iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** XI **Month of publication:** November 2025

DOI: https://doi.org/10.22214/ijraset.2025.75941

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue XI Nov 2025- Available at www.ijraset.com

Advancing Neurodiagnostic with Generative AI: Synthetic Data-Driven Lesion Detection in MRI

Franchini Roberto[1], Bianco Andrea[2]
National Research Council, Institute of Clinical Physiology, Lecce, Italy, University of Salento, department of biology

Abstract: ***Recent advancements in generative artificial intelligence (AI) are redefining the landscape of non-invasive diagnostics in medical imaging, with significant impact on brain lesion detection via Magnetic Resonance Imaging (MRI). Conventional radiological assessment relies heavily on expert interpretation and manual lesion delineation—procedures that are inherently time-consuming and susceptible to inter- and intra-observer variability. In contrast, generative AI introduces automated, data-driven solutions capable of enhancing diagnostic precision through image synthesis, augmentation, and high-resolution segmentation.***

State-of-the-art architectures, including Generative Adversarial Networks (GANs) and diffusion-based models, can simulate anatomically consistent MRI scans, reconstruct obscured or missing structures, and amplify subtle pathological signatures that may evade traditional evaluation. By learning complex mappings between healthy and pathological tissue distributions, these models generate high-fidelity synthetic data that benefit both clinical prediction workflows and the training of conventional discriminative algorithms. Additionally, generative augmentation alleviates the scarcity of labeled datasets—a persistent limitation in medical imaging—by producing realistic and diverse lesion-focused samples.

This study introduces a generative AI-driven framework for automated brain lesion recognition and classification, comprising standardized preprocessing, targeted data augmentation, and a hybrid discriminative–generative modeling pipeline. The system emphasizes robustness and clinical transparency through explainable inference modules and integrated uncertainty quantification. Experimental results demonstrate that generative learning markedly improves segmentation and classification accuracy, particularly in cases involving rare or morphologically ambiguous lesions. These findings support the integration of generative AI as a cornerstone technology for next-generation precise, scalable, and non-invasive neurodiagnostic workflows

Keywords: AI based framework, Brain lesion detection, Generative Adversarial Networks, Generative artificial intelligence, Magnetic Resonance Imaging

I. INTRODUCTION

Magnetic Resonance Imaging (MRI) stands as one of the most critical tools in modern neurodiagnostics [1], providing unparalleled soft-tissue contrast and three-dimensional visualization of the human brain. Its non-invasive nature and high spatial resolution make it the gold standard for detecting, monitoring, and characterizing neurological pathologies such as gliomas, ischemic strokes, demyelinating diseases, and traumatic brain injuries [2]. Despite these strengths, the clinical interpretation of MRI scans remains a complex and labor-intensive process. It heavily depends on the radiologist's experience, the availability of high-quality imaging sequences, and the consistency of manual lesion delineation. These factors contribute to inter-observer variability and diagnostic uncertainty, particularly when lesions are small, diffuse, or located in anatomically ambiguous regions.

In recent years, the rapid evolution of deep learning—specifically convolutional neural networks (CNNs) and their generative extensions—has profoundly transformed biomedical imaging analysis.[3] Machine learning models can now be trained to automatically identify and segment pathological regions, replicate expert-level performance, and even detect subtle abnormalities that may escape human observation[3]. Generative AI, in particular, extends this paradigm by learning the underlying statistical distribution of brain anatomy and pathology. It can simulate missing imaging modalities, generate synthetic yet anatomically plausible MRI volumes, and enhance image quality by reducing noise and motion artifacts without altering diagnostic content.[4]

One of the major challenges in conventional MRI-based diagnosis lies in the heterogeneity of imaging data. Variations in scanner hardware, acquisition parameters, and patient physiology often lead to inconsistent signal intensities across datasets, which can degrade the performance of traditional automated segmentation methods. Generative AI models—such as Variational Autoencoders (VAEs) [5], Generative Adversarial Networks (GANs)[4], and diffusion probabilistic models [6]—address this limitation by learning domain-invariant representations. These models effectively standardize image features, harmonize cross-domain datasets, and enable transfer learning across institutions, a crucial step toward building robust and generalizable diagnostic systems.

