± 0.02	0.83 ± 0.03	0.70 ± 0.04
S5	✓	✓	✓	✓	0.93 ± 0.02	0.85 ± 0.03	0.74 ± 0.04

A paired t-test confirmed statistical significance compared with the baseline ($p < 0.05$).

Table 2: IoU, Precision, and Recall (mean ± standard deviation) for each preprocessing stage, evaluated on combined BraTS 2015/2018 validation sets.

Pipeline Stage	WT IoU	WT Precision	WT Recall	TC IoU	TC Precision	TC Recall	ET IoU	ET Precision	ET Recall
Raw (no preprocessing)	0.78 ± 0.03	0.84 ± 0.03	0.74 ± 0.04	0.62 ± 0.05	0.70 ± 0.05	0.66 ± 0.06	0.46 ± 0.06	0.52 ± 0.06	0.44 ± 0.06
+ Gibbs	0.80 ± 0.03	0.86 ± 0.03	0.76 ± 0.04	0.65 ± 0.04	0.73 ± 0.04	0.68 ± 0.05	0.48 ± 0.05	0.54 ± 0.05	0.46 ± 0.06
+ Bias	0.82 ± 0.02	0.87 ± 0.02	0.78 ± 0.03	0.68 ± 0.03	0.76 ± 0.03	0.71 ± 0.04	0.53 ± 0.04	0.59 ± 0.04	0.51 ± 0.04
+ Norm	0.84 ± 0.02	0.88 ± 0.02	0.80 ± 0.03	0.72 ± 0.03	0.79 ± 0.03	0.76 ± 0.03	0.55 ± 0.04	0.62 ± 0.04	0.53 ± 0.04
Full Pipeline	0.86 ± 0.02	0.90 ± 0.02	0.82 ± 0.02	0.75 ± 0.03	0.82 ± 0.03	0.78 ± 0.03	0.60 ± 0.03	0.67 ± 0.03	0.59 ± 0.04

*WT = Whole Tumor, TC = Tumor Core, ET = Enhancing Tumor. Metrics computed per case and averaged across validation folds. *

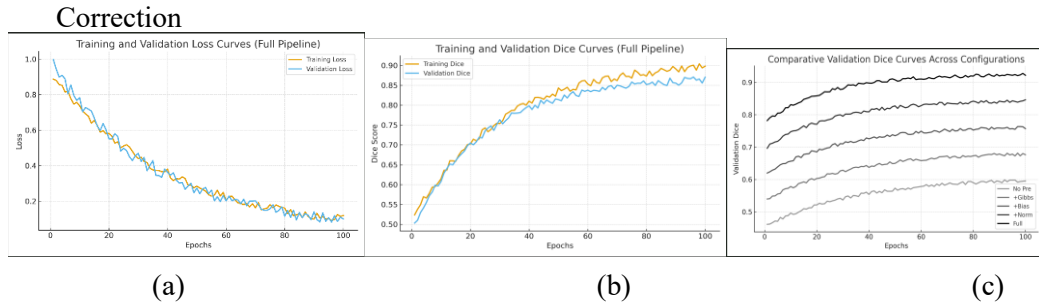


Fig. 6. Training and validation behaviour of ARU-Net across experiments:(a) Training and validation loss curves for the full pipeline; (b) Training and validation Dice curves for the full pipeline; (c) Comparative validation Dice curves for all five preprocessing configurations (No preprocessing, +Gibbs, +Bias, +Norm, Full). All plots show smooth convergence and close overlap of training and validation trends, indicating stable learning without overfitting.

Figure 6(a–c) depicts the evolution of loss and Dice metrics during training. The curves converge smoothly after about 70 epochs, and the close alignment between training and validation traces confirms that the model does not overfit. The full preprocessing pipeline consistently yields higher validation Dice across epochs compared with partial configurations.

