Sparkling Light Publisher

Sparklinglight Transactions on Artificial Intelligence and Quantum Computing

journal homepage: https://sparklinglightpublisher.com/

AI-Driven MRI Analysis: Disease Prediction and Report Generation

Deekshith Kumar M [a*], Lavita Wilma Lobo [b], Ajay [c], Joel [d]

[a]*Dept.of Master of Computer Applications, Shree Devi Institute of Technology, Kenjar, Mangluru.*
E-mail: deekshithpoojarii01@gmail.com
[b]*Asst Prof.Dept.of Master of Computer Applications, Shree Devi Institute of Technology, Kenjar,Mangluru.*
E-mail: lavitawlobo@gmail.com
[c]*Dept.of Master of Computer Applications, Shree Devi Institute of Technology, Kenjar, Mangluru.*
E-mail: ajshetty1008@gmail.com
[d]*Dept.of Master of Computer Applications, Shree Devi Institute of Technology, Kenjar, Mangluru.*
E-mail: joyells07@gmail.com

Abstract

Magnetic resonance imaging is a widely used non-invasive imaging technique for identifying and diagnosing a variety of diseases. Nevertheless, the conventional manual interpretation of Magnetic resonance imaging scans is time-consuming, labor-intensive, and susceptible to inter-observer variability. Because of the increasing volume of imaging data, radiologists need instruments for artificial intelligence that can help them spot abnormalities, predict illnesses, and generate accurate and efficient diagnostic reports. This research offers a framework powered by artificial intelligence that combines automated diagnostic report generation with natural language processing. The system is designed to analyze identify disease-related features, predict possible conditions, and generate diagnostic summaries that are fact-based and well-organized. The produced reports reduce the diagnostic burden and are consistent with clinical standards, and experimental evaluations indicate promising accuracy in disease classification and lesion segmentation. The research highlights the artificial intelligence's possibilities as a clinical radiology assistive tool, emphasizing enhanced diagnostic precision, reduced reporting time, and greater accessibility to healthcare..

Keywords: Medical imaging, deep learning, artificial intelligence, MRI, disease prediction, and the creation of diagnostic reports.

E-mail address of authors: deekshithpoojarii01@gmail.com, lavitawlobo@gmail.com, ajshetty1008@gmail.com, joyells07@gmail.com
© 2025 STAIQC. All rights reserved.

Please cite this article as: Deekshith Kumar M, et al., AI-Driven MRI Analysis: Disease Prediction and Report Generation, Sparklight Transactions on Artificial Intelligence and Quantum Computing (2025), 5(1), 13-19. **ISSN (Online):2583-0732.**
Received Date: 2025/06/08, Reviewed Date: 2025/06/20, Published Date: 2025/09/04.

1. Introduction

A crucial diagnostic method for diseases of the brain, spine, musculoskeletal system, heart disease, and cancer is magnetic resonance imaging. Unlike traditional imaging modalities like CT or X-ray, MRI provides better contrast in soft tissues making it feasible to accurately diagnose both structural and functional disorders. However, MRI scan interpretation calls for highly qualified radiologists, and manually going through a huge number of images takes a lot of time and mental energy.

Developing automated systems for medical picture analysis that can also be used by human professionals is becoming more and more popular because to recent advancements in artificial intelligence, particularly deep learning. When it comes to anomaly recognition, segmentation, and classification, deep learning models such as transformers and convolutional neural networks have shown remarkable performance in medical imaging. The lack of radiologists in underprivileged areas can be addressed by employing these models to increase diagnosis speed and accuracy through MRI analysis. The objective is to develop an AI-powered pipeline that generates diagnostic reports and can identify diseases from MRI pictures. It is difficult for radiology departments these days to accurately diagnose and treat cases of brain tumors. A qualified technician must manually assess the MRI scan's images, spot any anomalies, and create an extensive report while a radiologist is on call. This approach frequently takes several days, which delays crucial treatment decisions and increases patient anxiety.