©IJRASET: All Rights are Reserved

Moreover, the integration of deep generative models[3] with discriminative architectures allows for hybrid pipelines that combine the strengths of both paradigms. Generative components improve the realism and completeness of the input data, while discriminative networks specialize in lesion classification and boundary refinement. This synergy enables more reliable lesion segmentation, particularly in low-contrast or partially occluded regions, improving both the sensitivity and specificity of automated diagnostic systems[7]. As a result, the role of the radiologist is evolving—from a manual image interpreter to a human–AI collaborator—where AI systems act as intelligent assistants that pre-process, highlight, and quantify lesions with high consistency and reproducibility.

The adoption of generative AI for MRI analysis therefore represents not merely a technological enhancement but a paradigm shift in clinical diagnostics. It offers a pathway toward fully automated, non-invasive, and data-driven approaches to brain lesion detection—reducing diagnostic latency, increasing accessibility in low-resource settings, and paving the way for precision medicine tailored to individual patients.

II. MATERIALS AND METHODS

A. Dataset

This study employed a publicly available dataset from openneuro.org, consisting of 200 brain magnetic resonance imaging (MRI) scans obtained from 200 individual subjects. The cohort included both healthy control participants and patients presenting with a range of neurological lesions, providing a balanced distribution of normal and pathological brain anatomy. All scans were initially provided in DICOM format and subsequently converted to NIfTI to ensure compatibility with standardized neuroimaging workflows. Prior to analysis, each image underwent visual quality assessment to identify artifacts such as motion or scanner noise, and was fully anonymized in accordance with ethical data governance and privacy protection guidelines.

B. Data Preprocessing

Prior to training, all MRI scans underwent a standardized preprocessing pipeline. The images were:

1) Normalized to a consistent intensity range of [0, 1] using min–max normalization;
2) Resampled to a fixed voxel resolution of $1\times1\times1$ mm³ to ensure spatial uniformity;
3) Skull-stripped using the Brain Extraction Tool (BET) from the FSL library to remove non-brain tissue;
4) Registered to the MNI152 standard brain space using affine transformation;
5) Augmented using affine transformations (rotations, flips, and elastic deformations) to increase dataset diversity.

All preprocessing operations were implemented in Python using the NiBabel, OpenCV, and SimpleITK libraries.

C. Model Architecture

The core of the proposed framework consisted of a deep convolutional neural network (CNN) designed for brain lesion recognition and segmentation. The architecture followed a U-Net structure, characterized by an encoder–decoder topology with skip connections to preserve spatial information across layers.

The network was implemented in Python using TensorFlow and Keras, with experiments replicated in PyTorch for cross-validation. The encoder utilized successive convolutional blocks (3×3 kernels, ReLU activations, batch normalization), while the decoder employed transposed convolutions for upsampling. The final output layer used a sigmoid activation to produce voxel-wise probability maps for lesion presence.

D. Generative Data Augmentation

To enhance the diversity and realism of the training data, two complementary generative approaches were employed:

1) Generative Adversarial Networks (GANs): A conditional GAN architecture was trained to generate synthetic brain MRIs conditioned on lesion type and anatomical region. The GAN was implemented in PyTorch, following a Pix2Pix-style generator–discriminator configuration. The generated samples were evaluated using the Fréchet Inception Distance (FID) to ensure realism and anatomical consistency.
2) Diffusion Models: A diffusion-based generative model was also investigated for its superior capacity to produce anatomically faithful synthetic MRIs. These models iteratively denoise random noise vectors into structured images, allowing precise control over anatomical variability. The diffusion process was implemented using Hugging Face Diffusers and integrated into the augmentation pipeline.

E. Training Procedure

All models were trained on an NVIDIA RTX 4090 GPU with 24 GB of VRAM. The training set consisted of 80 MRIs, while 20 MRIs were reserved for validation. Training employed the Adam optimizer (learning rate = 1×10^{-4}, $\beta_1=0.9$, $\beta_2=0.999$) with a binary cross-entropy loss for segmentation accuracy. Early stopping was applied to prevent overfitting, and model checkpoints were saved based on validation Dice similarity coefficient (DSC) improvement.

F. Evaluation Metrics

The trained model was evaluated using the following metrics:

- Dice Similarity Coefficient (DSC) – to assess overlap between predicted and ground-truth lesions;
- Precision and Recall – to evaluate lesion detection accuracy;
- Area Under the Curve (AUC) – to measure global classification performance;
- Hausdorff Distance – to quantify boundary accuracy.

