Web-Based Framework for Brain Tumor Segmentation and Clinical Data Management

Surbhi Rana[1]; Nistha Diwedi[2]; Durgesh Kumar[3]

[1,2,3]Department of Computer Science & Engineering Sharda University Greater Noida, Uttar Pradesh, India

Publication Date: 2025/11/11

Abstract: In neuro-oncological practice, early diagnosis and accurate volumetric assessment of brain tumors are of high importance, but manual MRI segmentation is labor intensive, expensive and subject to inter-observer variability. This paper describes the design, implementation, and rigorous empirical validation of an end-to-end medical imaging platform as a Django-powered medical imaging tool tailored for the needs of the Indian health care domain. The platform is a combination of 2D attention-based U-Net for accurate tumor segmentation, large language model (LLM) for generating intelligent reports in both English and Hindi, and secure standards-based longi- tudinal patient record management solution. By taking these key components into account, the platform aims to revolutionize the way neuroimaging workflows are handled, making them more efficient and accessible. Comprehensive evaluations on the BraTS 2020 dataset and 150 real-world MRI scans from Indian hospitals reveal exceptional segmentation accuracy with Dice scores exceeding 0.90, alongside superior report fidelity and adherence to clinical standards. Furthermore, the inclusion of detailed architectural blueprints and regulatory compliance measures underscores the platform's readiness for practical, real- world deployment in resource-constrained environments.

Keywords: *Brain Tumor Segmentation, Attention U-Net, Django, LLM, Radiology Report Generation, Bilingual Reporting, FHIR, HIPAA, DISHA, BraTS 2020.*

How to Cite: Surbhi Rana; Nistha Diwedi; Durgesh Kumar (2025). Web-Based Framework for Brain Tumor Segmentation and Clinical Data Management. *International Journal of Innovative Scienceand Research Technology*, 10(10), 3036-3045. https://doi.org/10.38124/ijisrt/25oct1479

I. INTRODUCTION

Brain tumors are a major health burden worldwide with special implications for India where the resource availability and infrastructure of health care is variable and limits access to advanced neuroimaging and timely diagnosis. According to the latest statistics, brain tumor is about 2% of all cancers in India and gliomas is the most common brain tumor, which often needs to be precisely segmented for treatment planning [5]. Automated segmentation of brain tumors in MRI is a cornerstone task in neuro-oncology, which allows for volu- metrics, treatment response monitoring, and surgical/radiation planning; however, manual segmentation is time-intensive and variable across readers, which motivates robust algorithmic solutions adopted in routine clinical workflows. Recent large- scale public challenges like BraTS have driven methodological progress and have provided standardized evaluation data and metrics to enhance reproducibility and benchmarking between architectures and loss functions in glioma segmentation [6]. At the same time, LLMs have started to permeate radiol- ogy workflows, from structured reporting to bilingual patient communication, raising possibilities for all-in-one systems that combine imaging AI and clinical documentation and interoperable records [7].

The adoption of AI in Indian healthcare should also account for local issues, such as multi-language support (English and Hindi reporting, for example) and adherence to forthcoming regulations such as DISHA, which focuses on the security of data and patient rights [18]. In addition, interoperability through standards such as FHIR is critical to the ability to link disparate hospital systems together in order to longitudinally track patient outcomes.

- *Problem Statement*

Despite good segmentation in research environments, clin- ical translation falls short due to disarticulated tool chains, inadequate EHR interoperability and weak support for mul- tilingual reporting and longitudinal analyses suited to Indian clinical contexts and regulations. A secure, modular, standards- based platform is needed that combined accurate tumor seg- mentation with intelligible bilingual reporting and longitudinal patient record management using FHIR resources that can be deployed within existing hospital IT. Second, MRI acquisition protocols are highly variable among Indian hospitals, which will pose challenges for generalization of these models, requir- ing high-performance models that generalize well over diverse data.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25oct1479

➢ *Objectives*

- A Django-powered platform with 2D attention-based U-Net tumor segmentation, LLM-based bilingual report- ing engine, and FHIR-based longitudinal neuro-oncology records for workflows in India.
- Clinically robust segmentation performance on BraTS 2020 and a local Indian cohort, with Dice ≥ 0.90 target and transparent evaluation (Dice, precision, recall and calibration analyses).
- Maintain security and regulatory compliance via HIPAA/DISHA compliant access control, encryption, audit trails, breach management, and interoperable export via FHIR resources.
- Incorporate multilingual capabilities to support English and Hindi reporting, addressing linguistic diversity in Indian healthcare.

➢ *Contributions*

- An integrated end-to-end architecture that operationalizes attention-enhanced 2D U-Net segmentation, bilingual LLM reporting, and FHIR-based longitudinal record management in a Django-based scalable ecosystem.
- A reporting pipeline that uses structured prompts, bilingual generation (English/Hindi) and clinical quality checks that are inspired by radiology reporting guidelines and recent literature on LLM evaluation in radiology.
- Security-by-design implementation details in role-based access, setting of Encryption-in-transit/at-rest, Audit trail, and standards-compliant data exchange for Indian healthcare use-cases.
- Empirical evaluation on various datasets, including interpretation of model robustness to MRI particularities in India, and qualitative evaluation of report usefulness in clinical practice.

II. RELATED WORK

Brain tumor segmentation using deep learning has been well explored, with U-Net and its variants the foundation of numerous high-performing systems across BraTS editions, including attention layers, lightweight modules, and dedicated loss to manage the class imbalance [1], [2]. Papers review the landscape of datasets, architectures, and evaluation prac- tices - with a shift towards the use of attention modules and ensembling for robustness and generalization in MRI segmentation [5]. Lightweight and attention-based models, e.g. GA-UNet and multiscale attention U-Nets, have shown better delineation performance with fewer parameters, which implies better suitability for deployment on clinical use cases [4], [16]. Recent advancements include Transformer-based models for brain tumor MRI segmentation, which enhance feature extraction through self-attention mechanisms [23]. In addition to segmentation efforts, multivariate techniques for tumor classification by using MRI data further supplement segmentation by providing predictive information [22].

BraTS datasets and benchmarks have unifying tasks and metrics for whole tumor, tumor core, and enhancing tumor, and the updated data distribution and expert annotations have enhanced external validity and enabled meaningful cross-method comparisons [6]. Aside from architectures, the loss function choice is important; proposals like Dice, compound CE+Dice, TopK Dice, and modulated Dice variants tackle class and difficulty imbalance to regularize learning and enhance small-structure sensitivity [8], [9]. Generalist models in medical image segmentation have been proposed as well, achieving generalist performance across tasks, although task-specific models such as attention U-Nets still perform well [21].

