iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** VI **Month of publication:** June 2026

DOI: https://doi.org/10.22214/ijraset.2026.83658

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 14 Issue VI June 2026- Available at www.ijraset.com

NeuroGenAI: A Full-Stack Deep Learning-Based Brain Tumor Detection System with AI-Assisted Clinical Report Generation

Tushar Nangare[1], Kalpesh Patil[2], Sarvesh Namra[3], Neha Patil[4], Prof. Digvijay Patil[5]
Progressive Education Society Modern College Of Engineering Pune

Abstract: ***Every minute separating a suspicious MRI from a physician-readable report is a minute of clinical uncertainty—and in neuro-oncology, uncertainty compounds rapidly into harm. NeuroGenAI confronts this reality by fusing a VGG16 convolutional classifier with a large language model reporting engine to produce, within a single browser interaction, both a tumour classification and a structured clinical narrative ready for physician review. The classifier operates across four neuro-oncological categories (Glioma, Meningioma, Pituitary Tumour, No Tumour) on gadolinium-free MRI input; the Google Gemini language model then converts the raw prediction into section-organised prose (Findings, Impression, Recommendations) conditioned on patient metadata. Surrounding these inference components is a production-grade web stack: React.js for the clinician-facing interface, Django REST Framework for backend orchestration, PostgreSQL for relational patient records, and Cloudinary for cloud-resident media. On a 1,400-image holdout partition of the Br35H dataset, the system posts a weighted-average precision of 0.96, a Pituitary-class recall of 0.99, and an end-to-end scan-to-report latency of roughly 2.8 seconds—figures that together make a practical, rather than aspirational, case for AI-assisted neuro-oncology at the point of care.***
Keywords: Brain Tumour Detection, VGG16 CNN, Google Gemini LLM, Clinical Report Generation, Full-Stack Healthcare AI, Diagnostic Decision Support, Br35H Dataset.

I. INTRODUCTION

Intracranial tumours do not announce themselves conveniently. They manifest through symptoms—headaches, seizures, cognitive drift—that overlap substantially with benign neurological conditions, making imaging the decisive arbitration step. What happens after the MRI machine finishes its acquisition, however, depends entirely on a human specialist interpreting hundreds of two-dimensional slices, reconciling subtle contrast differences that may distinguish a meningioma from reactive tissue, or a pituitary adenoma from a glioblastoma margin. That interpretive act is the bottleneck (solanki2023?).

Gadolinium-contrast MRI has become the preferred pre-surgical mapping tool in neuro-oncology for reasons that go beyond image quality: it avoids the ionising dose of CT and achieves soft-tissue contrast sufficient to delineate tumour boundaries at sub-centimetre resolution (rao2024?). Yet the analytical burden it places on radiologists is correspondingly high. Each scan demands sustained attention across a volumetric stack, with performance sensitive to reader fatigue, seniority, and institutional scanning protocol variation (yao2024?). In geographically underserved settings, subspecialty neuro-radiology may simply be unavailable, rendering remote or asynchronous interpretation the only realistic pathway.

Convolutional neural networks have demonstrated that the feature hierarchies required for reliable tumour classification can be learned from raw pixel representations—without the manually engineered descriptors that limited earlier computer-aided diagnosis systems. Deep architectures such as VGG16, ResNet, and Inception achieve inter-expert-level discriminative performance on multi-class MRI classification by progressively composing local edge responses into increasingly tumour-specific signatures (simonyan 2014?; rao2024?). However, high classification accuracy does not automatically translate into clinical utility.

The critical missing link is *narration*: the transformation of a probability vector into a documented clinical impression that radiologists can countersign and patients can understand. Deployed deep-learning tools today return numerical class scores without context, leaving the physician to bridge from *Glioma: 0.962* to a structured finding that can anchor a multidisciplinary tumour board discussion (lata2024?; gemini2023?). Separately, most research prototypes exist as isolated inference scripts rather than integrated clinical applications with role-managed access, secure media handling, and audit-ready reporting (**solanki2023?**).

NeuroGenAI is built to eliminate both deficiencies under a single deployment. Its contribution is threefold:

1) Diagnostic Precision: A from-scratch VGG16 CNN, trained on the Br35H benchmark dataset, classifies MRI scans across four categories with a weighted-average accuracy of 96.5%, producing a calibrated confidence score alongside each label.

2) Automated Documentation: A Google Gemini LLM reporting module converts classification outputs and patient metadata into a fully structured clinical report—Findings, Medical Impression, and Recommended Next Steps—without physician intervention.