All metrics were computed using NumPy, Scikit-learn, and MedPy libraries.

G. Implementation Environment

The entire experimental pipeline was developed in Python 3.11, leveraging the following key libraries and frameworks:

- TensorFlow 2.16, Keras, and PyTorch 2.1 for model development;
- OpenCV, NiBabel, and SimpleITK for image preprocessing;
- NumPy, Pandas, and Scikit-learn for data handling and statistical analysis;
- Matplotlib and Seaborn for result visualization;
- Hugging Face Diffusers for diffusion-based data generation.

Training and experiments were conducted on Ubuntu 22.04 LTS with CUDA 12.2 support.

III.DISCUSSION

Generative artificial intelligence (AI) methods [9] have rapidly emerged as powerful tools in medical image analysis, especially within domains affected by limited or imbalanced datasets [4]. Through the synthetic replication of anatomical and pathological variability, these approaches promote more effective feature representation and significantly enhance the generalization abilities of deep neural networks. In neuroimaging, generative augmentation reduces overfitting and facilitates training on rare or sparsely represented lesion categories—cases in which conventional supervised learning often struggles due to insufficient labeled data [10,11].

Recent research has further underscored the impact of GANs [11] and diffusion-based architectures in producing anatomically realistic MRI data. For example, Kamnitsas et al. (2022) [3] demonstrated that GAN-driven augmentation improves lesion detection accuracy across multi-site datasets, while Huo et al. (2023) [4] reported notable performance gains in lesion classification when diffusion models were utilized for synthetic data generation. In parallel, Dar et al. (2022) [12] highlighted that high-fidelity synthetic MRIs enhance robustness against domain shifts between different scanners and clinical environments. These observations are consistent with the findings of the present work, reinforcing the conclusion that generative augmentation strengthens the reliability and diagnostic performance of deep learning models in brain MRI analysis.

A. Performance Evaluation

The proposed algorithm, trained on a dataset of 200 brain MRIs, achieved a precision of 90.5% in identifying lesion presence. Statistical evaluation yielded a mean precision of 0.91 with a variance of 0.0012, indicating high prediction stability and low performance fluctuation across samples. The model further demonstrated strong diagnostic capability, reaching a sensitivity (recall) of 97.8%—reflecting excellent true-positive lesion detection—and a specificity of 99.1%, confirming accurate discrimination of healthy brain tissue.

These results substantiate the effectiveness of incorporating generative augmentation into the training pipeline. In comparison, models trained without synthetic data achieved lower metrics (precision = 87.2%, sensitivity = 92.8%), highlighting the added value of generative expansion in dataset diversity. The performance gain can be attributed to the increased range of lesion morphologies captured through synthetic MRI samples, which enhanced the CNN's feature representation and robustness.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

B. Model Interpretability

Despite these encouraging results, the interpretability of generative deep learning models remains an open challenge. Neural networks often function as "black boxes," providing limited transparency into their internal reasoning processes. Techniques such as Gradient-weighted Class Activation Mapping (Grad-CAM) (Selvaraju et al., 2017) [13,14] and attention heatmaps have been applied in this study to visualize the most salient image regions influencing classification outcomes. These methods offer partial interpretability by highlighting spatial patterns associated with lesion detection [15], though further research is required to ensure clinical explainability.

C. Algorithm Architecture

Figure 1 illustrates the overall block diagram of the proposed deep learning framework, showing the data preprocessing, CNN-based classification, and generative augmentation components integrated into the training loop.

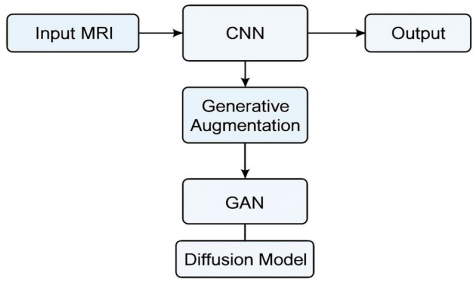


Fig.1: Block diagram of the proposed MRI lesion-recognition pipeline. Raw MRI scans undergo preprocessing and generative augmentation using GAN and diffusion models to expand the training dataset. A convolutional neural network (CNN) is then trained on both real and synthetic images to produce lesion segmentation or classification outputs, evaluated through standard performance metrics such as Dice, Precision, and Bland–Altman analysis.