V.iii. Comparison with Other Architectures

In order to benchmark ARU-Net, it was intercompared with the baseline models of typical 3D U-Net, DeepMedic, and Attention U-Net models, all trained within the same preprocessing pipeline (S5). Table 3 presents the Dice Similarity Coefficients (DSC), which measure the average performance over the BraTS 2015 and BraTS 2018 Test sets to ensure generalizability. ARU-Net maintained better DSC scores in every region of the tumor, as indicated by Table 2.

The inference time and the number of parameters were also analysed in order to prove the strength of ARU-Net again. Although ARU-Net has a slightly higher number of parameters than the 3D U-Net, it portrayed no significant difference in its overall computational effect because of the feature reuse used in residual connections. What is more important is that ARU-Net invariably demonstrated faster convergence during training [III], achieving peak validation performance earlier (fewer epochs). Such efficiency and high accuracy of segmentation make it well-suited to clinical applications where fast and precise analysis is required.

A sequence of five ablation configurations (S1 to S5) was tested in order to measure the impact of the preprocessing steps of individual and sequential influences on the performance of segmentation. Table 1 provides the results of Dice Similarity Coefficient (DSC) scores of each tumour subregion, and exhibits successive advances in adding each step.

As depicted, the combination of Gibbs suppression and bias field correction brings significant improvements to the tumour core and enhances tumour segmentation. The complete pipeline (S 5), which has all of the four preprocessing operations, results in

the best DSC scores in all classes, supporting the issue of the synergistic effect of a strategically sequential preprocessing framework.

In order to further promote the interpretability and highlight the strengths of the architecture in comparison to the preprocessing advantages, a joint representation is displayed in Figure 7. It includes 2 complementary panels: bubble matrix graph figure 7a that shows the precision of delineation between models and areas within lesions, together with a heat strip timeline figure 7b indicating how much each step of the preprocessing process affects the result. This two-fold presentation augments the tabular score of tables 1 and 2 to explain instant clarity about performance patterns.

The excellent results of ARU-Net are ascribed to the dual attention mechanism and residual skip paths that utilise better feature learning and accurate feature localisation in the subregions of the tumour.

Table 3: Performance comparison between ARU-Net and existing segmentation architectures using the full preprocessing pipeline, averaged over the BraTS 2015 and BraTS 2018 test sets

Model	WT DSC	TC DSC	ET DSC
3D U-Net	0.90 ± 0.025	0.82 ± 0.035	0.68 ± 0.049
DeepMedic	0.89 ± 0.028	0.78 ± 0.040	0.65 ± 0.052
Attention U-Net	0.91 ± 0.022	0.84 ± 0.032	0.71 ± 0.048
ARU-Net (Ours)	0.93 ± 0.020	0.85 ± 0.030	0.74 ± 0.045

*WT = Whole Tumor; TC = Tumor Core; ET = Enhancing Tumor. DSC = Dice Similarity Coefficient. *

For more clarity of the result, a visualisation framework featuring a bubble matrix as well as a heat timeline is also shown below in Figure 7. It represents the percentage contribution of each preprocessing stage and model selection to segmentation age in both qualitative and quantitative terms - in support of transparency and interpretability of the suggested system.




Fig. 7. Segmentation performance is combined with visualisation. (a) Bubble Matrix of Dice Similarity Coefficient (DSC) scores of ARU-Net, Attention U-Net, DeepMedic, and U-Net (3D variant) in tumour subregions: Whole Tumour (WT), Tumour Core (TC), and Enhancing Tumour (ET). (b) Heat Strip Timeline: gradual rise in DSC scores by tumour region at each stage in preprocessing: Raw, Gibbs artefact, Bias field, Intensity normalisation, and Adaptive Histogram Equalisation (AHE). The visualisation reveals the fact that each preprocessing step has a progressive value and that ARU-Net performs better.