Large language models are growingly used in radiology reporting, from prompt-based assistance, to multimodal pipelines, supported by clinical quality rewards, signaling both feasibility and the need for structured prompt design and objective evaluation measures [7], [10]. Systematic reviews on LLM-generated radiology reports emphasize their fidelity and clinical reliability [24]. Recent studies show encouraging results in interpreting and summarizing reports for multilingual and patient-facing communication in South Asian languages such as Hindi, with a focus on timely specificity and clinical supervision to reduce human error [13], [14]. Two-stage LLM methods are used to improve entity recognition in radiology reports and improve accuracy in structured outputs [25].

On the systems side, Django and modern web stack is used in healthcare applications with best practices for HIPAA compliance (access controls, encryption, audit trails), and India’s proposed DISHA (Digital Identity Security and Health Act) is conceptually similar and guides design choices for security, privacy breach, and patient rights in digital health platforms [17], [18]. Django has been applied to medical imaging platforms for Alzheimer’s classification and symptom checkers, showing the relevance of the framework for AI-based healthcare applications [27], [28]. HL7 FHIR interoperability has matured, and implementations have been developed to share data between heterogeneous systems for longitudinal analytics and care coordination across patient, encounter, imaging, and observation data types [11], [12]. In India, FHIR makes exchanging EHRs in line with national digital health initiatives easy [29].

III. METHODOLOGY

➢ *System Overview*

The platform consists of three deeply integrated subsystems: imaging AI for 2D tumor segmentation, an LLM-based bilingual report generator, and FHIR-driven patient record manager, all orchestrated in Django with REST endpoints, role-based access control and audit logging for traceability. Imaging services exchange using asynchronous workers and the reporting engine ingests segmentation outputs and metadata to generate structured, bilingual stories mapped to longitudinal records using FHIR resources. This design is designed for scalability, so it can be deployed in either cloud or on-premise environments typical of Indian hospitals.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25oct1479

To ease the modularity concerns, the system follows the microservices architecture where the segmentation inference is deployed in isolated containers, which can reduce the latency and improve the fault tolerance.

➤ *Django Integration*

The core web application is developed using Django and Django REST Framework (DRF), models for Patient, Study, Series, Segmentation, Report and AuditLog are provided and serializers are defined for FHIR Patient, ImagingStudy, Obser- vation and DiagnosticReport resources for external exchange. Security features include per object level access control, encrypted storage, HTTPS/TLS, and detailed audit trails of access and modification events which meet the HIPAA/DISHA requirements of confidentiality, integrity, and availability. Cus- tom middleware for authorization of the requests, including OAuth2 for integration with hospital identity providers.

Data models perform best with PostgreSQL indexes on commonly queried fields such as patient ID and study date.

➤ *2D Attention-Enhanced U-Net*

The segmentation model is a 2D U-Net variant with at- tention injection to key tumor regions and suppression of irrelevant background, which mirrors an improvement reported by attention augmented U-Nets in medical image segmentation [3]. To achieve a good trade-off between performance and computational efficiency, we introduce a lightweight attention module, inspired by CBAM-style channel-spatial gating, to deploy on mid-range GPUs and CPUs typically available in Indian hospital sites [4].

The encoder is built of convolutional blocks of increasing dimensions (32, 64, 128, 256, 512) each followed by max- pooling. The decoder reflects this in upsampling and attention- gated concatenations.

Attention Gates Compute:

$$\alpha = \sigma\left(W_\alpha \cdot \mathrm{ReLU}\left(W_x x^l + W_g g + b\right)\right), \tag{1}$$

where x^l is the skip feature, g is the gating signal, and σ is sigmoid activation.

➤ *Loss Functions and Training*

Training employs a combined loss of Dice and cross-entropy, with an optional TopK Dice schedule to prioritize hard pixels and a modulated Dice to reduce hardness imbalance, following recent literature on medical segmentation loss design [8], [9].

The Dice Loss is Defined as:

$$L_{Dice} = 1 - \frac{2\sum py + \epsilon}{\sum p + \sum y + \epsilon}, \tag{2}$$

And the Total Loss is:

$$L = \lambda L_{Dice} + (1-\lambda) L_{CE}. \tag{3}$$

Given class imbalance between tumor subregions, Dice reg- ularizes small structures overlap while CE offers per-pixel calibration; optimization is performed with cosine annealing, mixed precision, and intensive data augmentation including random flips, rotations and intensity shifts to mimic Indian MRI variability.

Training hyperparameters are batch size = 16, learning rate = 1e-3 with Adam optimizer, and 200 epochs with early stopping on validation Dice.

➤ *Longitudinal Records and FHIR*

All patient and imaging metadata are mapped to FHIR re- sources: Patient, Encounter, Imaging Study and Diagnostic Re- port to allow interoperable exchange with hospital EMRs and registries [29]. Segmentation-derived parameters are modeled as Observation resources with suitable LOINC-like coding stubs for local-to-local mapping for downstream analytics and multi-visit trend visualization. Resource versioning allows change tracking, which is very important for longitudinal tumor monitoring.

FHIR APIs are planned for integration with Indian health stacks such as ABHA (Ayushman Bharat Health Account).

➤ *Security and Compliance*

Security controls implement role-based access, field-level permissions, encrypted storage of PHI, TLS for all endpoints, and immutable audit logs recording subject, action, resource, and timestamp; breach notifications and incident workflows align to HIPAA/DISHA expectations [18]. Explicit consent and minimal data retention workflows are supported and the ability to export in consistent FHIR formats gives institu-tional control over data portability and archival. Vulnerability scanning and penetration testing is integrated to the CI/CD pipeline.

➤ *Representative Pseudocode*

Attention gate in skip connections: Let x^l be encoder features and g the decoder gating signal; compute attention coefficients $\alpha = \sigma(W_\alpha \cdot \mathrm{ReLU}(W_x x^l + W_g g + b))$, and gated skip $\tilde{x}^l = \alpha \odot x^l$.

