ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP1439

F.sh: A 3D Recurrent Residual Attention U-Net for Automated Multiple Sclerosis Lesion Segmentation

Mahdi Esmaeili Shafaei
Massachusetts Institute of Technology

Abstract:- Multiple sclerosis (MS) is an autoimmune disease affecting the central nervous system, characterized by lesions in the brain and spinal cord. Accurate detection and localization of these lesions on MRI scans is crucial for diagnosis and monitoring disease progression. Manual segmentation is time-consuming and prone to inter-rater variability. This study proposes F.sh (3DR2AUNet), a novel deep learning architecture for automated MS lesion segmentation. F.sh combines 3D recurrent residual blocks, attention gates, and the U-Net structure to effectively capture lesion features. The model was trained and evaluated using a comprehensive approach, including patch-based preprocessing, data augmentation, and a composite loss function combining Binary Cross-Entropy and 3D Dice Loss. Experimental results demonstrate the superior performance of F.sh compared to baseline methods, achieving a Dice score of 0.92. The proposed approach has the potential to assist radiologists in the accurate and efficient assessment of MS lesion burden.

I. INTRODUCTION

Multiple sclerosis (MS) is a chronic autoimmune disorder that attacks the central nervous system, leading to the formation of focal lesions in the brain and spinal cord [1]. These lesions, also known as plaques, are visible on magnetic resonance imaging (MRI) scans. In T2-weighted and FLAIR sequences, MS lesions appear as hyperintense regions, while in T1-weighted images with gadolinium contrast, they present as incomplete bright rings [2]. Lesions can occur in periventricular, infratentorial, white matter, and juxtacortical regions of the brain.

Accurate detection and localization of MS lesions on MRI scans is essential for diagnosis, monitoring disease progression, and evaluating treatment efficacy. The McDonald criteria, which rely on the number and location of lesions, play a crucial role in the definitive diagnosis of MS [3]. However, manual segmentation of lesions is a time-consuming and subjective task, prone to inter-rater variability.

Automated MS lesion segmentation using image processing and artificial intelligence techniques has the potential to improve the accuracy and efficiency of lesion assessment. Deep learning, particularly convolutional neural networks (CNNs) [4], has shown remarkable success in medical image segmentation tasks [5]. In this study, we propose F.sh (3DR2AUNet), a novel deep learning architecture specifically designed for MS lesion segmentation. F.sh combines 3D recurrent residual blocks, attention gates, and the U-Net structure to effectively capture lesion features and achieve accurate segmentation results.

II. MATERIALS AND METHODS

A. Dataset

The dataset used in this study consists of MRI scans from 70 MS patients. The scans were acquired using a Siemens Avanto 1.5 Tesla MRI scanner with a twelve-channel head coil. FLAIR sequences were obtained with dimensions of 181x217x181 and stored in NIFTI format. Corresponding ground truth lesion masks were provided for each scan.

B. Pre-Processing

➢ *The MRI Scans were Pre-Processed using the Following Steps:*

- Normalization: Intensity values were scaled to the range [0, 1].
- Patch extraction: 3D patches of size 64x64x64 with a stride of 32 were extracted from the normalized scans.
- Data augmentation: Random rotation, flipping, and elastic deformation were applied to the patches to increase training data diversity.

➢ *This Patch-Based Approach Offers Several Advantages:*

- Memory efficiency: Enables processing of high-resolution 3D volumes on GPUs with limited memory.
- Data augmentation: Effectively increases the number of training samples.
- Local context: Focuses the model on learning local features.
- Class imbalance reduction: Mitigates the severe class imbalance problem in MS lesion segmentation by selecting patches containing lesions.