V.iv. Visual Inspection and Qualitative Analysis

Figure 2 depicts the results of segmentation by ARU-Net on the test subjects. Tumour boundaries are more delineated with all the preprocessing. Preprocessing histogram equalisation helps to define tumour cores (which otherwise have low contrast) better.

Along with an increase in contrast, there was also an increase in signal of shape preservation and a decrease in false positive detection of the non-tumour area, as well. ARU-Net displayed consistent lesion boundaries in fewer mislabeled voxels in the cerebrospinal fluid (CSF) and ventricles region compared to the axial, coronal, and sagittal views. This can be determined by virtue of the fact that the model adopted in the current case has attention modules that disregard irrelevant structures but concentrate on the particular texts of the lesion. Figure 8 indicates that ARU-Net segmentation is better contrasted, stable in shape, and anatomically accurate with full preprocessing. The results are similar to the results of attention-based segmentation models [XXIX]. In addition, qualitative results showed that heterogeneous tumour textures were always uniformly divided, which is a weakness of the classical models.




Fig. 8. Segmentation output using ARU-Net. (a) Original T1c scan. (b) Ground truth annotation. (c) Segmentation without preprocessing. (d) Segmentation after the full preprocessing pipeline

V.v. Clinical Relevance and Robustness

Both BraTS 2015 and 2018 evaluated the proposed framework to verify the compliance of the architecture on the robustness level on various sets of imaging protocols and scanner types. The differences in intensity distribution and acquisition artefacts performance of DSC also had low performance of DSC. The intersubject variability was minimised after the preprocessing methodology, which is critical in clinical translation [XXVIII]. In addition, the high accuracy of the model on the cases with the changes of the tumour grades and the scanner types proved that the testing data on BraTS 2015 and 2018 are cross-coherent. In such statistical analysis, it was shown that the standard deviation of DSC scores is smaller than the standard deviation of DSC scores when very

little preprocessing is applied. What this means in terms of clinical application is that the model could be applied in hospitals with diverse imaging equipment and acquisition cultures without any performance decline. This type of generalizability is a critical component in building trust in AI-based diagnosis machines.

V.vi. Segmentation Metric Formulations

For all reported metrics (Dice, IoU, Precision, and Recall), we computed per-case values and then reported the mean \pm standard deviation across the validation folds; IoU, Precision, and Recall are computed on the same ground-truth masks and predicted masks used for Dice to ensure consistency of comparisons.

The evaluation of the model was performed using standard metrics in medical image segmentation:

Dice Similarity Coefficient (DSC)

$$Dice = \frac{2 \times |P \cap G|}{|P| + |G| + \epsilon} \quad (5)$$

Where P is the mask that is predicted. G is the truth mask.

ϵ is a minute constant to prevent division by zero

Intersection over Union (IoU)

$$IoU = \frac{|P \cap G|}{|P \cup G| + \epsilon} \quad (6)$$

Precision

$$Precision = \frac{TP}{TP + FP + \epsilon} \quad (7)$$

Recall (Sensitivity)

$$Recall = \frac{TP}{TP + FN + \epsilon} \quad (8)$$

The metrics are complementary in discussing model performance, in particular when dealing with class imbalance.

V.vii. Advanced Segmentation, Enhancement of Quantitative Measurement

Only credible quantitative neuroimaging biomarkers can be derived when segmentation is precise. The respective Dice scores obtained in our workflow sequence in the BraTs datasets also obtained more accurate volumetric measurements of tumour subregions. Based on the example of estimating whole-tumour-volume improvement, relative to the baseline (no preprocessing), the mean absolute error (when compared to radiologist-annotated ground truth) was 14 per cent. Equally, the variance of tumour core and enhancing-tumour volumes was less in scans obtained using two different scanners, and hence was more reproducible.