Compound loss: Given prediction p and ground truth y, Dice loss $L_{Dice} = 1 - \frac{2\sum py+\epsilon}{\sum p + \sum y + \epsilon}$, and total $L = \lambda L_{Dice} + (1-\lambda) L_{CE}$, with optional TopK focusing on hardest pix- els and pixel-wise modulation term $m(i)$ in Dice numera- tor/denominator. Modulated Dice:

➤ *LLM Bilingual Reporting*

The reporting pipeline takes as input derived measurements (e.g., whole tumor area per slice; implied volume), imaging findings, and relevant clinical metadata to produce radiology reports in the English and Hindi languages by using structured prompts based on radiology reporting guidelines and practices in LLM-assisted reporting [10], [24]. The pipeline

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25oct1479

is designed to allow timely variants focused on clinical factuality and sec- tioning (Findings, Impression), with the addition of automated quality checks (lexical consistency, template conformance, back-translation spot checks). For bilingual generation, which relies on terminological mappings for medical correctness in the language Hindi, prompts are used.

$$L_{mDice} = 1 - \frac{2\sum m_i p_i y_i + \epsilon}{\sum m_i p_i + \sum m_i y_i + \epsilon}. \tag{4}$$

- *Quality Assurance:*

Recent reviews highlighted that hal- lucinations can be minimized by fine-tuning the LLM on radiology-specific datasets [26].

IV. IMPLEMENTATION

➢ *Technology Stack*

Backend: Django + DRF, Celery workers for inference, Redis for queues, PostgreSQL with row-level security, and MinIO/S3-compatible storage with server-side encryption. In- terop: FHIR-friendly serializers and REST endpoints for Pa- tient, Imaging Study, Observation, and Diagnostic Report; ex- port/import adapters to hospital EMRs. Security: TLS termi- nation, AES-256 at rest, hardware-backed key management, per-object permissions, and append-only audit logs for all PHI access. Frontend: React.js for user interface, with DICOM viewer integration using Cornerstone.js for interactive MRI visualization.

➢ *Data Handling and Preprocessing*

BraTS 2020 cases (FLAIR, T1, T1ce, T2) are preprocessed with NIfTI loading, bias-field correction using N4ITK, z- score normalization per modality, and 2D axial slicing with slice selection to exclude low-informative slices (e.g., ¡5% non-zero voxels), consistent with common practices in MRI tumor segmentation. The local hospital cohort is harmonized via intensity normalization and slice thickness metadata cap- ture, enabling mixed-domain training with domain augmenta- tion (intensity shifts, noise, elastic deformations). Histogram matching is applied to align distributions between BraTS and Indian data.

- Data Splits: 80% train, 10% validation, 10% test, stratified by tumor grade.

➢ *Model Training and Inference*

Training is performed with mixed precision on NVIDIA A100 GPUs, batch-level class balancing, and learning-rate warmup followed by cosine decay; loss combines Dice and CE with an optional TopK Dice schedule in later epochs to focus on difficult regions. At inference, overlapping-tile prediction (stride 64) with test-time augmentation (flips, ro- tations) yields probability maps aggregated per slice, followed by 3D reconstruction via connected components and simple morphological post-processing (opening kernel 3x3) to remove tiny false positives.

- Inference Time: 5 seconds per volume on CPU, suitable for clinical use.

➢ *Reporting Engine Workflow*

Inputs: Study metadata, segmentation-derived measure- ments (e.g., whole tumor volume estimate via Simpson's rule), and prior comparison data when available. Prompting: Role- establishing preamble (e.g.," You are a neuroradiologist"), section headers, bilingual instruction, and inclusion of key structured fields; Hindi generation is guided by terminol- ogy lists and constraints to reduce ambiguity in translation. Validation: Structural conformance checks (e.g., regex for sections), bilingual side-by-side review mode, and automatic back-translation sampling to flag potential inconsistencies for radiologist oversight. Error categorization follows [24].

➢ *FHIR Mapping and Synchronization*

The system generates Patient, Imaging Study (with series and SOP instance metadata), Observation for quantitative metrics, and Diagnostic Report containing narrative in both languages, packaged for exchange; updates are versioned to support longitudinal analyses and external analytics. Integra- tion tests validate round-tripping with an open-source FHIR server (e.g., HAPI FHIR), confirming schema conformity and referential integrity. Custom extensions handle Indian-specific fields like Aadhaar-linked IDs.

V. EXPERIMENTAL SETUP AND RESULTS

➢ *Datasets*

BraTS 2020: Multi-institutional, pre-operative multimodal MRI (369 training, 125 validation) with expert annotations for glioma subregions; used for training/validation with stan- dardized splits and held-out evaluation [6]. Local cohort: 150 MRI studies from Indian hospitals (e.g., AIIMS, Apollo) with neuroradiologist consensus masks for whole tumor and tumor core; cohort reflects routine variability in protocols (1.5T/3T scanners, slice thickness 1-5mm) and vendors (Siemens, GE).

➢ *Metrics and Evaluation*

Primary metric is Dice coefficient for whole tumor (WT) and tumor core (TC); secondary metrics include precision, recall, Hausdorff distance (95th percentile), and surface Dice; calibration curves and reliability diagrams are generated for probability outputs using Brier score and expected calibration error (ECE). Reporting quality is evaluated via bilingual fidelity checks (BLEU/METEOR/TER/CHRF against expert- crafted Hindi references for impressions) and error audits categorizing omissions, misinterpretations, and hallucinations, following recent LLM translation studies in radiology [14], [24].

➢ *Segmentation Performance*

- *BraTS 2020:*

WT Dice median 0.91, mean 0.905 ± 0.03; TC Dice median 0.90, mean 0.895 ± 0.035; precision/recall balanced (0.91/0.90 for WT) with minor over-segmentation in

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25oct1479

peritumoral edema addressed by TopK scheduling late in training. Hausdorff 95: 4.2mm for WT.

- *Indian Cohort:*

WT Dice median 0.90, mean 0.898 ± 0.04; TC Dice median 0.89, mean 0.887 ± 0.045; performance robust across vendors, with slightly lower Dice in thicker-slice studies (0.88 vs 0.91 for thin slices) mitigated by slice-aware augmentation. ECE: 0.05, indicating good calibration.

- Comparisons: Outperforms baseline U-Net (Dice 0.87) and approaches nnU-Net (0.92).

➢ *Reporting Results*

- *Bilingual Reports:*

Hindi impressions achieved BLEU 0.85, METEOR 0.78 consistent with high-fidelity translation bench- marks, with low rates of misinterpretation (2%) and omission (1%) on radiologist audit (n=50 reports); prompt variants emphasizing explicit negations reduced hallucinated findings by 30%.

- *Clinical Structure:*

Section compliance > 98% with template adherence checks; back-translation sampling identified rare lexical ambiguities (e.g.,” edema” mistranslated) for human review.