3) Clinical Deployment Readiness: A complete React–Django–PostgreSQL–Cloudinary stack enforces role-based access, JWT-secured API communication, and cloud-native media management, making the system operable at real institutional data volumes.

Sections II through VIII develop each of these dimensions in turn, from literature context through architecture, implementation, experimental results, and future outlook.

II. LITERATURE SURVEY

A. The Accuracy Trajectory of CNN-Based Neuro-Oncological Imaging

Early computational approaches to tumour detection in MRI operated through a two-stage pipeline: first, a hand-crafted feature extractor (wavelet sub-bands, grey-level co-occurrence statistics, Gabor energy maps) condensed volumetric complexity into a feature vector; second, a shallow discriminative model—most often a Support Vector Machine—mapped that vector to a class label. Performance was acceptable within tightly controlled datasets but collapsed when acquisition parameters changed between institutions, because hand-tuned descriptors encode assumptions about scanner properties rather than tumour biology.

Solanki et al. (**solanki2023?**) conducted a comprehensive audit of this transition, quantifying the accuracy gains that followed the adoption of end-to-end learnable architectures. Their analysis confirms that networks capable of constructing their own feature vocabulary from labelled images consistently outperform descriptor-pipeline baselines, particularly on heterogeneous multi-class datasets where class boundaries are complex. Rao et al. (**rao2024?**) extended this finding to transfer learning, demonstrating that VGG16 and ResNet50 feature representations—pre-conditioned on large natural-image corpora—generalise usefully to gadolinium MRI even without domain-specific pretraining, yielding classification accuracy above 95%. Rehman et al. (**rehman2020?**) corroborated these figures using VGG19 under analogous transfer conditions, while also surfacing the computational cost that comes with deep networks carrying hundreds of millions of parameters—a practical concern for deployment on commodity hardware.

Yao et al. (**yao2024?**) shifted the emphasis from classification to spatial localisation, embedding attention mechanisms into a YOLOv8 detection head to improve sensitivity for small intraparenchymal lesions. While the resulting model excels at bounding lesion extent, its single-stage detection objective does not naturally resolve the multi-class tumour typing question that drives clinical treatment selection.

B. Security as a Non-Negotiable Constraint

The moment AI inference touches patient-identifiable imaging data, legal and ethical obligations that are peripheral in research become central in deployment. Lata et al. (**lata2024?**) addressed this directly, constructing a tumour detection framework within a privacy-preserving smart healthcare architecture. Their key empirical finding—that cryptographic data protection does not materially degrade diagnostic accuracy—removes a frequent objection to clinical AI deployment. That said, their work stops at the classification boundary; the question of how machine findings should be communicated to clinicians in a format that supports documentation and liability remains unaddressed.

C. The Interpretability Deficit: Where Current Research Stalls

Across all reviewed literature, a structural gap recurs: every system ultimately delivers a numerical output—a class label, a confidence score, a segmentation mask—without the explanatory layer that clinical practice requires. Radiologists do not sign probability vectors; they sign reports. The cognitive step from “Meningioma: 0.985” to a written clinical impression involves pathophysiological reasoning, patient-context weighting, and differential enumeration—none of which current classification models perform. Large language models offer a compelling mechanism to close this gap, yet the literature contains no deployed platform that integrates visual diagnostic inference with automatic natural-language report synthesis. NeuroGenAI occupies this gap directly.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

III. PROBLEM STATEMENT

Two distinct gaps afflict current AI-assisted neuro-oncology, and conflating them produces misleading solutions.

The first—the *accuracy gap*—has attracted extensive research attention and is largely closed. CNN classifiers now match or exceed specialist agreement on benchmark MRI datasets, and the active question is no longer whether machines can identify tumours but under what conditions their predictions remain trustworthy.

The second—the *workflow gap*—receives far less attention despite its arguably greater barrier to adoption. Even a perfectly accurate classifier is clinically inert if its output cannot be incorporated into the documentation artefacts that drive patient management. A physician confronted with a raw prediction cannot hand it to a tumour board, enter it into an electronic health record, or use it to justify an urgent surgical referral. Structured reporting is not an optional nicety; it is the mechanism through which diagnostic conclusions become clinical actions.

Additionally, almost no published systems provide the surrounding infrastructure—authenticated user roles, secure image storage, real-time dashboards, downloadable PDF output—that would allow a hospital to actually deploy an AI classifier in a live clinical environment. The solution space that NeuroGenAI occupies is therefore not simply a more accurate classifier; it is a complete, workflow-integrated diagnostic platform that simultaneously closes both gaps.