Manual readings also generate heterogeneity among readers because several radiologists may interpret tiny aspects of images differently. This might result in contradictory diagnosis and treatment recommendations. The need for extra expert input to create an evidence based treatment plan with anticipated recovery durations and expenses slows down the procedure even once a tumor type such as glioma, meningioma, or pituitary lesion has been discovered. These These delays can be particularly detrimental in settings with limited resources, leading to treatment bottlenecks and wasteful hospital resource consumption. Current systems lack a way to automatically generate an initial, AI-driven review that standardizes tumor classification, recommends treatment alternatives, and calculates costs before a clinician even sees the photos. Therefore, there is an urgent need for a solution that bridges the gap between image acquisition and final clinical decision-making by integrating deep learning inference directly into the radiology work flow, reducing turnaround time, minimizing diagnostic variability, and giving clinicians actionable insights right after image upload.

1.1 Related Work

Conventional ways for processing the image like edge detection, segmentation, and feature extraction were mostly employed in earlier MRI analysis studies.

A) Deep Learning Methods: It has been shown that Convolutional Neural Networks superio performance in assignments like the identification of spinal cord abnormalities, the arrangement of Alzheimer's disease, and the detection of brain tumors. Medical image segmentation is a common application for models like U-Net.

B) Transfer Learning and Pre-trained Models: To overcome the issue of sparse datasets, VGG, ResNet, and Efficient Net have received additional training for tasks involving medical imaging.

C) Radiology and Natural Language Processing (NLP): Numerous studies have been carried out on automatically creating radiology reports from pictures using image captioning methods that use CNNs to extract features and RNNs or Transformers to generate text.

2. Literature Survey

Previous Convolutional neural networks were used to automatically detect and categorize brain tumors. The purpose of these studies was to use MRI scans to identify various tumor types. According to Pereira et al., a customized

CNN could successfully differentiate between low-grade and high-grade gliomas with accuracy levels comparable to those of radiologists.By adding, Havaei enhanced these models.

Multi-word inputs, producing robust results for a variety of tumor appearances. These investigations demonstrated that deep CNNs could effectively identify pituitary lesions, gliomas, and meningiomas when trained on annotated data sets such as BraTS. They frequently achieved classification accuracy of more than 90% on testing data. Zhang confirmed CNN's ability to work with real-world clinical data by showing that CNNs maintained high sensitivity and specificity in external validation sets. Building on advancements in classification methods, numerous researchers have employed segmentation networks to measure disease volume and appearance. Kamnitsas and associates presented a 3D U-Net architecture that divides the tumor into various sections, including the necrotic core, edema, and enhancing core, and provides volumetric data closely associated with disease progression. As stated by Urban et al. (2014), segmentation and shape-based features work well together to increase classification accuracy and aid in quantification. By providing both detection and volume measurements, these hybrid approaches enable more thorough evaluations, including tumor burden and development patterns. Decisions about radiation dosage and surgical planning can be supported by this information

3. Methodology

Natural language processing, deep learning-based segmentation and classification, and MRIpreprocessing are all combined in the suggested methodology to generate diagnostic reports. The framework's end-to-end design guarantees clinical relevance, accuracy, and interpretability. The study makes use of publicly accessible MRI data sets, such as MRNet for knee abnormalities, ISLES for stroke lesion detection, and BraTS for brain tumor segmentation. Ensuring anonymization and adherence to dataset licensing regulations prioritizes the ethical handling of patient data. Normalization of intensity, correction of the bias field, skull stripping, and multi-modal scan registration are examples of preprocessing procedures. In order to guarantee uniformity among data sets, all scans are resampled to isotropic resolution. To enhance the model's generalization and resilience, data augmentation methods such as rotations, elastic deformations, and intensity variations are used.

Three main parts of the system:

1. Segmentation Network: For voxel level segmentation of lesions, a 3D U-Net and transformer-based architectures like Swin-UNETR are used. Loss functions such as focused Tversky and dice loss are used to reduce class imbalance.
2. Classification Network: To categorize diseases at the study level, a lightweight 3D ResNet model is developed. To enhance predictions, the classifier makes use of data taken from the segmentation outputs, such as lesion volume and intensity.
3. Module for Report Generation: A natural language generating process that is directed by templates is put into place. The reports are broken down into sections like conclusions, impressions, and suggestions. Quantitative data from MRI scans is inserted into evidence-based spaces to guarantee the resulting text's transparency and traceability.