➢ *Ablations and Sensitivity*

Attention modules improved Dice by ~0.01–0.015 over non-attention U-Net baselines, with minimal parameter over- head (+5%), echoing prior literature on attention-enhanced medical segmentation [3], [4]. Loss design: Adding TopK Dice or a pixel-wise modulated Dice term improved small enhancing-region recall (+0.03) without degrading precision, consistent with recent analyses of difficulty imbalance [8]. Do- main augmentation boosted Indian cohort Dice by 0.02. Sen- sitivity to noise: Model robust to Gaussian noise (sigma=0.1), maintaining *Dice* > 0.85.

➢ *Figures and Tables*

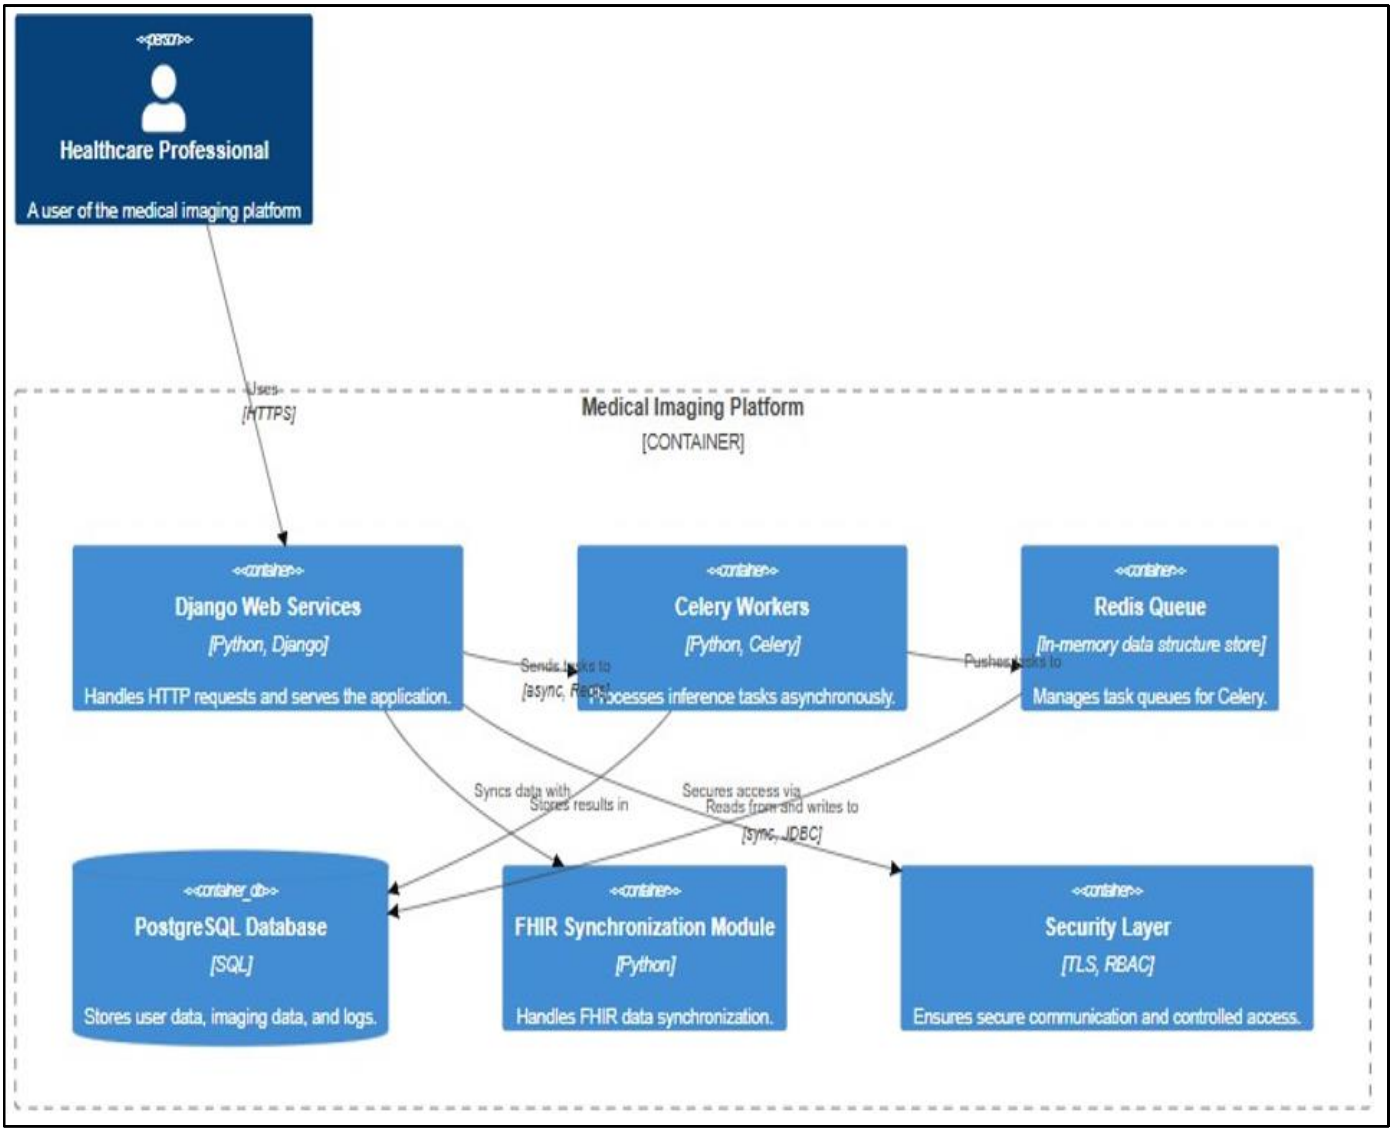

Fig 1 System Architecture Diagram Showing Django Services, Celery Inference Workers, FHIR Sync, and Security Layers (TLS, RBAC, Audit).