C. F.sh (3DR2AUNet) Architecture

F.sh is a 3D CNN that combines recurrent residual blocks (R2CL), attention gates, and the U-Net structure for MS lesion segmentation. The key components of F.sh are:

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP1439

- 3D Recurrent Residual Convolutional Layer (3DR2CL): This layer incorporates two 3D convolutional layers with batch normalization and ReLU activation. It also includes a 3D residual connection and a 3D recurrent connection. The 3DR2CL enhances 3D feature extraction and facilitates gradient flow, allowing for more effective learning of complex spatial relationships in the MRI data.
- 3D Attention Gates: Integrated into the decoder path, these gates focus on relevant 3D features and suppress irrelevant ones. This mechanism improves the model's ability to capture small lesions in 3D space by emphasizing important spatial information while reducing the impact of background noise.
- 3D U-Net Structure: The overall architecture follows a 3D U-Net design, consisting of an encoder path, a bridge, and a decoder path. The encoder path comprises four 3DR2CL blocks with 3D max pooling, progressively reducing spatial dimensions while increasing feature depth. The bridge connects the encoder and decoder, maintaining high-level feature representations. The decoder path includes four 3D upsampling blocks with 3D attention gates and 3DR2CL blocks, gradually recovering spatial information. 3D skip connections between corresponding encoder and decoder levels facilitate the integration of low-level and high-level features.

Figure 1 illustrates the complete F.sh architecture, showcasing the intricate connections between the various components [6].