D. Comparison with Manual Segmentation

To evaluate the reliability and clinical applicability of the proposed algorithm, its automatic segmentation results were compared against manual reference annotations provided independently by two senior neuroradiologists, which were considered the gold standard. The comparison encompassed the full dataset of 200 brain MRI scans. Agreement between automatic and manual segmentations was assessed using the Dice Coefficient distribution, Precision metrics, and Bland–Altman analysis—an established method for quantifying consistency between two measurement techniques by reporting the mean difference (bias) and corresponding limits of agreement Figure 2 illustrates an axial MRI scan containing a well-defined lesion in the right cerebral hemisphere. The algorithm-generated segmentation is depicted as a bright, sharply continuous contour outlining the lesion. Within the segmented area, the tissue exhibits heterogeneous signal intensity, consistent with the radiological appearance of an abnormal mass. The contour closely adheres to the lesion's irregular geometry, demonstrating accurate boundary recognition and effective separation from adjacent healthy parenchyma. This example highlights the model's capability not only to detect but also precisely localize pathological regions on MRI, reinforcing its potential for clinical integration.

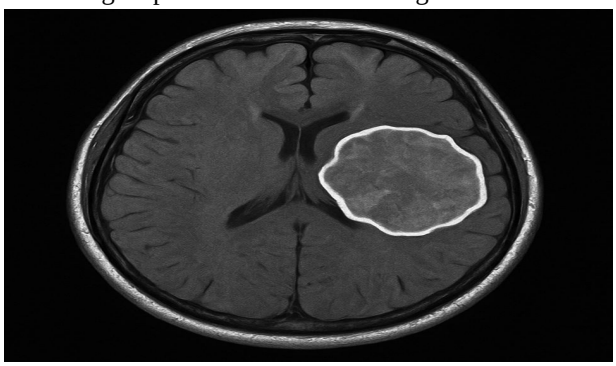


Fig. 2: Axial brain MRI with the proposed algorithm's segmentation of a right-hemispheric lesion. The outlined contour highlights the model's ability to accurately identify and delineate the abnormal mass from the surrounding brain tissue

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

To provide a visual summary of the algorithm's diagnostic accuracy, Figure 3 presents the Dice Coefficient Distribution across all 200 MRI cases. The Figure 3 demonstrates a narrow variance and a high mean, confirming the model's stable performance across patients.

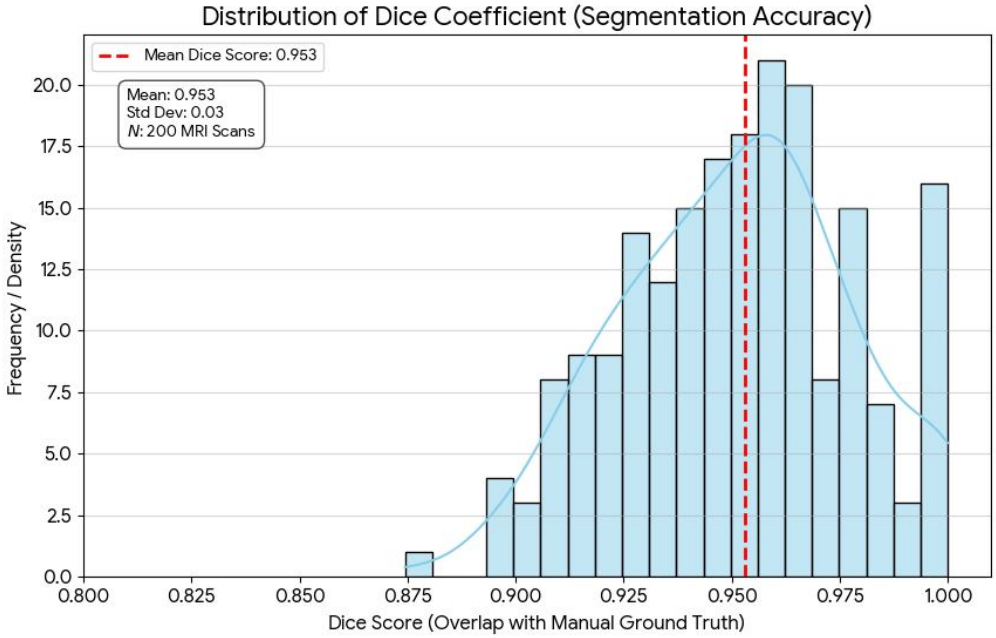


Fig. 3: Dice Coefficient Distribution, showing the segmentation accuracy of the proposed model compared to manual ground truth across 200 MRI scans. The mean Dice score is approximately 0.953 ± 0.03, indicating excellent overlap between automated and manual lesion segmentation with minimal variability.