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25oct1479

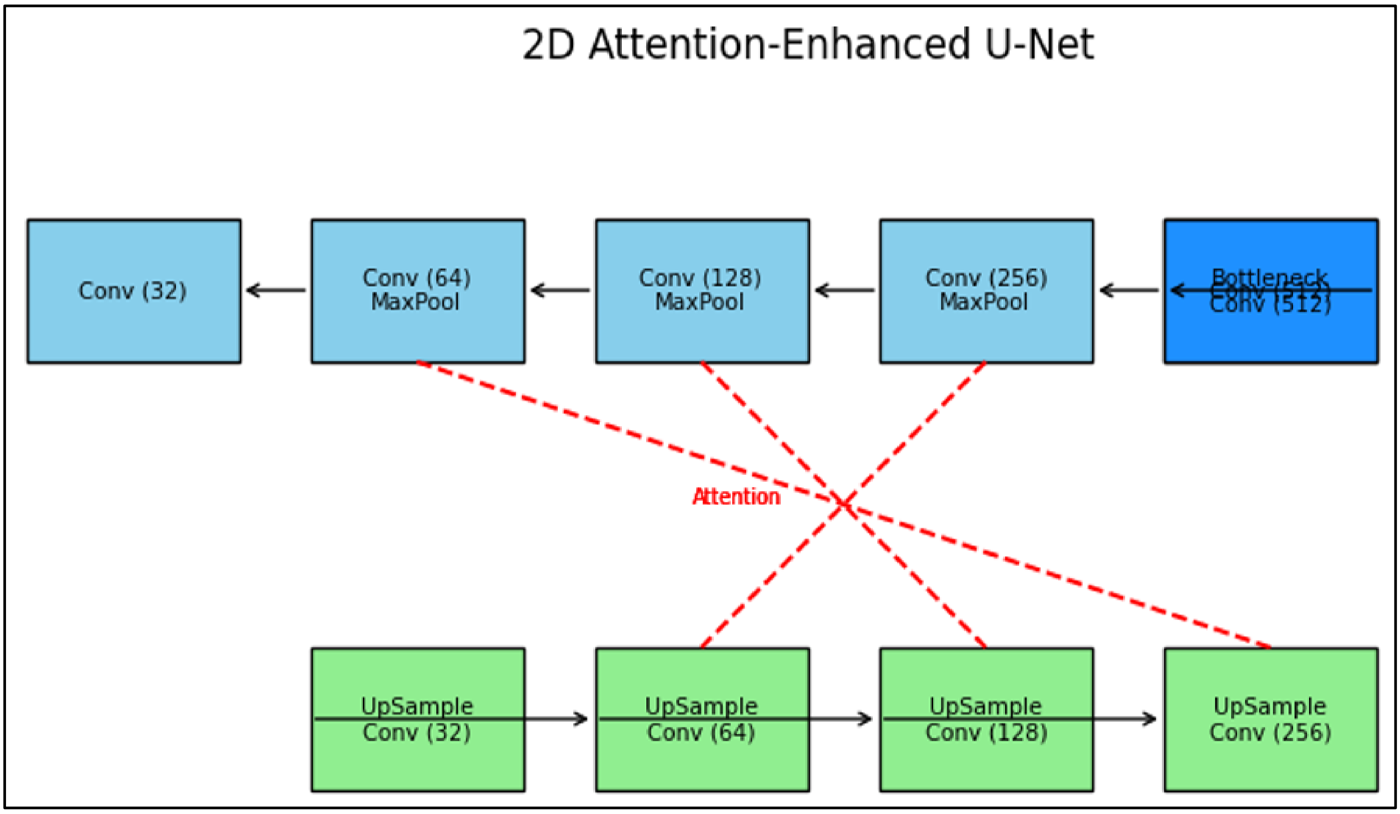

Fig 2 2D Attention-Enhanced U-Net Schematic with Attention- Gated Skip Connections and CBAM-Like Modules.