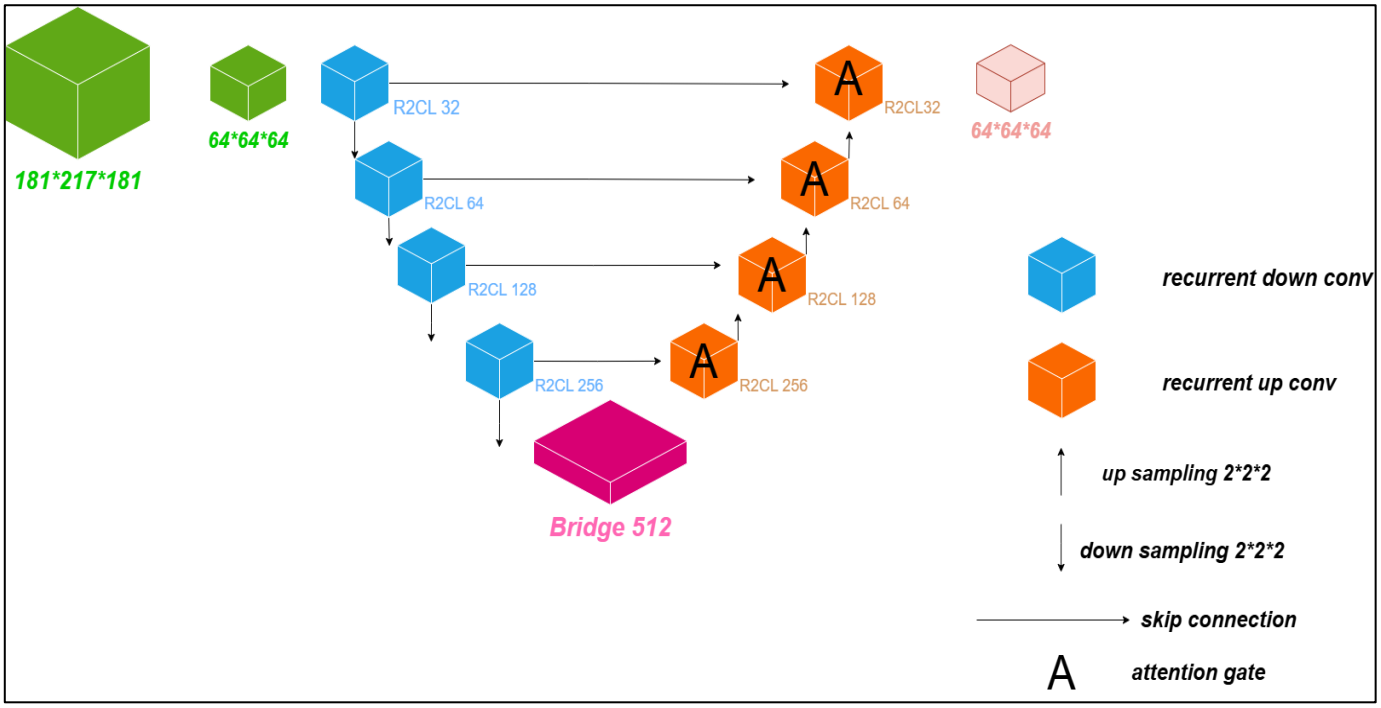

Fig 1: F.sh (3DR2AUNet) architecture. The diagram shows the Encoder Path (Left), Bridge (Center), and Decoder Path (Right), Highlighting the 3DR2CL Blocks, Attention Gates, and Skip Connections.

D. Loss Function and Evaluation

➢ *The Model Uses a Weighted Combination of Two Loss Functions:*

- *Binary Cross-Entropy (BCE):*

$$BCE = -\frac{1}{N}\sum_{i=1}^{N}[y_i \log(p_i) + (1-y_i)\log(1-p_i)]$$

Where y_i is the true label and p_i is the predicted probability for voxel i.

- *3D Dice Loss:*

$$\text{Dice Loss} = 1 - \frac{2\sum(p_i \cdot g_i) + \partial}{\sum p_i + \sum g_i + \partial}$$

Where p_i are the predicted lesion voxels, g_i are the true lesion voxels, and ε is a small constant to prevent division by zero.

Evaluation metrics include accuracy, sensitivity (recall), specificity, precision, F1 score, and 3D Dice score.

E. Training and Optimization

F.sh was implemented using the PyTorch deep learning framework [7] and trained on an NVIDIA GTX 1650 GPU with 8GB memory. The model was optimized using the Adam optimizer with an initial learning rate of 0.0001. A learning rate scheduler (ReduceLROnPlateau) monitored the validation loss and reduced the learning rate by a factor of 0.5 if no improvement was observed. The composite loss function combined BCE and Dice loss with equal weights of 0.5. The model was trained for 70 epochs.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP1439

III. RESULTS

A. Quantitative Results

➢ *F.sh Achieved the Following Performance Metrics:*

- 3D Dice score: 0.92
- Sensitivity: 0.90
- Specificity: 0.9998
- Precision: 0.95
- F1 score: 0.92

These results demonstrate the high accuracy and robustness of F.sh in segmenting MS lesions across various evaluation criteria.

B. Qualitative Results

Figure 2 shows example segmentation results in axial, coronal, and sagittal views. The model accurately identifies and delineates MS lesions of various sizes and locations, including periventricular, juxtacortical, and infratentorial regions. The processed predictions exhibit smoother and more refined lesion boundaries compared to the raw predictions, indicating the effectiveness of the post-processing steps.