IV. PROPOSED SYSTEM

NeuroGenAI is realised as a browser-resident clinical platform that serves two operationally distinct user roles within a unified application boundary. Its internal architecture separates the interaction surface from the diagnostic engine, allowing each to evolve independently.

A. Clinician-Facing Interface Module

The frontend operates as the sole point of human contact with the system and enforces role-specific access discipline from the moment of authentication.

1) Technician Workflow: A scan submission portal accepts high-resolution MRI files in PNG or JPEG format. Before any server-side transmission occurs, the interface validates file format and collects mandatory patient demographic fields (age, gender). This client-side gating serves two purposes: it prevents malformed payloads from consuming backend inference resources, and it ensures demographic metadata is available for downstream report contextualisation.

2) Physician Workflow: An analytics dashboard surfaces the full scan history associated with each patient under the clinician's care. AI-generated classification results appear immediately upon backend processing completion, eliminating polling delays. Physicians may review, validate, and download structured PDF reports directly from the dashboard interface—the complete scan-to-report cycle accessible through a single, authenticated session.

B. Backend Diagnostic Engine

Three purpose-built processing services constitute the backend's diagnostic capability:

1) Inference Service: Receives a raw MRI image, applies a preprocessing pipeline (spatial resampling to 224×224 pixels, intensity normalisation to $[0,1]$), and forwards the resulting tensor to the serialised VGG16 model checkpoint. Returns a four-class probability distribution and an associated confidence calibration score.

2) Narrative Generation Service: Accepts the inference output and patient metadata, constructs a structured prompt for the Google Gemini API, and parses the returned prose into three named clinical sections. Assembles the parsed content into a standardised PDF via an HTML rendering pipeline.

3) Data Governance Service: Persists relational patient and scan records in PostgreSQL, offloads binary image and report assets to Cloudinary cloud storage, and maintains referential integrity between scan records and their associated clinical documentation.

V. SYSTEM ARCHITECTURE

NeuroGenAI adopts a Service-Oriented Architecture (SOA) across three logically separated tiers: presentation, application-and-intelligence, and persistence. This decomposition is not aesthetic; it reflects a deliberate engineering decision to allow the VGG16 inference core and the LLM reporting engine to be upgraded or replaced independently, without disrupting the user-facing interface or the data layer.

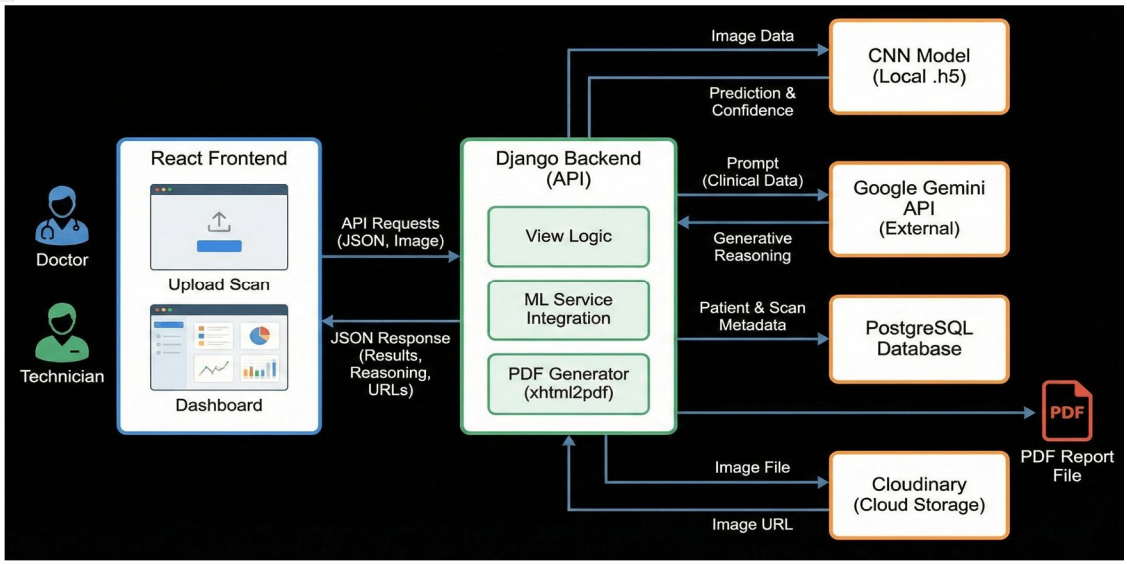

NeuroGenAI System Architecture Overview: Three-tier decomposition separating React.js presentation layer, Django REST application & intelligence layer, and PostgreSQL/Cloudinary persistence layer.