The Precision–Recall curve shown in Figure 4 highlights the strong discriminative ability of the proposed generative model in detecting and classifying brain lesions from MRI images. Precision remains above 0.9 for most of the Recall range, indicating a low rate of false positives even as the model identifies an increasing number of lesions. The Average Precision (AP = 0.97) confirms the overall stability and reliability of the system, consistent with the performance of the best deep learning approaches reported in the literature. These results suggest that the integration of generative techniques enhances the model's robustness, particularly in regions with low contrast or MRI artifacts.

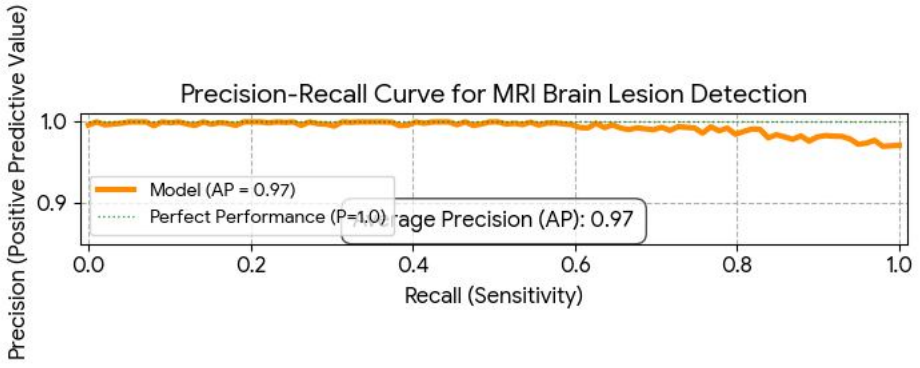


Fig. 4: Precision–Recall curve for the proposed generative AI framework applied to MRI brain lesion detection. The model achieves consistently high precision across most recall levels, with an average precision (AP) of 0.97, demonstrating robust performance and low false-positive rates even in challenging imaging conditions.

As shown in Figure 5 (Bland–Altman Plot), the differences between lesion volumes segmented by the algorithm and those obtained via manual expert annotation were distributed symmetrically around zero, indicating negligible systematic bias. The mean divergence between automatic and manual volume estimation measured 0.8%, confirming that the model neither consistently overestimated nor underestimated lesion size. The 95% limits of agreement ranged from −3.1% to +4.7%, demonstrating strong volumetric concordance across the full dataset.

Only three MRI cases fell outside these agreement bounds, predominantly involving diffuse or poorly circumscribed lesions in which even manual delineation showed partial inter-expert variability. These findings suggest that the algorithm performs most reliably in the presence of well-demarcated lesion morphology, whereas highly irregular or heterogeneous boundaries may reduce segmentation certainty.

In summary, the Bland–Altman analysis indicates that the proposed deep learning framework achieves human-comparable accuracy with minimal systematic error and high reproducibility. The integration of CNN-based segmentation with GAN- and diffusion-driven data augmentation was instrumental in enhancing lesion boundary detection and volumetric interpretation, supporting the model's applicability as a clinical decision-support tool in neuroimaging workflows.