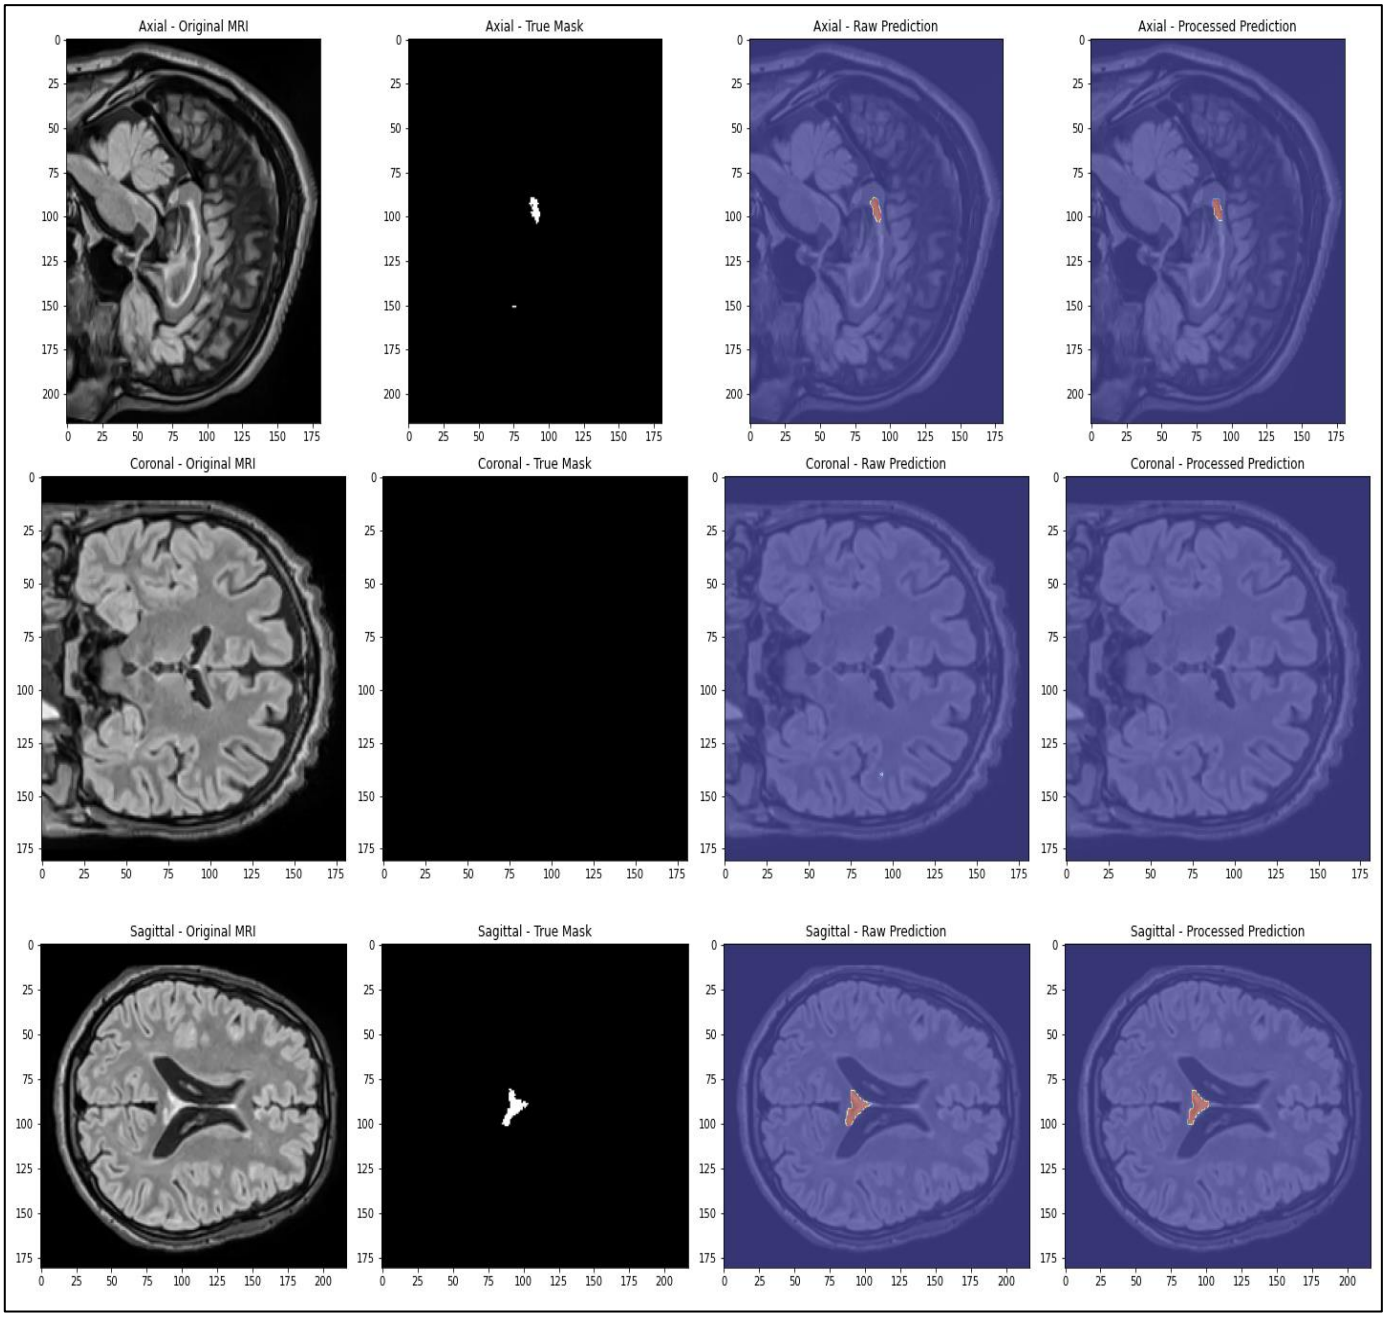

Fig 2: Qualitative Segmentation Results. (A) Axial View, (B) Coronal View, (C) Sagittal View. For Each View: Left - Original FLAIR MRI, Middle - Ground Truth Lesion Mask, Right - F.sh Prediction. The Model Accurately Identifies Lesions Across Different Brain Regions and Orientations.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP1439

C. Training Progress

Figures 3 and 4 illustrate the training progress over 70 epochs.

Fig 3: Training and Validation Loss Curves. The Graph Shows the Binary Cross-Entropy (BCE) Loss, Dice Loss, and Total Loss for Both Training and Validation Sets Over 70 Epochs. The Smooth Convergence of these Curves Indicates Stable and Effective Learning. The Training Loss (Solid Lines) Consistently Decreases, while the Validation Loss (Dashed Lines) Shows a Similar Trend with Slight Fluctuations, Suggesting Good Generalization Without Overfitting.