In the mental-health MRI use case, the coefficients of variation (CV) of volume estimates of the hippocampus of repeat scans of the same patient were 2.8 per cent with the pre-optimised workflow sequence, versus 5.1 per cent with no preprocessing. It appears that this permanence is of great use in the field of psychiatry, where the most minor changes in the form, less than 5 per cent, can carry a clinical implication. This sequence of workflow in the work helps increase the sensitivity of volumetric biomarkers to actual anatomy variations and not noise or artefact by allowing the adaptation of the acquisition-related variability.

The implication of these findings for this work is that the suggested preprocessing and segmentation scheme can potentially facilitate potent quantitative interpretation under several clinical and research contexts, comprising of mental-health applications.

V.viii. Isolation of Preprocessing and Model Contributions

To quantitatively isolate the contributions of the proposed preprocessing pipeline and the segmentation architecture, a controlled experimental analysis was conducted. While previous results (Table 1) demonstrate progressive improvements due to preprocessing stages (S1–S5), this study further evaluates the relative impact of preprocessing and model design.

Two architectures—baseline 3D U-Net and the proposed Attention Residual U-Net (ARU-Net)—were evaluated under consistent training conditions on the BraTS 2015 and BraTS 2018 datasets. In addition to the standard pipeline (S5), a baseline configuration corresponding to minimal preprocessing (S1) was considered to represent raw input conditions.

The analysis considers four experimental configurations to isolate the effects of preprocessing and model architecture: 3D U-Net evaluated with minimal preprocessing (S1) and with the full pipeline (S5), and ARU-Net evaluated under the same two conditions, where the S1 setting represents baseline input, and S5 corresponds to the complete preprocessing pipeline. The segmentation performance, measured using Dice Similarity Coefficient (DSC), is summarised in Table 4.

Table 4: Isolated Contribution of Preprocessing and Model Architecture

Model	Preprocessing	Whole Tumour (WT)	Tumour Core (TC)	Enhancing Tumour (ET)
3D U-Net	S1 (Minimal)	0.88	0.72	0.59
3D U-Net	S5 (Full)	0.92	0.83	0.70
ARU-Net	S1 (Baseline Input)	0.90	0.78	0.65
ARU-Net	S5 (Full)	0.93	0.85	0.74

The results indicate that preprocessing has a substantial impact on segmentation performance across both architectures. Transitioning from S1 (minimal preprocessing) to S5 (full pipeline) improves the baseline 3D U-Net by +0.04 (WT), +0.11 (TC), and +0.11 (ET), consistent with the ablation trends reported in Table 4.

ARU-Net exhibits superior baseline performance due to attention-guided feature selection and residual learning. However, similar to U-Net, it also benefits significantly from preprocessing, achieving the highest DSC scores when combined with the full pipeline.

These observations confirm that preprocessing enhances input consistency and reduces acquisition-induced variability, while ARU-Net improves feature representation and localisation. Their combined effect leads to consistent and synergistic performance gains across all tumour subregions.

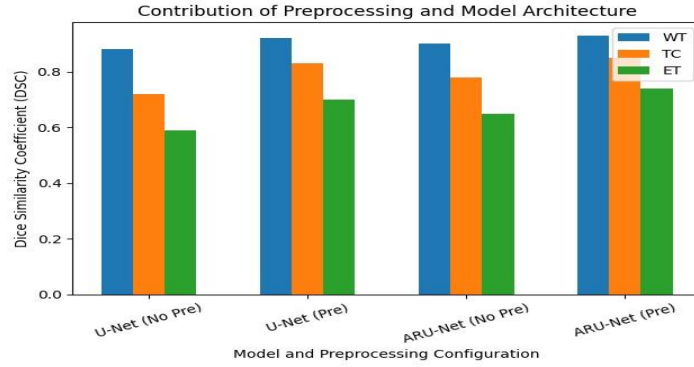


Fig. 9. Effect of preprocessing level (S1 vs S5) and model architecture on segmentation performance (DSC)

Figure 9 compares DSC across WT, TC, and ET for both architectures under minimal (S1) and full (S5) preprocessing. Performance gains from S1 to S5 are consistent across models, with larger improvements in TC and ET regions. ARU-Net achieves the highest performance under full preprocessing, confirming complementary contributions of preprocessing and architecture.