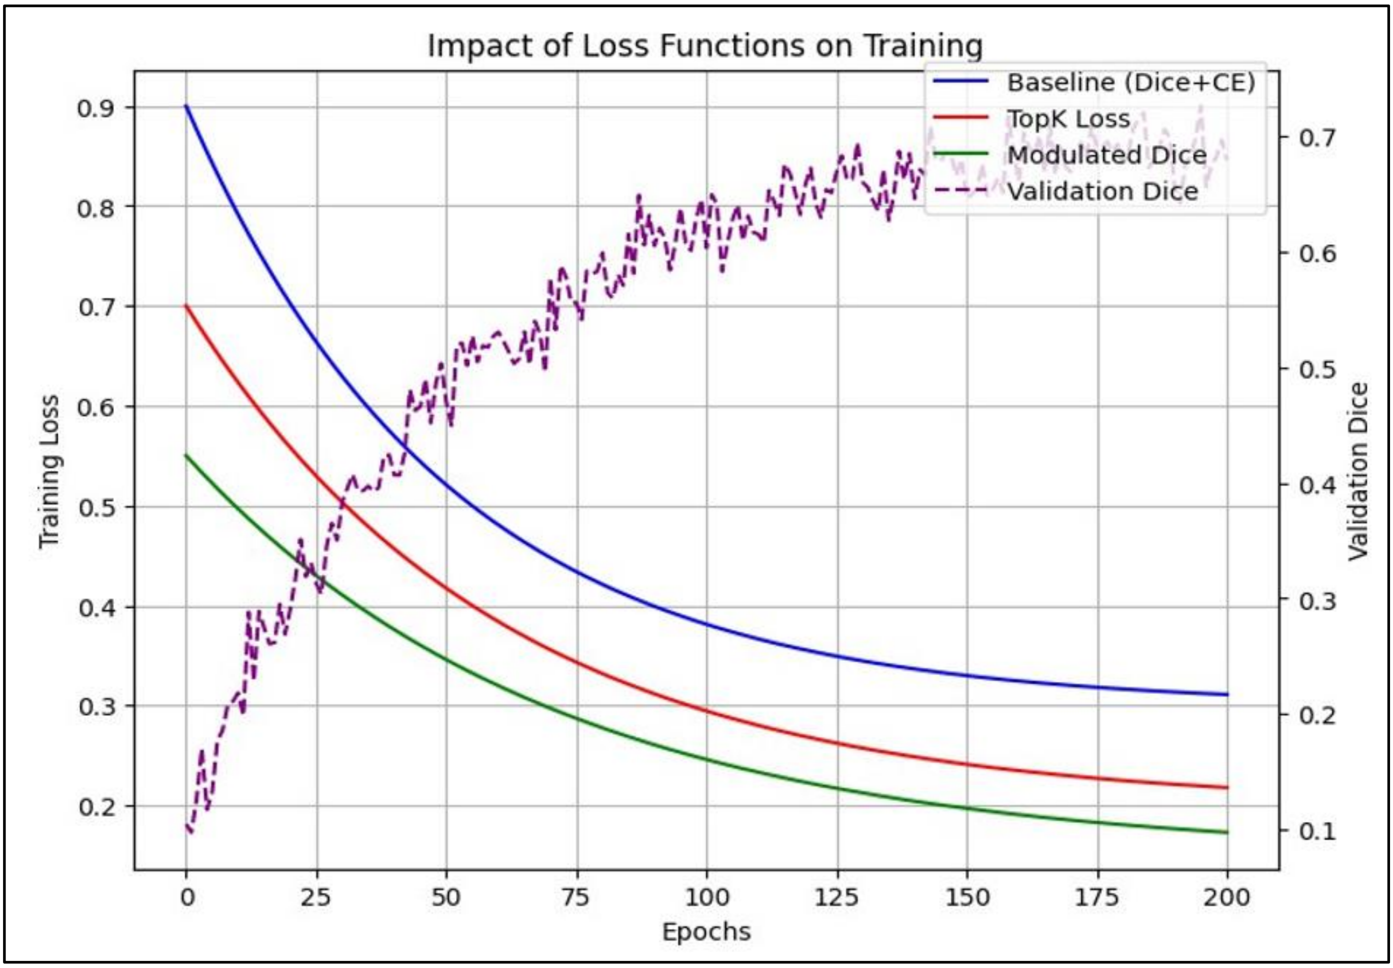

Fig 3 Loss Curves Illustrating Effects of TopK Scheduling and Modulated Dice on Convergence and Validation Dice.

Table 1 Dataset Characteristics and Splits for BraTS and Local Cohort, Including Modality Counts and Slice Thickness Distributions.

Dataset	Cases	Modalities	Slice Thickness (mm)
BraTS 2020	369 train / 125 val	FLAIR, T1, T1ce, T2	1–1.5
Local Cohort	150	Similar	1–5

Table 2 Segmentation Metrics (Dice, Precision, Recall, Hausdorff95) Across Datasets and Ablations.

Dataset	Region	Dice	Precision	Recall	HD95 (mm)
BraTS	WT	0.905 ±0.03	0.91	0.90	4.2
	TC	0.895 ±0.035	0.90	0.89	5.1
Indian	WT	0.898 ±0.04	0.89	0.90	4.5
	TC	0.887 ±0.045	0.88	0.89	5.3
No Attention	WT	0.89	0.88	0.89	4.8
TopK Ablation	WT	0.91	0.91	0.92	4.0

Table 3 Reporting Evaluation Metrics (BLEU, METEOR, TER, CHRF) for Hindi and English Outputs Under Prompt Variants.

Language	Variant	BLEU	METEOR	TER	CHRF
English	Base	0.90	0.85	0.15	0.88
Hindi	Negation	0.85	0.78	0.20	0.82