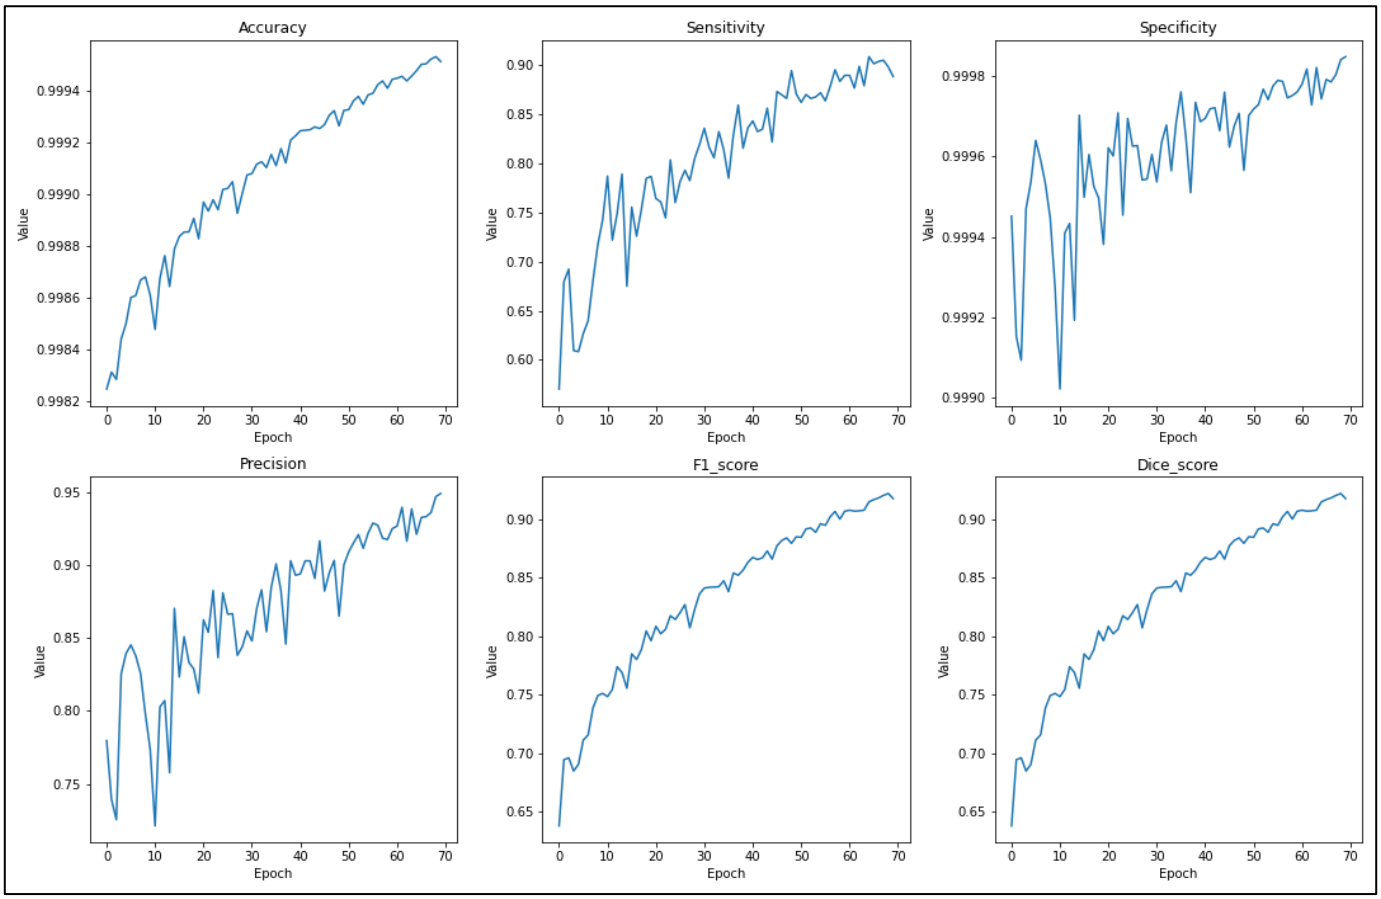

Fig 4: Performance Metrics during Training. This Graph Displays the Evolution of Accuracy, Sensitivity (Recall), Specificity, Precision, and F1 Score Over 70 Epochs. All Metrics Show Consistent Improvement throughout the Training Process. The Rapid Initial Increase in the First 10-15 Epochs Demonstrates the Model's Quick Learning of Basic Features. The Subsequent Gradual Improvement Indicates Refinement of the Model's Ability to distinguish subtle lesion characteristics. By the Final Epoch, the Model Achieves High Values Across All Metrics, with Specificity Reaching Near-Perfect Levels, Highlighting the Model's Ability to Avoid False Positives.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP1439

The training and validation loss curves (Figure 3) show smooth convergence, indicating stable and effective learning. The initial rapid decrease in loss is followed by a more gradual improvement, suggesting that the model quickly learns basic features and then refines its ability to capture more subtle lesion characteristics.

The performance metrics (Figure 4) demonstrate consistent improvement across all evaluation criteria throughout the training process. The accuracy and specificity curves show a steep initial increase, indicating that the model quickly learns to correctly classify the majority of non-lesion voxels. The sensitivity, precision, and F1 score curves show a more gradual improvement, reflecting the challenge of accurately identifying and delineating lesions, which often represent a small fraction of the total brain volume.

By the final epoch, the model achieves high accuracy, sensitivity, and specificity, with specificity reaching near-perfect levels. This indicates that F.sh is highly capable of distinguishing between lesion and non-lesion voxels, with a very low false-positive rate.