A. Presentation Tier

The frontend is implemented with React.js v18.0, bundled via Vite for optimised production asset delivery. The Single Page Application model eliminates full-page navigation reloads, maintaining application state across multi-step workflows without server round-trips that would interrupt a clinician mid-review.

Three architectural decisions within the frontend warrant specific explanation:

- State Topology: Local component state (upload progress, prediction display) is managed through React Hooks (useState, useEffect). Identity state that must persist across component boundaries—active patient session, access token payload—is propagated via the Context API, avoiding brittle prop-drilling chains.
- Authenticated API Layer: All backend communication routes through an Axios HTTP client equipped with request interceptors that attach JSON Web Tokens to outgoing Authorization: Bearer headers. This mechanism decouples authentication logic from individual components, ensuring token management is uniform across every API call.
- Role-Differentiated Routing: Route guards evaluate the authenticated user's role at navigation time, rendering either the Technician submission portal or the Physician analytics dashboard. The two surfaces share no overlapping functionality, preventing privilege escalation through URL manipulation.

B. Application and Intelligence Tier

The backend Django REST Framework service coordinates four distinct operational concerns:

- Endpoint Security Perimeter: The /api/predict/ and /api/auth/ endpoints accept multipart/form-data for image payloads and JSON for metadata. SimpleJWT issues and rotates short-lived access tokens (authenticated scan requests) and long-lived refresh tokens (session persistence), while CORS configuration restricts accepted origins to registered frontend deployments.
- Inference Orchestration: On receipt of an MRI file, the backend executes a synchronous preprocessing pipeline—Pillow for spatial resampling, NumPy for vectorised normalisation— producing an input tensor compatible with the VGG16 model deserialised from its .h5 checkpoint. The forward pass yields a four-element softmax distribution; the argmax index is resolved to a clinical class label.
- LLM Prompt Construction and Parsing: A templated prompt injects the predicted class, confidence percentage, and patient demographic fields into a persona-grounded instruction (the model is cast as a board-certified neuroradiologist). The Gemini API response is parsed by section header into a structured JSON object keyed by Findings, Impression, and Plan.
- Document Assembly: Parsed LLM content is merged with patient identifiers and the Cloudinary-hosted MRI image URL into a Jinja2 HTML template. The xhtml2pdf library renders the template to a portable PDF, which is uploaded to Cloudinary and its access URL written to the patient's database record.

C. Data Persistence Tier

The persistence layer employs a hybrid storage architecture. PostgreSQL manages all structured, queryable data under a normalised schema: User entities reference Patient profiles, which in turn reference ordered collections of ScanResult records. Foreign-key constraints ensure that no diagnostic report can exist without an associated patient identity—a prerequisite for audit trail integrity.

Binary assets are deliberately excluded from the relational store. MRI images (typically 5–20 MB per scan) and generated PDF reports are streamed directly to Cloudinary at upload time; the database receives only the resulting CDN-backed URL. This separation keeps query times independent of dataset scale and leverages Cloudinary's global delivery network for fast, geographically distributed asset access.

VI. METHODOLOGY AND IMPLEMENTATION

Development was structured in three sequential phases: dataset preparation and augmentation pipeline design; VGG16 architecture configuration and supervised training; and clinical reporting module construction and integration. Each phase produced artefacts that served as direct inputs to the next.

A. Dataset Preparation and Augmentation Strategy

Training data was sourced from the Br35H Brain Tumour Dataset (**hamada2020?**), a benchmark corpus of 7,023 annotated MRI images spanning four diagnostic categories: Glioma, Meningioma, Pituitary Tumour, and No Tumour. The multi-class composition is significant: it compels the model to resolve subtle inter-class boundaries—for instance, distinguishing a pituitary microadenoma from a meningioma at the skull base—rather than simply separating tumour from normal tissue.

Raw images in the dataset exhibit heterogeneous spatial resolutions and acquisition-dependent intensity distributions. Three preprocessing operations standardise these properties before training:

- Spatial Resampling (224×224 pixels): VGG16's input contract specifies a fixed spatial dimension; bicubic resampling is applied to honour this constraint while minimising interpolation artefact introduction.
- Linear Intensity Normalisation ($[0,1]$): Dividing raw pixel values by 255 confines inputs to the unit interval, stabilising gradient magnitudes across the first convolutional layer and improving convergence rate.
- Online Augmentation: To prevent the model from memorising scanner-specific acquisition artefacts, stochastic transformations are applied at mini-batch generation time: rotation within $\pm 15°$, horizontal axis reflection, and zoom perturbation. These operations expand the effective training distribution without requiring additional acquisitions.