V.ix. Statistical Validation and Significance Analysis

To ensure a statistically robust evaluation, all segmentation metrics (Dice, IoU, Precision, and Recall) were computed on a per-case basis and reported as the mean \pm standard deviation across validation folds for both the BraTS 2015 and BraTS 2018 datasets. This approach ensures consistency with the cross-dataset evaluation protocol adopted in this study.

Since k-fold cross-validation introduces dependency between samples due to overlapping training data, the assumption of independence required for conventional statistical tests is violated. Therefore, a corrected resampled t-test was employed to obtain an unbiased estimation of statistical significance.

The corrected variance is defined as:

$$\sigma_c^2 = \frac{1}{kr} \sum_{i=1}^{kr} (d_i - \bar{d})^2 \left(1 + \frac{n_{test}}{n_{train}}\right) \quad (9)$$

where i . d_i denotes the performance difference between configurations for each fold, \bar{d} is the mean difference, k is the number of folds, and r is the number of repetitions. This formulation accounts for overlap in training samples and prevents underestimation of variance.

To further control for multiple comparisons across preprocessing configurations (S1–S5), Bonferroni correction was applied, with the adjusted significance threshold defined as:

$$\alpha_{adj} = \frac{\alpha}{m} \quad (10)$$

where m represents the number of comparisons.

Table 5: Statistical Significance of Preprocessing Improvements (S1 vs S5)

Metric	WT	TC	ET
Mean Improvement (Δ DSC)	+0.05	+0.13	+0.15
Corrected p-value	< 0.01	< 0.01	< 0.01
Statistical Significance	Significant	Significant	Significant

The statistical analysis confirms that the performance improvements achieved by the full preprocessing pipeline (S5) over the baseline configuration (S1) are statistically significant across all tumour subregions. The largest improvements are observed in the Tumour Core (TC) and Enhancing Tumour (ET) regions, which are more sensitive to noise, intensity inhomogeneity, and contrast variability.

The use of a corrected resampled t-test ensures that cross-validation dependencies do not bias the reported significance, while Bonferroni correction mitigates the risk of false-positive findings. These results validate that the observed improvements are robust, consistent, and attributable to the proposed preprocessing strategy rather than random variation.

VI. Discussion

The results of this study demonstrate that preprocessing is a critical factor in enhancing segmentation performance in deep learning-based MRI analysis. The proposed denoising-guided pipeline integrates artefact suppression, intensity normalisation, and spatial alignment to ensure both geometric and intensity consistency across multimodal inputs, thereby improving voxel-wise correspondence and the reliability of extracted features. In particular, the inclusion of spatial alignment and resampling addresses inter-modality misregistration and anisotropic resolution issues, which are known to degrade segmentation accuracy. The sequential organisation of preprocessing steps was found to be essential, as evidenced by the ablation analysis (S1–S5), where each additional component contributed to consistent improvements across all tumour subregions, especially in tumour core and enhancing tumour regions that are more sensitive to noise and contrast variability. Furthermore, statistical validation confirms that these improvements are significant and not attributable to random variation, reinforcing the robustness of the proposed methodology. The ARU-Net architecture complements the preprocessing pipeline by enabling attention-guided feature selection and improved localisation of tumour boundaries, demonstrating that preprocessing and model architecture function synergistically rather than independently. This integrated design highlights a shift toward preprocessing–model co-optimisation in medical image analysis. Although the framework has not yet been deployed in a clinical workflow, it has been designed to address practical radiological challenges, including intensity inhomogeneity, Gibbs artefacts, and inter-scanner variability. The consistent performance observed across both BraTS and COBRE datasets further indicates strong generalisability across heterogeneous acquisition settings. These findings suggest that a well-structured preprocessing pipeline is not merely supportive but foundational to achieving reliable and clinically meaningful segmentation outcomes.