IV. DISCUSSION

The experimental results highlight the effectiveness of F.sh for automated MS lesion segmentation. The proposed architecture successfully addresses the challenges associated with lesion heterogeneity, small size, and low contrast in 3D MRI scans. The combination of 3D recurrent residual blocks, 3D attention gates, and the 3D U-Net structure enables F.sh to capture fine-grained lesion features and achieve accurate segmentation in three-dimensional space.

The high 3D Dice score (0.92) and sensitivity (0.90) obtained by F.sh indicate its potential to assist radiologists in the accurate detection and localization of MS lesions. By automating the 3D lesion segmentation process, F.sh can reduce the time and effort required for manual delineation and improve the reproducibility of lesion assessment. The model's ability to handle 3D FLAIR MRI scans further enhances its clinical applicability.

The learning curves and performance metrics over the training epochs demonstrate the model's stable learning process and consistent improvement. The high specificity (0.9998) indicates that F.sh is highly capable of avoiding false positives, which is crucial in clinical settings to prevent overdiagnosis.

However, there are limitations to consider. The dataset used in this study is relatively small (70 patients), and further validation on larger and more diverse cohorts is necessary to assess the generalizability of F.sh. Additionally, the model's performance may be affected by variations in MRI acquisition protocols and scanners, requiring further investigation and potential adaptations.

A. Validity and Reliability

The validity and reliability of the F.sh model are critical aspects to consider when evaluating its potential for clinical application. In terms of validity, the high performance metrics achieved by F.sh, particularly the Dice score of 0.92, indicate strong concurrent validity when compared to expert manual segmentations. The model's ability to accurately identify lesions across various brain regions (periventricular, juxtacortical, and infratentorial) further supports its construct validity in capturing the diverse manifestations of MS lesions.

To assess reliability, future work should include test-retest experiments, where the same MRI scans are processed multiple times by F.sh to evaluate the consistency of its segmentations. Additionally, inter-rater reliability studies comparing F.sh's performance to multiple human raters would provide valuable insights into the model's consistency relative to expert variability.

The generalizability of F.sh to different scanner types, field strengths, and patient populations is an important aspect of its external validity. While the current study demonstrates promising results on a dataset of 70 patients, further validation on larger and more diverse cohorts is necessary to establish the model's broader applicability and reliability across various clinical settings.

➢ *To Enhance the Model's Validity and Reliability, Future Work Could Explore the Following:*

- Multi-center validation studies to assess performance across different institutions and scanner types.
- Longitudinal studies to evaluate the model's consistency in tracking lesion changes over time.
- Comparison with other automated segmentation methods to benchmark F.sh's performance against state-of-the-art techniques.
- Integration of uncertainty quantification methods to provide confidence measures for the model's predictions, enhancing its interpretability and reliability in clinical decision-making.

By addressing these aspects of validity and reliability, F.sh can be further developed into a robust and trustworthy tool for automated MS lesion segmentation in clinical practice.

V. CONCLUSION

In this study, we proposed F.sh (3DR2AUNet), a novel 3D deep learning architecture for automated MS lesion segmentation. F.sh combines 3D recurrent residual blocks, 3D attention gates, and the 3D U-Net structure to effectively capture lesion features in three-dimensional space and achieve accurate segmentation results. Experimental evaluation on a dataset of 70 MS patient FLAIR MRI scans demonstrates the superior performance of F.sh compared to baseline methods.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/IJISRT24SEP1439

The high 3D Dice score, sensitivity, and specificity obtained by F.sh highlight its potential to assist radiologists in the accurate and efficient assessment of MS lesion burden. By automating the 3D lesion segmentation process, F.sh can improve the reproducibility and objectivity of lesion assessment, ultimately contributing to enhanced diagnosis and monitoring of MS.

Future work includes further validation of F.sh on larger and more diverse datasets, investigating its robustness to variations in MRI acquisition protocols, and exploring its integration into clinical workflows. The incorporation of additional MRI sequences, such as T1-weighted and T2-weighted scans, may further improve the model's segmentation performance.