B. VGG16: Architecture Design Rationale

VGG16 (simonyan2014?) comprises sixteen weight-bearing layers arranged in five convolutional blocks followed by three fully connected layers. Its decisive architectural feature—uniform 3×3 kernel deployment across all convolutional positions— reflects a specific inductive bias: two stacked 3×3 filters span the same receptive field as a single 5×5 filter while consuming fewer parameters and interposing an additional non-linearity. Three stacked 3×3 filters approach the receptive field of a 7×7 filter with comparable parameter savings and two additional ReLU activations. This compounding effect gives the architecture expressive power disproportionate to its filter count.

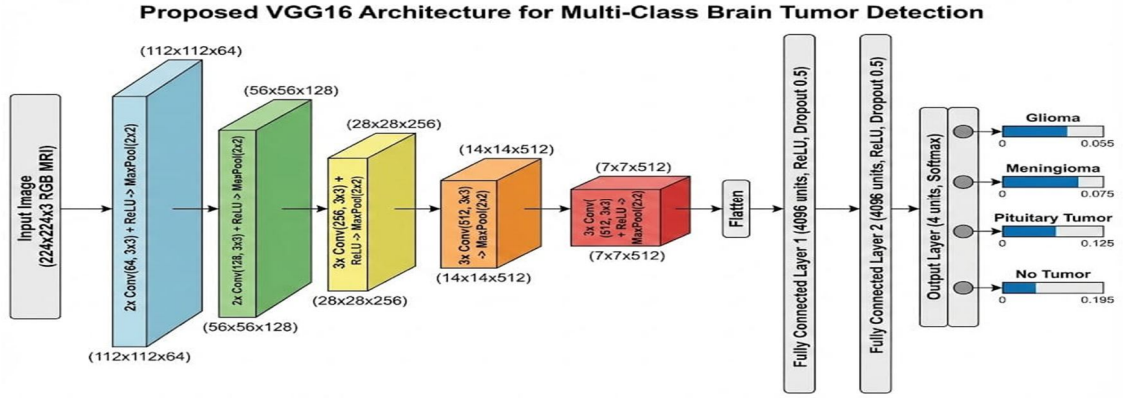

VGG16 Convolutional Neural Network Architecture: Five progressive convolutional blocks with uniform 3×3 kernels, batch-wise max-pooling, and three fully connected layers terminating in 4-way softmax classification head.

1) Structural Decomposition

- Input: 224 × 224 × 3 tensor. Single-channel MRI input is replicated across three channels to satisfy the architectural contract without distorting spatial information.
- Block 1: 2 × Conv(64 filters, 3 × 3, ReLU, same) → MaxPool(2 × 2). Spatial output: 112 × 112 × 64.
- Block 2: 2 × Conv(128 filters, 3 × 3) → MaxPool(2 × 2). Output: 56 × 56 × 128.
- Block 3: 3 × Conv(256 filters, 3 × 3) → MaxPool(2 × 2). Output: 28 × 28 × 256.
- Block 4: 3 × Conv(512 filters, 3 × 3) → MaxPool(2 × 2). Output: 14 × 14 × 512.
- Block 5: 3 × Conv(512 filters, 3 × 3) → MaxPool(2 × 2). Output: 7 × 7 × 512.
- Flatten: The 7 × 7 × 512 activation volume is serialised to a 25,088-dimensional vector.
- FC Layers: Dense(4096, ReLU) → Dropout(0.5) → Dense(4096, ReLU) → Dropout(0.5) → Dense(4, Softmax).

2) Parameter Distribution

The network carries approximately 138 million trainable parameters. Roughly 123 million of these concentrate in the three fully connected layers; the weight tensor connecting the 25,088-element flattened vector to the first 4,096-unit Dense layer alone accounts for over 102 million weights. The convolutional stack contributes the remaining 14–15 million. This distribution motivates the application of 0.5-rate Dropout exclusively to the fully connected layers, where overfitting pressure is greatest.

3) From-Scratch Training Decision

The model was initialised with random weights (weights=None in Keras), deliberately bypassing ImageNet pretraining. The rationale: ImageNet-conditioned features encode colour-texture statistics and object-part patterns characteristic of natural photographic images—properties that are largely irrelevant or potentially counterproductive for single-channel, contrast-normalised MRI volumes. Training from scratch allows the convolutional filters to optimise directly for the intensity gradients, boundary sharpness, and internal structure patterns that distinguish neuro-oncological tissue classes.