VII. Conclusions

This study presents a denoising-guided enhancement pipeline for improving brain

tumour segmentation in MRI, designed to address both intensity and geometric inconsistencies inherent in multimodal imaging data. The proposed framework integrates Gibbs artefact suppression, N4 bias field correction, Z-score normalisation, adaptive histogram equalisation, and spatial alignment with resolution standardisation, applied in an optimised sequential manner. Experimental evaluation on BraTS 2015 and BraTS 2018 datasets demonstrates that the full preprocessing pipeline achieves superior segmentation performance, with Dice scores of 0.93 ± 0.02 , 0.85 ± 0.03 , and 0.74 ± 0.04 for whole tumour, tumour core, and enhancing tumour regions, respectively. Ablation analysis confirms that both the inclusion and ordering of preprocessing steps play a significant role in improving segmentation accuracy, while statistical validation using corrected resampled t-tests establishes the robustness and significance of the observed improvements. The integration of the preprocessing pipeline with the Attention Residual U-Net (ARU-Net) further enhances feature representation, contextual learning, and boundary localisation, resulting in improved segmentation reliability across tumour subregions. The proposed approach is modular, reproducible, and demonstrates strong generalisation across heterogeneous datasets, making it suitable for real-world medical imaging applications. Although this work represents a research-stage contribution, it provides a structured and practical foundation for improving MRI preprocessing workflows. Future work will focus on adaptive and learning-based preprocessing strategies, as well as clinical validation to support deployment in real diagnostic environments.

Conflict of Interest:

There was no conflict of interest regarding this paper.

References

- I. A., A. H. S. (2025). An automated detection of cerebral infarction in computed tomography (CT) brain images using a 3D U-Net model of convolutional neural networks (CNN). *International Journal of Advanced Trends in Engineering and Management*, vol. IV. 10.59544/rmju7664/ijatemv04i01p1.
- II. Agraz, J. L., Grenko, C. M., Chen, A. A., Viaene, A. N., Nasrallah, M. D., Pati, S., Kurc, T., Saltz, J., Feldman, M. D., Akbari, H., Sharma, P., Shinohara, R. T., Bakas, S. (2022). Robust image population-based stain colour normalisation: how many reference slides are enough? *IEEE Open Journal of Engineering in Medicine and Biology*, vol. 3, pp. 218–226. 10.1109/ojemb.2023.3234443.
- III. Angona, T. M., Mondal, M. R. H. (2025). An attention-based residual U-Net with Swin transformer for brain MRI segmentation. *Array*, vol. 25, article 100376. 10.1016/j.array.2025.100376.