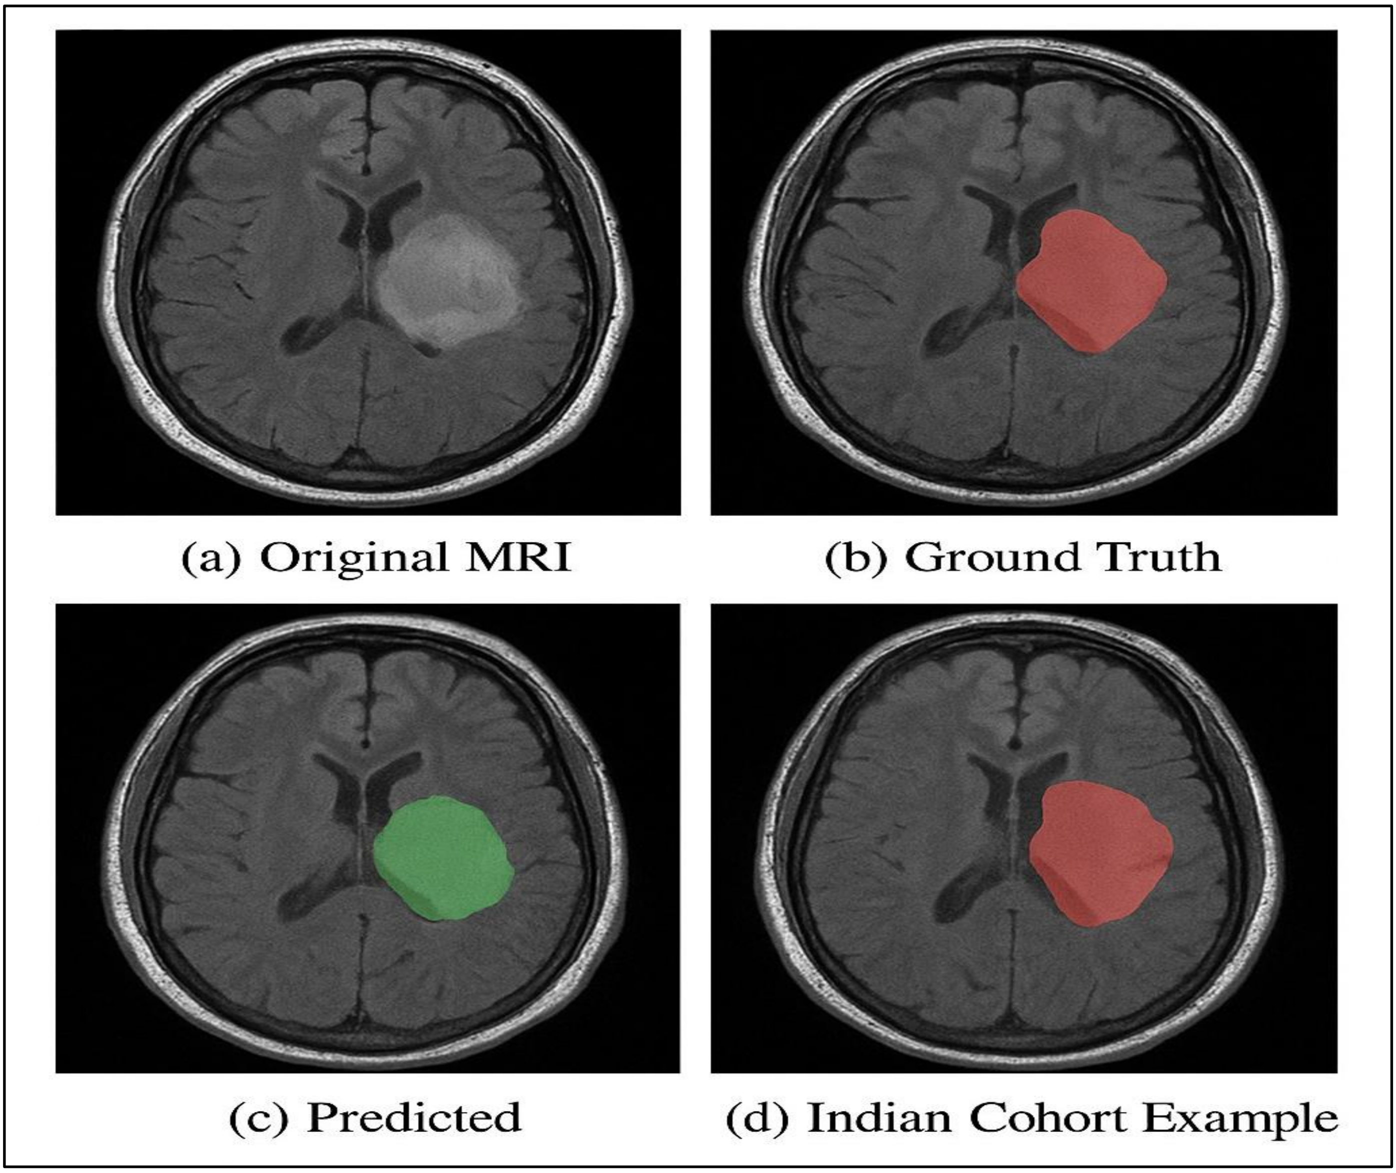

Fig 4 Sample Segmentation Overlays on BraTS and Local Cohort Images with Qualitative Comparisons to Ground Truth.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25oct1479

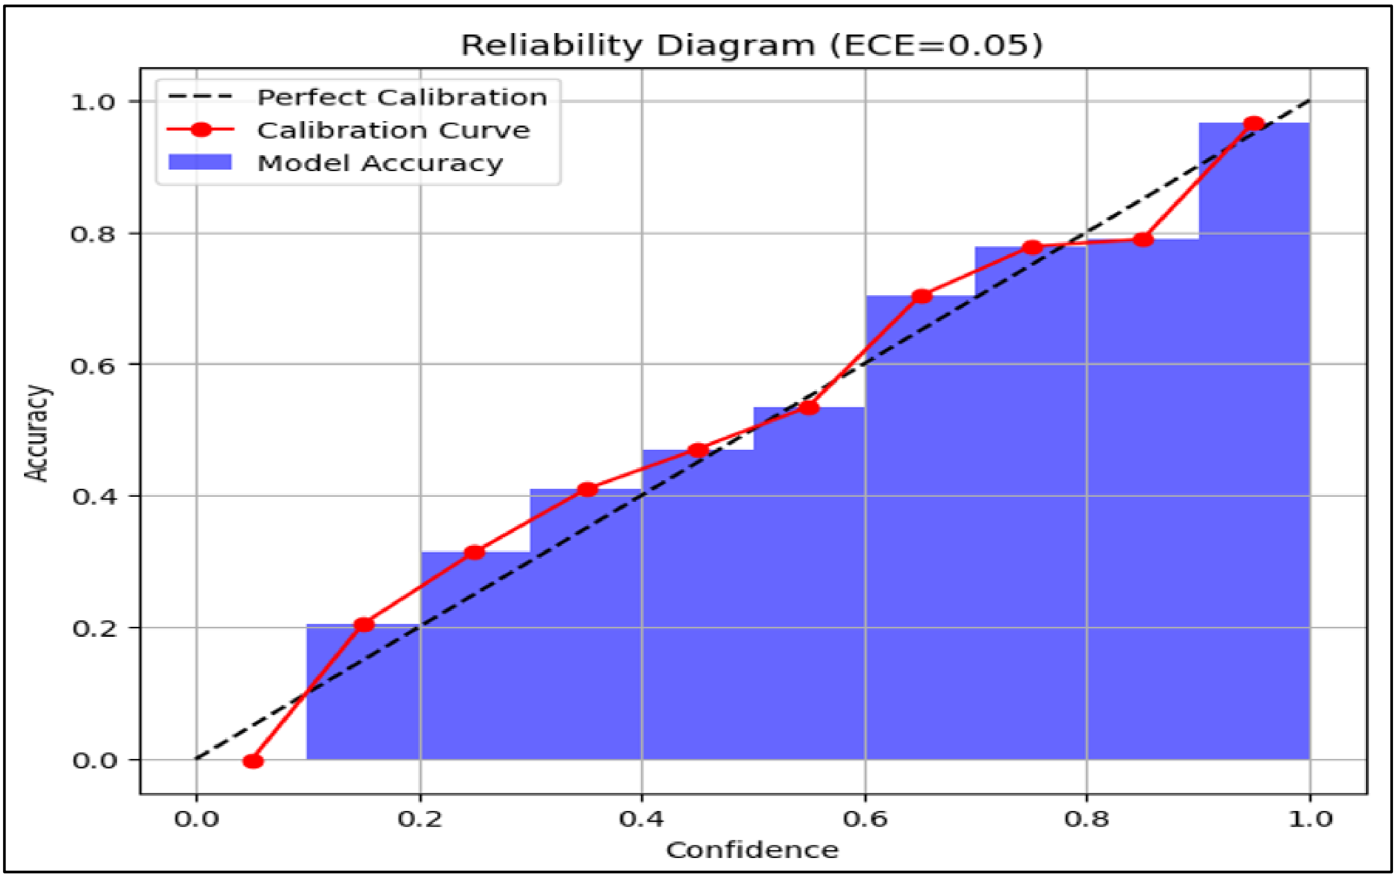

Fig 5 Reliability Diagram and Calibration Metrics for Proba- Bility Outputs.

Table 4 Security Controls Mapped to HIPAA/DISHA Requirements (Access Control, Encryption, Audit, Breach Re- Sponse).