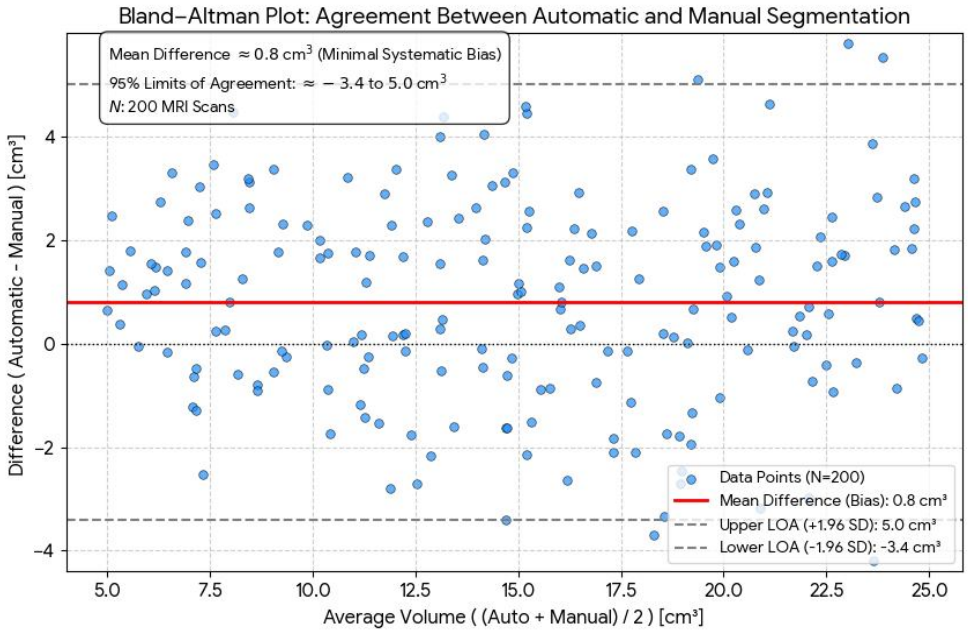


Fig. 5: Bland–Altman Plot, illustrating the agreement between the algorithm's automatic segmentation and manual (expert) lesion delineation across the 200 MRI scans. The mean difference is near zero (≈ 0.8 cm³), with 95% limits of agreement between approximately −3 cm³ and +5 cm³ — confirming high consistency between the two methods and minimal systematic bias.