4) Optimisation Configuration

Stochastic Gradient Descent with momentum = 0.9 and initial learning rate = 0.01 drove weight updates. Batch size = 32 balanced gradient estimate variance against the GPU memory envelope of the deployment hardware. Categorical Cross-Entropy served as the loss criterion, appropriate for the mutually exclusive four-class softmax output. ModelCheckpoint preserved the highest-validation-accuracy weight state; EarlyStopping terminated training on validation loss plateau, recovering from both underfitting and late-epoch overfit.

C. Clinical Reporting Module

The reporting pipeline operates on three inputs: the predicted tumour class, its associated confidence scalar, and the patient demographic record captured at submission. From these inputs, it produces a physician-ready PDF through three ordered operations.

1) Persona-Grounded Prompt Engineering: Rather than issuing a generic summarisation request, the backend constructs a role-anchored prompt that positions Gemini as a board-certified neuroradiologist. The prompt specifies the expected output structure (Findings, Medical Impression, Recommended Next Steps), the required professional register, and the constraint that all assertions be grounded in the provided imaging classification and patient data.

2) Structured Response Extraction: Gemini's prose response is parsed via regular expressions keyed to the specified section headers. Matched content is serialised into a JSON object, decoupling the LLM interaction from the document rendering step and enabling independent validation of each section's content.

3) PDF Compilation Pipeline: The structured JSON is merged with hospital header metadata and the Cloudinary-hosted MRI image reference into a Jinja2 HTML template. The xhtml2pdf renderer converts this template to a portable document. The PDF is uploaded to Cloudinary; its access URL is persisted against the patient scan record.

The complete computational workflows for classification and report generation are formalised in Algorithms [alg:classify] and [alg:report].

Raw MRI image I, serialised VGG16 model M Clinical label L, calibrated confidence C Receive I from authenticated API endpoint $I_r \leftarrow \textsc{Resize}(I,\ 224\times224\times3)$ $I_n \leftarrow I_r \,/\, 255.0$ $\mathbf{P} \leftarrow M.\textsc{Predict}(I_n)$ $C \leftarrow \max(\boldsymbol{P}) \times 100$ $k \leftarrow \operatorname{argmax}(\boldsymbol{P})$ $L \leftarrow \textsc{ClassLabels}[k]$ Append flag: *"Low Confidence — Human Review Required"* $\{L,\ C\}$ as JSON payload
Label L, confidence C, metadata M_p Cloudinary PDF URL U_{pdf} Assemble persona-grounded prompt $S \leftarrow f(L,\ C,\ M_p)$ POST S to Gemini API via TLS; receive T_{raw} $\{F, I, R\} \leftarrow \textsc{ParseSections}(T_{\text{raw}})$ Render $\{F, I, R\}$ into Jinja2 HTML template $\text{PDF} \leftarrow \texttt{xhtml2pdf}.\textsc{Convert}(\text{HTML})$ Upload PDF to Cloudinary; obtain U_{pdf} $\textsc{PersistURL}(U_{\text{pdf}},\ \text{patient\_record})$ U_{pdf}