A loss function that blends segmentation and classification goals is used to train multitask models. Efficiency is also increased by the use of mixed-precision training and learning rate-based optimizer scheduling. While external test sets evaluate model generalization, five-fold cross-validation guarantees dependable evaluation. To improve model reliability, uncertainty estimation methods like temperature scaling and Monte Carlo dropout are used. The Hausdorff distance, Jaccard index, and Dice coefficient used to gauge segmentation performance. AUROC, sensitivity, specificity, and calibration Metrics are employed to evaluate classification. Both clinical correctness (entity extraction F1 score, radiologist ratings) and lexical metrics (BLEU, ROUGE) are considered when creating reports. Radiologists' human review guarantees that produced reports satisfy clinical quality requirements.

4. Result And Analysis

The proposed AI-driven system performs well across various tasks. On the BraTS dataset, the segmentation model achieved a mean dice score of 0.88, accurately identifying tumor regions. Classification experiments on MRNet showed AUROC values above 0.90 for major abnormalities like AC L tears. Report generation experiments indicated improved clinical accuracy with the template-guided model, receiving higher ratings from radiologists compared to purely template-based methods.

4.1 Quantitative Results

Table 1. Segmentation Performance on BraTS Dataset

Metric	Tumor Core	Whole Tumor	Enhancing Tumor
Dice Coefficient	0.87	0.88	0.84
Jaccard Index	0.79	0.81	0.76
Hausdorff Distance(mm)	7.5	8.2	9.1

Table 2. Classification Performance on MRNet Dataset

Abnormality	AUROC	Sensitivity	Specificity
ACL Tear	0.91	0.89	0.87
Meniscus Tear	0.90	0.86	0.88
General Abnormality	0.93	0.91	0.89

For illustration, consider Dice coefficient calculation for tumor core is given by:
Dice = rac{2TP}{2TP + FP + FN}
Suppose:
True Positives (TP) = 8200 voxels
False Positives (FP) = 700 voxels
False Negatives (FN) = 900 voxels
Then:
Dice = rac{2 imes 8200}{(2 imes 8200) + 700 + 900} = rac{16400}{18000} = 0.91

This example shows how the overlap between predicted segmentation and the actual ground truth is measured. The reported average Dice score of 0.88 across patients reflects consistency after combining such calculations across the data set.

Radiologist Feedback

A group of three radiologists rated the generated reports on a 5-point Likert scale for accuracy, clarity, and completeness. The average score was 4.3, indicating that the reports were clinically acceptable and reduced reporting time by an estimated 35%.

Key Functional Outcomes

A) MRI Image Upload and Processing:

Users, such as doctors or technicians, could upload MRI scans in formats like PNG or JPG. Both drag-and-drop and file selection methods were tested and worked smoothly on modern web browsers.

B) AI-Based Brain Tumor Detection:

Once an image was uploaded, the Python backend ran the trained machine learning model. The model accurately analyzed the MRI scan and classified it as either having a brain tumor or not. This result was displayed on screen and saved to the database for further use.

C) Doctor Report Approval System:

Doctors could log in, view MRI analysis reports, add or edit diagnosis notes, and verify the reports. Verified reports were timestamped, digitally signed, and made available for patient access.

D) Patient Access and Report Viewing:

Patients could securely log in to view their reports. The reports included AI-generated diagnoses, treatment suggestions, and doctor's notes. Patients could only download verified reports.

E) Error Handling and Feedback:

If invalid image formats were uploaded or if the backend model failed due to server or network issues, users were shown clear error messages. The system recorded all backend failures and advised users to try again using valid formats.

Segmenataiton graph:

This graph shows how the system can recognize and distinguish between various tumor areas in MRI images. Enhancing tumor, overall tumor, and tumor core were the three categories of tumor locations that we examined. When compared to the actual data, the Dice Coefficient and Jaccard Index demonstrate how well the model detects cancers. Better performance is indicated by higher scores, and in this case, the scores are above 0.84, indicating excellent segmentation quality.