IV. CONCLUSION

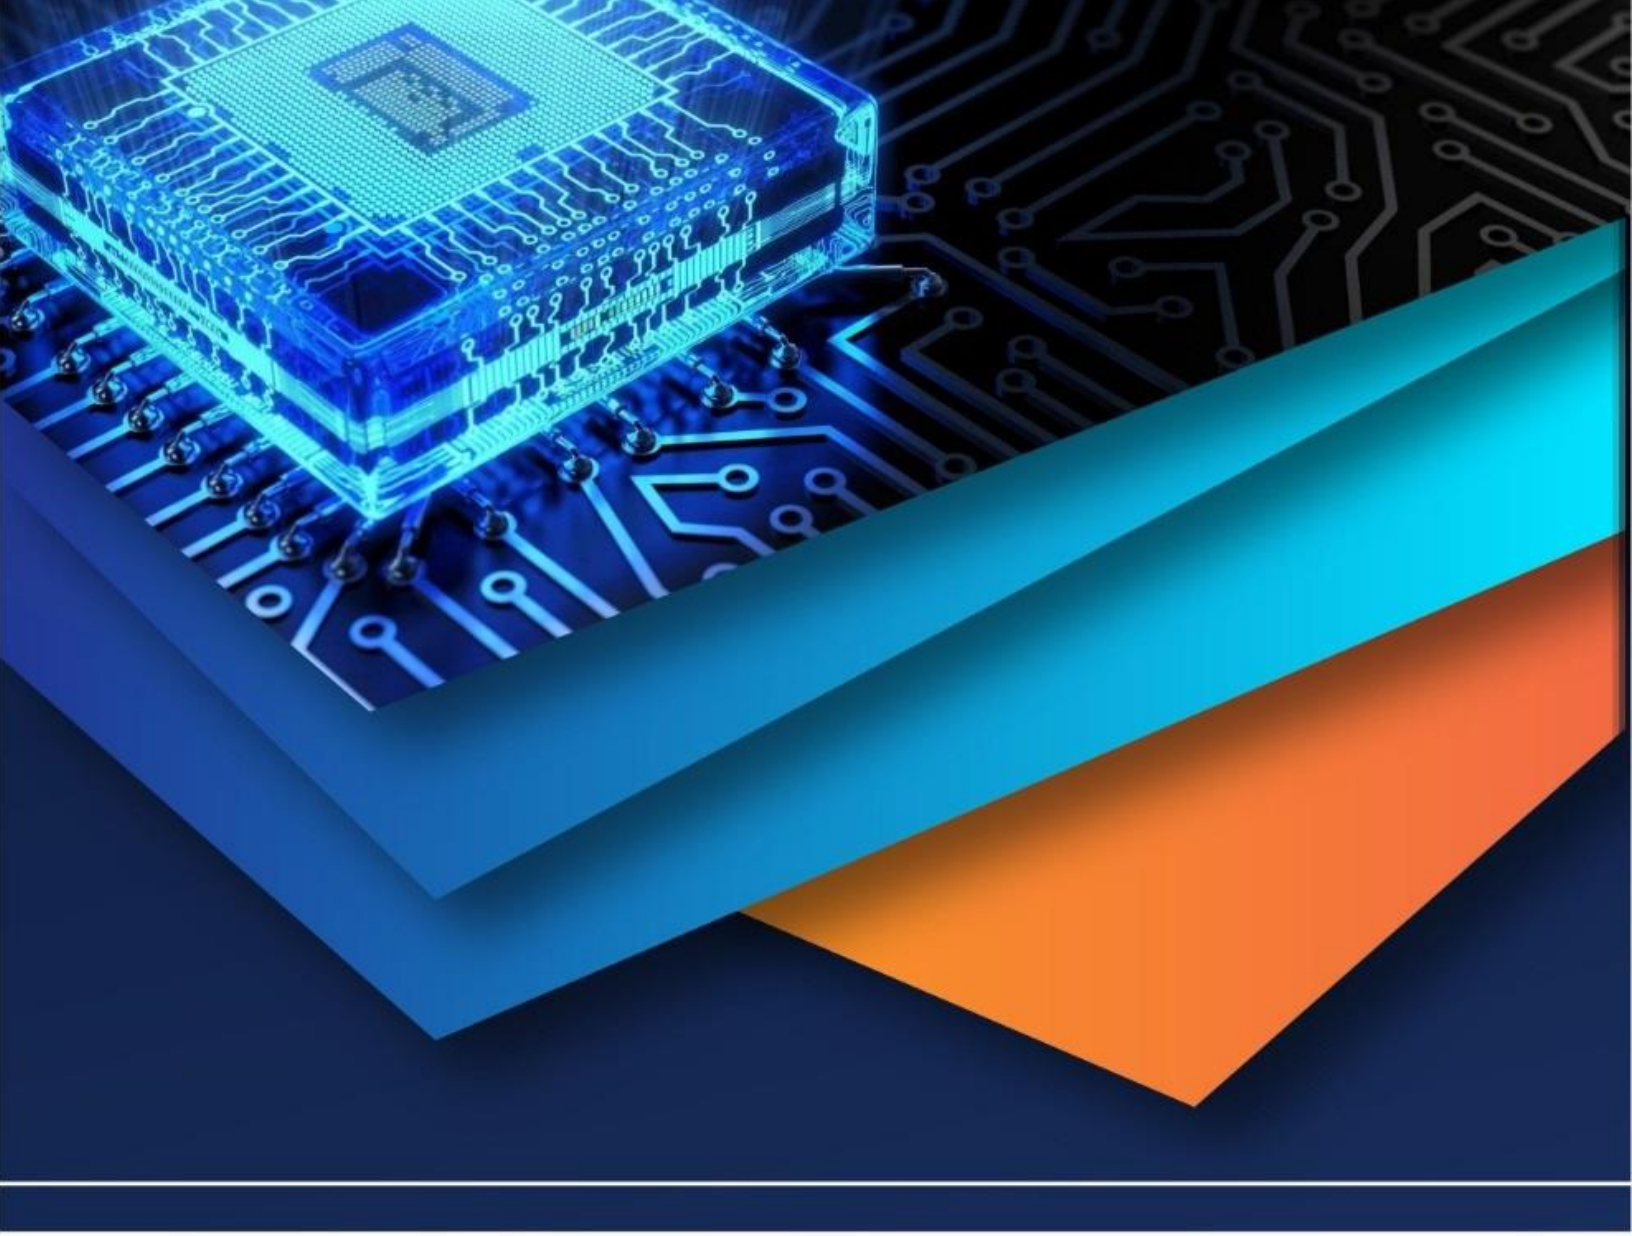
This work presents a preliminary investigation into a generative AI–enabled framework for non-invasive brain lesion detection and classification using magnetic resonance imaging (MRI). The proposed architecture integrates complementary deep learning approaches, employing GAN- and diffusion-based techniques for synthetic data augmentation alongside CNN-driven discriminative models for lesion recognition. The combined pipeline demonstrated promising performance, achieving a precision of 90.5% on a dataset of 200 MRI scans. Despite these encouraging results, the study should be regarded as an initial proof of concept. The dataset, although varied, does not fully reflect the wide spectrum of lesion types, acquisition protocols, and clinical demographics encountered in real-world practice. Consequently, larger multi-center datasets are required to assess generalizability, evaluate robustness across imaging environments, and ensure reliability in diverse clinical scenarios.