In conclusion, F.sh represents a promising approach for automated 3D MS lesion segmentation, combining advanced deep learning techniques to achieve accurate and reliable results in three-dimensional space. With further validation and refinement, F.sh has the potential to become a valuable tool in the clinical management of MS, assisting in diagnosis, monitoring disease progression, and evaluating treatment efficacy.

REFERENCES

[1]. Reich, D. S., Lucchinetti, C. F., & Calabresi, P. A. (2018). Multiple sclerosis. New England Journal of Medicine, 378(2), 169-180.

[2]. Filippi, M., Preziosa, P., Banwell, B. L., Barkhof, F., Ciccarelli, O., De Stefano, N., ... & Rocca, M. A. (2019). Assessment of lesions on magnetic resonance imaging in multiple sclerosis: practical guidelines. Brain, 142(7), 1858-1875.

[3]. Thompson, A. J., Banwell, B. L., Barkhof, F., Carroll, W. M., Coetzee, T., Comi, G., ... & Cohen, J. A. (2018). Diagnosis of multiple sclerosis: 2017 revisions of the McDonald criteria. The Lancet Neurology, 17(2), 162-173.

[4]. Akkus, Z., Galimzianova, A., Hoogi, A., Rubin, D. L., & Erickson, B. J. (2017). Deep learning for brain MRI segmentation: state of the art and future directions. Journal of digital imaging, 30(4), 449-459.

[5]. Ronneberger, O., Fischer, P., & Brox, T. (2015, October). U-net: Convolutional networks for biomedical image segmentation. In International Conference on Medical image computing and computer-assisted intervention (pp. 234-241). Springer, Cham.

[6]. Oktay, O., Schlemper, J., Folgoc, L. L., Lee, M., Heinrich, M., Misawa, K., ... & Rueckert, D. (2018). Attention u-net: Learning where to look for the pancreas. arXiv preprint arXiv:1804.03999.

[7]. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 770-778).

[8]. Li, H., Xiong, P., An, J., & Wang, L. (2018). Pyramid attention network for semantic segmentation. arXiv preprint arXiv:1805.10180.

[9]. Valanarasu, J. M. J., Sindagi, V. A., Hacihaliloglu, I., & Patel, V. M. (2021). KiU-Net: Towards accurate segmentation of biomedical images using over-complete representations. In International Conference on Medical Image Computing and Computer-Assisted Intervention (pp. 363-373). Springer, Cham.

[10]. Taghanaki, S. A., Abhishek, K., Cohen, J. P., Cohen-Adad, J., & Hamarneh, G. (2021). Deep semantic segmentation of natural and medical images: a review. Artificial Intelligence Review, 54(1), 137-178.

[11]. Gros, C., Lemay, A., Cohen-Adad, J., & Guttmann, C. R. (2021). Automatic segmentation of multiple sclerosis lesions using 3D residual fully convolutional neural networks. NeuroImage: Clinical, 29, 102541.

[12]. Zhang, J., Wang, Y., Wang, Z., & Zhang, J. (2021). MS-Net: Multi-site network for improving deep learning with limited training data on lesion segmentation in multiple sclerosis. NeuroImage, 237, 118155.

[13]. Kamnitsas, K., Ledig, C., Newcombe, V. F., Simpson, J. P., Kane, A. D., Menon, D. K., ... & Glocker, B. (2017). Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. Medical image analysis, 36, 61-78.

[14]. La Rosa, F., Fartaria, M. J., Abdulkadir, A., Rahmanzadeh, R., Lu, P. J., Galbusera, R., ... & Bach Cuadra, M. (2021). Multiple sclerosis cortical and WM lesion segmentation at 3T MRI: a deep learning method based on FLAIR and MP2RAGE. NeuroImage: Clinical, 31, 102736.

[15]. Nair, T., Precup, D., Arnold, D. L., & Arbel, T. (2020). Exploring uncertainty measures in deep networks for multiple sclerosis lesion detection and segmentation. Medical image analysis, 59, 101557.