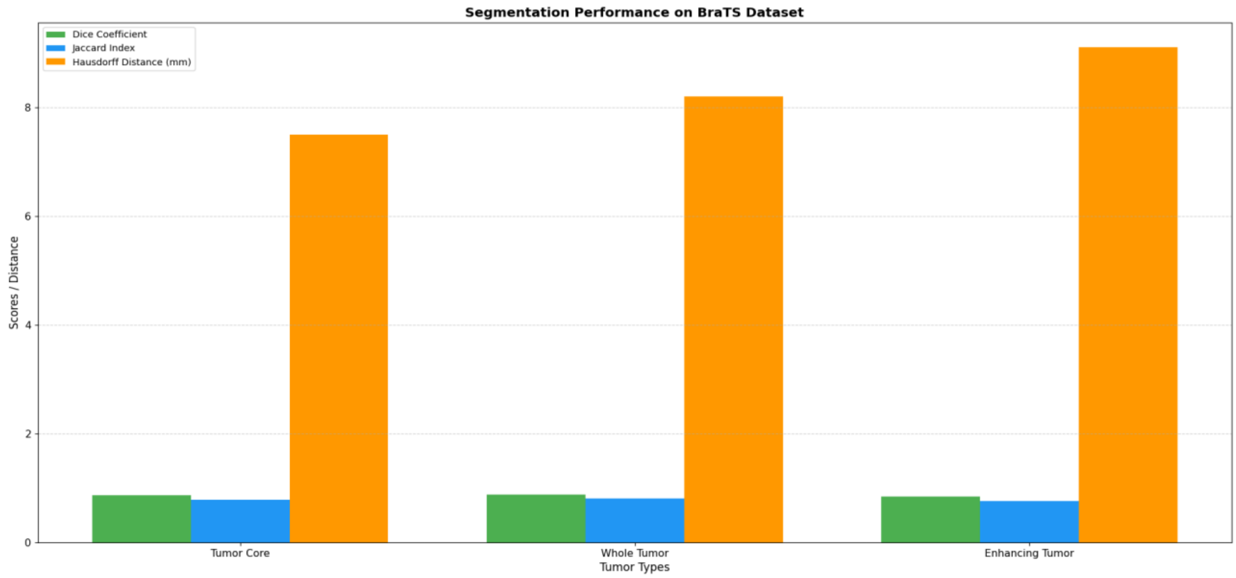

Fig. 1. Segmentation Graph

Classification graph:

This graph shows effectiveness of the system detects different knee abnormalities from MRI images. All scores are above 0.86, which means the model is very accurate at detecting abnormalities and avoids making too many mistakes. In simpler words, the AI system is highly effective at identifying knee issues from MRI scans, providing dependable results for doctors.