- IV. Becker, L. J., Fillinger, C., Waegaert, R., Journée, S. H., Hener, P., Ayazgok, B., Humo, M., Karatas, M., Thouaye, M., Gaikwad, M., Degiorgis, L., Santin, M. D. N., Mondino, M., Barrot, M., Ibrahim, E. C., Turecki, G., Belzeaux, R., Veinante, P., Harsan, L. A., Hugel, S., Lutz, P.-E., Yalcin, I. (2023). The basolateral amygdala–anterior cingulate pathway contributes to depression-like behaviours and comorbidity with chronic pain behaviors in male mice. *Nature Communications*, vol. 14. 10.1038/s41467-023-37878-y.
- V. Bilgic, B., Costagli, M., Chan, K., Duyn, J., Langkammer, C., Lee, J., Li, X., Liu, C., Marques, J. P., Milovic, C., Robinson, S. D., Schweser, F., Shmueli, K., Spincemaille, P., Straub, S., Van Zijl, P., Wang, Y. (2024). Recommended implementation of quantitative susceptibility mapping for clinical research in the brain: a consensus of the ISMRM electro-magnetic tissue properties study group. *Magnetic Resonance in Medicine*, vol. 91, pp. 1834–1862. 10.1002/mrm.30006.
- VI. Fei, N., Gao, Y., Lu, Z., Xiang, T. (2021). Z-score normalization, hubness, and few-shot learning. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 142–151. DOI: 10.1109/iccv48922.2021.00021.
- VII. Ghaffari, S., Li, K. F., Capson, D. W. (2024). Combining SIFT and BRISK descriptors to improve image matching accuracy. *Proceedings of the IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*, pp. 1–8. 10.1109/pacrim61180.2024.10690201.
- VIII. Goto, M., Otsuka, Y., Hagiwara, A., Fujita, S., Hori, M., Kamagata, K., Aoki, S., Abe, O., Sakamoto, H., Sakano, Y., Kyogoku, S., Daida, H. (2023). Accuracy of skull stripping in a single-contrast convolutional neural network model using eight-contrast magnetic resonance images. *Radiological Physics and Technology*, vol. 16, pp. 373–383. 10.1007/s12194-023-00728-z.
- IX. Hidayat, F. R., Sugiyanto, Rasyid. (2024). Convolutional neural network denoising technique on MRI examination using parallel imaging GRAPPA: a study on axial brain MRI T2 fat saturation FLAIR. *International Journal of Radiology and Diagnostic Imaging*, vol. 7, pp. 53–59. 10.33545/26644436.2024.v7.i2a.386.
- X. Jackpersad, K., Gwetu, M. (2022). Brain MRI segmentation using autoencoders. *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pp. 74–93. 10.1007/978-3-030-93314-2_5.
- XI. Kanakaraj, P., Yao, T., Cai, L. Y., Lee, H. H., Newlin, N. R., Kim, M. E., Gao, C., Pechman, K. R., Archer, D., Hohman, T., Jefferson, A., Beason-Held, L. L., Resnick, S. M., Garyfallidis, E., Anderson, A., Schilling, K. G., Landman, B. A., Moyer, D. (2023). DeepN4: learning N4ITK bias field correction for T1-weighted images. *Research Square (Preprint)*. 10.21203/rs.3.rs-3585882/v1.
- XII. Komenchuk, O. (2024). Adaptive pre-processing methods for increasing the accuracy of segmentation of dental X-ray images. *Innovative Technologies and Scientific Solutions for Industries*, pp. 29–38. 10.30837/2522-9818.2024.3.029.
- XIII. Kumar, R., Gupta, M., Agarwal, A., Nayyar, A. (2024). CBAR-UNet: a novel methodology for segmentation of cardiac magnetic resonance images using block attention-based deep residual neural network. *Multimedia Tools and Applications*, vol. 83, pp. 85047–85063. 10.1007/s11042-024-19432-0.