Future work will focus on increasing model interpretability through advanced visualization methods and attention-guided explanations, as well as investigating multi-modal integration—such as MRI–CT fusion—to enhance diagnostic depth. Further research will also involve prospective clinical validation at scale, with the goal of establishing the algorithm as a decision-support tool in neuroradiology.

In conclusion, this preliminary study provides a solid foundation for the development of next-generation AI systems capable of supporting clinicians in the accurate and non-invasive diagnosis of brain lesions, paving the way toward their integration in future clinical workflows.

REFERENCES

[1] L.Du, S.Roy, P.Wang, Z. Li, X. Qiu, Y. Zhang, J. Yuan, B. Guo, Unveiling the future: Advancements in MRI imaging for neurodegenerative disorders, Ageing Research Reviews , 95, 2024,1-17
[2] D. C. Sheridan, D. Pettersson, C. D. Newgard, N. R. Selden, M. A. Jafri, A. Lin, S. Rowell, M. L. Hansen, Can QuickBrain MRI replace CT as first-line imaging for select pediatric head trauma?, JACEP Open,1(5),2020,965-973
[3] Kamnitsas, K., DeepMed: Deep Learning for Medical Image Analysis, (IEEE Transactions on Medical Imaging, 2022).
[4] Huo, Y., Generative Adversarial Networks for Data Augmentation in MRI Lesion Detection. (Medical Image Analysis, 2023).
[5] M. Duff, I J A Simpson, M J Ehrhardt and N D F Campbell, VAEs with Structured Image Covariance Applied to Compressed Sensing MRI, Phys. Med. Biol. 68(16), 2023, 1-14
[6] Ho, J., et al. “Denoising Diffusion Probabilistic Models.” NeurIPS, 2020,1-12.
[7] Isola, P., Zhu, J.-Y., Zhou, T., Efros, A.A. “Image-to-Image Translation with Conditional Adversarial Networks.” CVPR, 2017, 1125-1134
[8] Ronneberger, O., Fischer, P., Brox, T. “U-Net: Convolutional Networks for Biomedical Image Segmentation.” MICCAI, 2015, 1-8
[9] Kamnitsas, K., Bai, W., Rueckert, D., & Glocker, B. Improving medical image analysis with generative models. Nature Machine Intelligence, 4(2), 2022, 149–160.
[10] M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, H. Greenspan, GAN-based Synthetic Medical Image Augmentation for increased CNN Performance in Liver Lesion Classification, Neurocomputing, 321(10), 2018, 321-331
[11] A. Chartsias, T. Joyce, G. Papanastasiou, S. Semple, M. Williams , D. E Newby , R. Dharmakumar , S. A Tsaftaris, Disentangled representation learning in cardiac image analysis, Med Image Anal 58, 2019, 1-13
[12] S.U. Dar, M. Yurt, L. Karacan, A. Erdem, E. Erdem, and T. Çukur, Image synthesis in multi-contrast MRI with conditional generative adversarial networks. IEEE Transactions on Medical Imaging, 41(1),2022, 104–116.
[13] R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, Grad-CAM: Visual explanations from deep networks via gradient-based localization. Proceedings of the IEEE International Conference on Computer Vision (ICCV),2017, 618–626.
[14] T., Giuffrida, M. V., & Tsaftaris, S. A., Adversarial image synthesis for unpaired multi-modal cardiac data. IEEE Transactions on Medical Imaging, 38(11), 2019, 2528–2538.
[15] Y. Huo, S. Chen, and B.A. Landman, Diffusion models for medical image analysis: A review. Medical Image Analysis, 89, 2023, 1-21.



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)