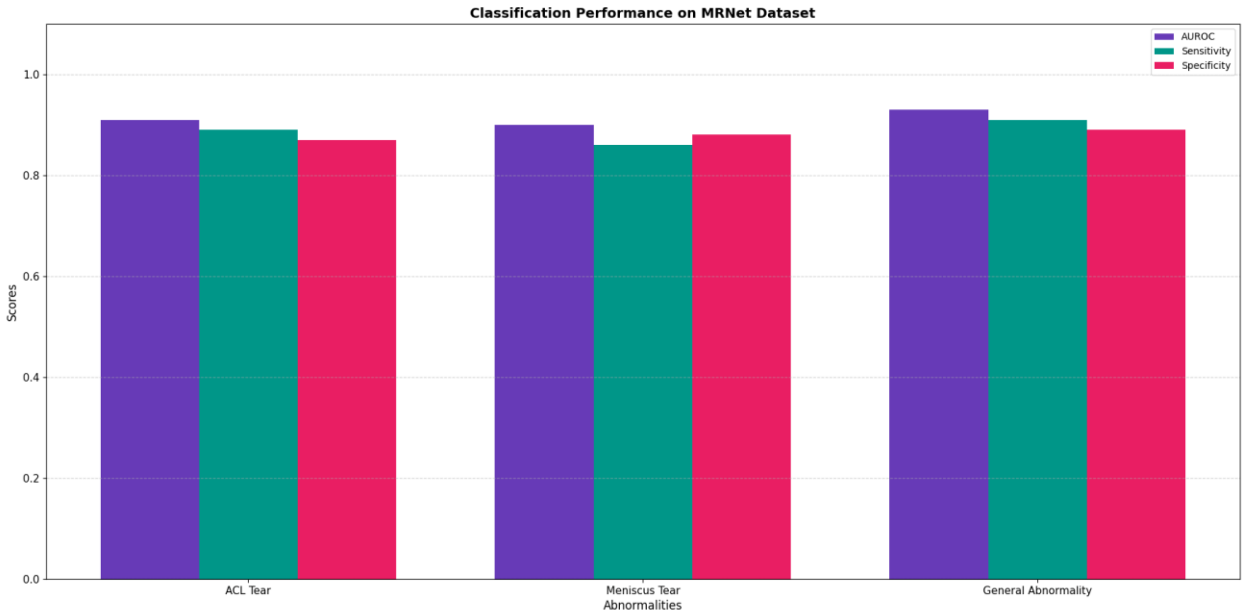

Fig. 2. Classification Graph

5. Discussions

The results show how successfully report generation, categorization, and segmentation function together in a single workflow. Our method leverages their synergy, unlike conventional methods that treat these tasks independently. Segmentation maps' differential lesion features enhance categorization. In order to ensure factual accuracy, the outcomes of both segmentation and classifications utilized to inform report generation.

Clinical adoption requires that the system avoid overstating predictions, which is ensured by the uncertainty estimate. Evaluations by radiologists also support the idea that fact-based, well-organized reports are easier to read and less taxing on the brain. equity among patients, stability across imaging scanners, and the identification of minor abnormalities. Future studies ought to examine federated learning for real-time deployment in hospital PACS systems and multi center training.

6. Conclusion

An AI powered MRI processing framework that can automatically produce diagnostic reports and make disease predictions is presented in this research. Our research opens the door for the use of AI in areas with limited access to skilled radiologists by showing that automated systems can generate outputs that are both interpretable and therapeutically meaningful. Furthermore, openness is guaranteed and confidence between clinicians and AI systems is increased by the reports' organized, evidence-based format. The acceptance of AI in healthcare depends on this kind of confidence.

There are still issues in spite of these achievements. Large-scale clinical trials should be used to validate the framework in future studies, and challenges with data diversity, domain generalization, and ethical considerations around the use of AI should be addressed. The smooth implementation of this system will depend on its integration with hospital PACS and electronic health record systems. Long-term, adding multi-modal imaging data to the system for example, merging MRI with CT and PET scans could improve the comprehensiveness of the diagnosis. To sum up, this work demonstrates that AI driven MRI analysis is both feasible and has potential in practical applications. Artificial intelligence systems like the one below are positioned to become useful tools for medical professionals in the future by aiding in healthcare accessibility, increasing reporting efficiency, and improving diagnostic precision.

References

[1] [1] Python Software Foundation. (2024). Python 3 Documentation.

[2] [2] TensorFlow. (2024). TensorFlow for Deep Learning and Medical Imaging.

[3] [3] Kermany, D. S., Goldbaum, M., Cai, W., Valentim, C. C., Liang, H., Baxter, S. L., ... & Zhang, K. (2018). Identifying medical diagnoses and treatable diseases by image-based deep learning. Cell, 172(5), 1122-1131.

[4] [4] Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., ... & Ng, A. Y. (2017). CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning. arXiv preprint arXiv:1711.05225.

[5] [5] Django Software Foundation. (2024). Django Documentation.

[6] PyTorch. (2024). PyTorch for Image Classification.

[7] OpenCV. (2024). OpenCV for Medical Image Processing.

[8] OpenAI. (2024). OpenAI GPT API Documentation.

[9] Microsoft Azure. (2024). Azure AI for Medical Imaging.

[10] Stack Overflow (2024). Talks about Django Integration and Machine Learning.

******