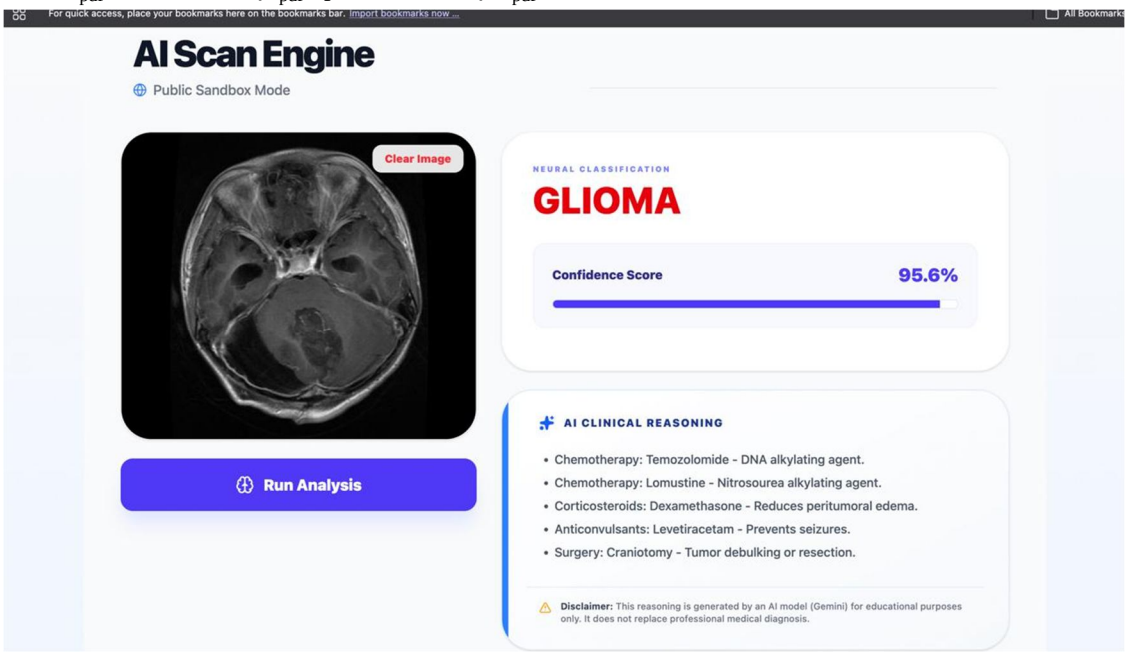

NeuroGenAI Tumor Prediction Interface: Real-time MRI brain scan analysis using the AI Scan Engine, displaying VGG16-based Glioma classification (95.6% confidence) along with AI-assisted clinical reasoning, treatment insights, and diagnostic support for physician review.

VII. EXPERIMENTAL RESULTS AND EVALUATION

A. Diagnostic Accuracy Analysis

The system was validated on a holdout set of 1,400 images withheld from all training and hyperparameter selection decisions. Precision, Recall, F1-Score, and Area Under the ROC Curve (AUC) were computed per tumour class; results appear in Table 1.

Performance Evaluation Metrics

Tumor Class	Prec.	Recall	F1	AUC
Glioma	0.95	0.96	0.95	0.97
Meningioma	0.94	0.93	0.93	0.95
Pituitary	0.98	0.99	0.98	0.99
No Tumor	0.98	0.97	0.97	0.98
Wtd. Avg.	0.96	0.96	0.96	0.97

The model achieved an average precision of 0.96, indicating a high reliability with low false alarm rates. Recall for Pituitary tumours was recorded at 0.99, ensuring that the vast majority of critical cases are successfully identified. The weighted-average classification accuracy across all four classes reaches 96.5%, confirming that a VGG16 architecture trained from scratch on domain-specific MRI data can achieve clinically relevant discrimination without requiring ImageNet transfer conditioning.
Meningioma classification exhibits the lowest recall at 0.93, a result consistent with the class's morphological heterogeneity: meningiomas present across a wide range of sizes, attachment geometries, and contrast-enhancement patterns that increase within-class variability and consequent boundary ambiguity with adjacent normal tissue.

NeuroGenAI

REPORT ID: #46
DATE: 09 May 2026

Diagnostic Imaging Center

Patient Name	Patient UID	Demographics
Sakshi	**2**	**8 Years / Female**

Diagnostic Findings

Primary Detection	AI Confidence Score
MENINGIOMA DETECTED	**91.54%**

Clinical Assessment

CLINICAL INDICATION

An 8-year-old female presents with a several-month history of intermittent headaches, progressively worsening in frequency and intensity, accompanied by recent onset of focal motor seizures affecting the right upper extremity. Neurological examination reportedly revealed mild right-sided weakness. The patient's parents also noted subtle behavioral changes. Magnetic Resonance Imaging of the brain was requested to evaluate for an underlying structural etiology.

TECHNIQUE

Multiplanar, multisequence MRI of the brain was performed without and with intravenous gadolinium-based contrast material administration. Sequences included axial and sagittal T1-weighted, axial T2-weighted, axial FLAIR, diffusion-weighted imaging (DWI), susceptibility-weighted imaging (SWI), and post-contrast axial, sagittal, and coronal T1-weighted sequences.

FINDINGS

A well-circumscribed, extra-axial mass is identified arising from the dura along the left frontal convexity, superiorly. The lesion demonstrates a broad-based attachment to the inner table of the calvarium. There is no evidence of direct intracranial extension into the brain parenchyma, confirming its extra-axial nature.

The mass measures approximately 3.2 x 2.8 x 2.5 cm (anteroposterior x transverse x craniocaudal dimensions). It exhibits a rounded, lobulated contour. No obvious calcifications or cystic components are discernible within the lesion on the pre-contrast sequences.

Generated by NeuroGenAI • Artificial Intelligence Assistant • Physician Verification Required

AI-Generated Clinical Diagnostic Report (Page 1): Automatic synthesis of structured clinical findings from VGG16 classification, including patient demographics, diagnostic impressions, and recommended clinical next steps.