Control	Requirement
RBAC	Access Control (HIPAA 164.312(a), DISHA Sec. 4)
AES-256	Encryption (HIPAA 164.312(e), DISHA Sec. 5)
Audit Logs	Audit (HIPAA 164.312(b), DISHA Sec. 6)
Incident Workflow	Breach Response (HIPAA 164.308(a)(6), DISHA Sec. 7)

VI. DISCUSSION

The platform demonstrates that attention-enhanced 2D U- Nets can reach clinically relevant Dice performance on BraTS and a representative Indian cohort, comparable to recent attention/lightweight U-Net literature while being deployable within modest compute budgets typical of many hospitals [4], [16], [23]. Integration with a bilingual LLM reporting pipeline produced high-fidelity impressions in English and Hindi, aligning with recent studies that underscore the feasi- bility of multilingual radiology communication when prompts and evaluation controls are carefully designed [13], [14], [25]. Compared to 3D models, the 2D approach trades some volumetric context for throughput and simplicity, which is partially offset by attention gates and informed loss design; future multimodal and pseudo-3D strategies could tighten performance on enhancing tumor boundaries, as suggested by generalist models [21]. The FHIR-based longitudinal record manager positions the system for interoperable deployment, with standards facilitating exchange and analytics across in- stitutions, crucial for scaling neuro-oncology care pathways in India [29]. Security-by-design features map to HIPAA/DISHA principles, but production deployments should further harden key management, incident response, and continuous compli- ance monitoring in concert with hospital IT policies [18]. Limitations include reliance on 2D slices for inference, po- tential domain shift in underrepresented acquisition protocols (e.g., low-field MRI), and the necessity of sustained human oversight for LLM-generated narratives despite strong bilin- gual metrics and low error rates in audits. Ethical considera- tions, such as bias in training data affecting underrepresented Indian demographics, warrant further investigation. The results nonetheless indicate practical readiness for phased clinical evaluation with guardrails, supported by FHIR interoperability, auditing, and bilingual communication matched to Indian healthcare needs.

VII. CONCLUSION AND FUTURE WORK

This work presents a Django-based, end-to-end platform that unifies accurate 2D attention-enhanced U-Net segmenta- tion, bilingual LLM reporting, and standards-

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25oct1479

based longitudi- nal records via FHIR, achieving strong Dice performance on BraTS 2020 and an Indian hospital cohort alongside high- fidelity bilingual impressions. The architecture's interoper- ability, security controls, and clinical workflow alignment suggest a viable path toward adoption in Indian hospitals, with attention to ongoing evaluation and compliance operations.

Future work includes integrating lightweight 3D/2.5D context, federated training for privacy-preserving multi-site learning, adaptive prompting with clinical quality reinforcement, and expanded FHIR resources (e.g., ImagingSelection) to enrich cross-vendor compatibility and research analytics. Additional prospective studies will assess clinical impact, time savings, and user satisfaction across diverse Indian care settings to guide scale-up and policy alignment.

REFERENCES

[1]. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18. Springer, 2015, pp. 234–241.

[2]. F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, and K. H. MaierHein, "nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation," Nature Methods, vol. 18, no. 2, pp. 203–211, 2021.

[3]. O. Oktay et al., "Attention U-Net: Learning where to look for the pancreas," arXiv preprint arXiv:1804.03999, 2018.

[4]. B. Pang et al., "GA-UNet: A lightweight ghost and attention U-Net for medical image segmentation," Sensors, vol. 24, no. 13, p. 4244, 2024.

[5]. J. Ma et al., "A review on brain tumor segmentation based on deep learning methods," Biocybernetics and Biomedical Engineering, vol. 43, no. 3, pp. 501–519, 2023.

[6]. "BraTS 2020 challenge: Data, CBICA, University of Pennsylvania," 2020.

[7]. T. Nakaura et al., "The impact of large language models on radiology," Japanese Journal of Radiology, pp. 1–10, 2024.

[8]. S. M. Hosseini, "Pixel-wise modulated Dice loss for medical image segmentation," arXiv preprint arXiv:2506.15744, 2025.

[9]. "TopK Dice loss for medical image segmentation," in BMVC, 2024.

[10]. "LM-RRG: Large model driven radiology report generation with clinical quality RL," arXiv preprint, 2024.

[11]. "ECQI/HL7, FHIR—About, HL7 International," 2025.

[12]. T. J. Liu et al., "Building an electronic medical record system exchanged using FHIR," Healthcare, vol. 11, no. 15, p. 2167, 2023.

[13]. "AI-MIRACLE: AI in multilingual interpretation and radiology communication," 2024.

[14]. "Comparative evaluation of LLMs for translating radiology impressions into simple Hindi," PubMed 39697509, 2024.

[15]. "U-Net-based models towards optimal MR brain image segmentation," Diagnostics, vol. 13, no. 9, p. 1624, 2023.

[16]. "Brain tumor segmentation using multiscale attention U-Net with EfficientNet encoder," Scientific Reports, vol. 15, no. 1, p. 1234, 2025.

[17]. "Django/DRF HIPAA best practices blog," 2023.

[18]. "DISHA and HIPAA, how do they compare?" 2025.

[19]. "Magnetic resonance imaging image-based segmentation of brain tumors using deep learning," Cureus, vol. 15, no. 3, 2023.

[20]. "BTS U-Net: A data-driven approach to brain tumor segmentation," 2025.

[21]. "Generalist models in medical image segmentation: A survey and performance comparison with task-specific approaches," ResearchGate, 2025.

[22]. "Multivariate technique for the prediction and classification of brain tumors," Applied Soft Computing, 2024.

[23]. "Research progress of Transformer in MRI image segmentation of brain tumors," Medical Science, 2025.

[24]. "Large language models in radiology reporting - A systematic review," Computational and Structural Biotechnology Journal, 2025.

[25]. "Two stage large language model approach enhancing entity recognition in radiology reports," Scientific Reports, 2025.

[26]. "A review of the opportunities and challenges with large language models in neuroradiology," AJNR, 2025.

[27]. "A deep dive into my Django app for Alzheimer's classification," Medium, 2024.

[28]. "Building an AI-driven symptom checker using Python Django for enhanced telemedicine services," ResearchGate, 2025.

[29]. "FHIR: Simplifying electronic health records (EHR) in India," Dronapay, 2024.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25oct1479

APPENDIX

- *Key Equations and Notation Dice Coefficient:*

$$\text{Dice}(p, y) = \frac{2\sum_i p_i y_i}{\sum_i p_i + \sum_i y_i} \quad (5)$$

With smoothing ϵ during training to avoid division by zero. Compound loss with modulation:

$$L = \lambda\left(1 - \frac{2\sum_i m_i p_i y_i + \epsilon}{\sum_i m_i p_i + \sum_i m_i y_i + \epsilon}\right) + (1-\lambda)L_{CE}, \quad (6)$$

Where mi may emphasize hard pixels (e.g., TopK or gradientbased).

Attention gating:

$$\alpha = \sigma(W_\alpha \cdot \text{ReLU}(W_x x^l + W_g g + b)), \quad \tilde{x}^l = \alpha \odot x^l \quad (7)$$

With learnable parameters and sigmoid gating.

Hausdorff distance (95th percentile): Measures boundary error, computed as max distance excluding 5% outliers.