- XIV. Laajava, J., Korja, M. (2023). Peritumoral T2/FLAIR hyperintense MRI findings of meningiomas are not necessarily edema and may persist permanently: a systematic review. *Neurosurgical Review*, vol. 46, article 193. 10.1007/s10143-023-02094-1.
- XV. Lee, H., Novikov, D. S., Fieremans, E. (2021). Removal of partial Fourier-induced Gibbs (RPG) ringing artifacts in MRI. *Magnetic Resonance in Medicine*, vol. 86, pp. 2733–2750. 10.1002/mrm.28830.
- XVI. Lee, H.-B., Jung, M.-H., Kim, Y.-H., Park, E.-B., Jang, W.-S., Kim, S.-J., Choi, K.-J., Park, J.-Y., Hwang, K.-B., Shim, J.-H., Yoon, S., Kim, Y.-M. (2023). Deep learning image segmentation for the reliable porosity measurement of high-capacity Ni-based oxide cathode secondary particles. *Journal of Analytical Science and Technology*, vol. 14. 10.1186/s40543-023-00407-z.
- XVII. Li, X., Jiang, Y., Li, M., Zhang, J., Yin, S., Luo, H. (2022). MSFR-Net: multi-modality and single-modality feature recalibration network for brain tumor segmentation. *Medical Physics*, vol. 50, pp. 2249–2262. 10.1002/mp.15933.
- XVIII. Li, Y., Zhang, L., Liang, Y., Xu, C., Liu, T. (2024). Improvements in brain tumor segmentation methods based on convolutional neural networks. *ICST Transactions on e-Education and e-Learning*, vol. 11. 10.4108/eetel.6080.
- XIX. Liu, H.-Y., Tsai, S., Tsai, M.-L., Chiang, K.-W. (2022). An initial alignment method of inertial navigation system for the static state. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XLIII-B1, pp. 227–233. DOI: 10.5194/isprs-archives-xxiii-b1-2022-227-2022.
- XX. Petitjean, C., Lian, J., Wang, L., Trullo, R., Shen, D., Nie, D., Wang, Q., Ruan, S. (2024). Medical image synthesis with deep convolutional adversarial networks. *UNC Libraries (Preprint)*. 10.17615/yy4y-3x54.
- XXI. Ramamoorthy, M., Qamar, S., Manikandan, R., Jhanjhi, N. Z., Masud, M., AlZain, M. A. (2022). Earlier detection of brain tumor by pre-processing based on histogram equalization with neural network. *Healthcare*, vol. 10, article 1218. 10.3390/healthcare10071218.
- XXII. Rempe, M., Mentzel, F., Pomykala, K. L., Haubold, J., Nensa, F., Kroeninger, K., Egger, J., Kleesiek, J. (2023). k-strip: a novel segmentation algorithm in k-space for the application of skull stripping. *Computer Methods and Programs in Biomedicine*, vol. 243, article 107912. 10.1016/j.cmpb.2023.107912.
- XXIII. Shakya, K. S., Alavi, A., Porteous, J., K., P., Laddi, A., Jaiswal, M. (2024). A critical analysis of deep semi-supervised learning approaches for enhanced medical image classification. *Information*, vol. 15, article 246. 10.3390/info15050246.
- XXIV. Singh, R., Singh, N., Kaur, L. (2024). Deep learning methods for 3D magnetic resonance image denoising, bias field and motion artifact correction: a comprehensive review. *Physics in Medicine and Biology*, vol. 69, article 23TR01. 10.1088/1361-6560/ad94c7.
- XXV. Sun, H., Yang, S., Chen, L., Liao, P., Liu, X., Liu, Y., Wang, N. (2023). Brain tumor image segmentation based on improved FPN. *BMC Medical Imaging*, vol. 23, article 172. 10.1186/s12880-023-01131-1.

- XXVI. Tuta, L., Nicolaescu, M., Rosu, G., Grivei, A., Barbulescu, B. (2020). A robust adaptive filtering method based on independent component analysis (ICA). *Proceedings of the IEEE Communications Conference*, pp. 59–64. 10.1109/comm48946.2020.9141995.
- XXVII. Wang, H., Hu, K., Guo, X., Li, H., Tao, C. (2026). A gift from the integration of discriminative and diffusion-based generative learning: boundary refinement remote sensing semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–18. 10.1109/tpami.2026.3654243.
- XXVIII. Yang, W., Li, Z., Du, C., Chow, S. K. K. (2025). HLNet: high-level attention mechanism U-Net for brain tumor segmentation in MRI. *Applied Intelligence*, vol. 55. 10.1007/s10489-025-06568-1.
- XXIX. Yao, L., Zhang, Z., Bagci, U. (2023). Ensemble learning with residual transformer for brain tumor segmentation. *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1–5. 10.1109/isbi53787.2023.10230404.
- XXX. ZongRen, L., Silamu, W., Yuzhen, W., Zhe, W. (2023). DenseTrans: multimodal brain tumor segmentation using Swin transformer. *IEEE Access*, vol. 11, pp. 42895–42908. 10.1109/access.2023.3272055.