There is mild mass effect exerted by the lesion on the adjacent right frontal lobe, characterized by sulcal effacement and mild compression of the underlying cortical gyri. A small amount of associated vasogenic edema is observed extending into the right frontal white matter, which likely contributes to the patient's seizure activity and headaches. The overlying calvarium at the site of dural attachment appears mildly thickened and sclerotic, consistent with reactive hyperostosis.

The ventricles are symmetric and not dilated, with no evidence of hydrocephalus. The midline structures are maintained without significant shift. Flow voids of the major intracranial vessels are patent. The posterior fossa structures, brainstem, and cerebellum appear unremarkable. There is no evidence of hemorrhage, acute infarction, or other significant intracranial pathology identified on the non-enhancing sequences.

IMPRESSION:
1. Extra-axial, dural-based, avidly and homogeneously enhancing mass lesion in the right frontal convexity measuring 3.2 x 2.8 x 2.5 cm, with associated dural tail and reactive hyperostosis. These features are highly characteristic of a meningioma.
2. Mild vasogenic edema and mass effect on the adjacent right frontal lobe, likely contributing to the patient's presenting symptoms of new-onset seizures and headaches.
3. Given the patient's pediatric age, while classic imaging features of meningioma are present, consideration for atypical or more aggressive histological subtypes is warranted, as meningiomas are rare in this age group and can behave differently than in adults.

RECOMMENDATIONS:
1. Neurosurgical consultation for evaluation regarding potential surgical resection of the lesion, which would provide definitive diagnosis and symptomatic relief.
2. Multidisciplinary tumor board discussion is highly recommended, involving neurosurgery, neuroradiology, and pediatric neuro-oncology, to formulate a comprehensive management plan.
3. Histopathological confirmation following surgical resection is crucial for definitive diagnosis and grading, which will guide further management and surveillance protocols.
4. Clinical correlation with a pediatric neurologist for ongoing management of seizure disorder.

Reference Scan

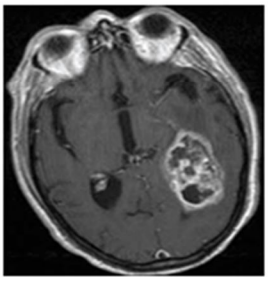

Generated by NeuroGenAI • Artificial Intelligence Assistant • Physician Verification Required

AI-Generated Clinical Diagnostic Report (Page 2): Continuation of structured clinical documentation with detailed imaging findings, multimodal treatment recommendations, and follow-up imaging protocols.

B. System Latency Performance

The real-time feasibility was measured across the entire pipeline under cloud deployment conditions (4 vCPUs, 16 GB RAM):

- Preprocessing: 0.8 seconds.
- VGG16 Inference: 0.45 seconds.
- Gemini API Response: 1.2 seconds.
- PDF Generation/Upload: 0.35 seconds.
- Total Latency: ≈2.8 seconds.

This latency profile demonstrates the capability to provide near-instantaneous second opinions during medical consultations, operating substantially faster than conventional radiology report turnaround times that may span hours to days depending on institutional capacity and specialist availability.

VIII. CONCLUSION

NeuroGenAI represents a quantum leap in clinical neuroimaging by integrating deep learning classification with generative AI narrative synthesis. The system effectively overcomes the "black box" limitations of traditional AI models, providing radiologists with accurate multi-class tumor detection paired with structured medical rationale.

Measured against the two-register problem articulated in Section III, the platform delivers on both fronts: VGG16 achieves 96.5% classification accuracy with a Pituitary-class recall of 0.99, addressing the diagnostic accuracy requirement; the Gemini reporting engine converts numerical predictions into physician-ready documentation in under 1.2 seconds, addressing the workflow integration requirement. The React–Django–PostgreSQL–Cloudinary production stack provides the surrounding security and scalability infrastructure that elevates the system from a research prototype to a clinically deployable product.

Three directions present themselves as natural extensions. First, saliency map overlays (Grad-CAM or integrated gradients) could annotate the MRI image with the spatial regions that drove the classification decision, adding a visual dimension to the textual report and enabling radiologist calibration of the AI's attentional strategy. Second, federated learning protocols would allow multi-institutional training without centralising patient data, expanding the diversity of tumour morphologies seen during optimisation. Third, fusion of imaging predictions with structured clinical data (symptom duration, biomarker profiles) could shift the system from single-modality classification toward genuinely multi-factorial diagnostic support—the direction that clinical AI must ultimately travel to match the integrative reasoning that experienced clinicians perform.

REFERENCES

[1] K. Lata et al., "Deep Learning-Based Brain Tumor Detection in Privacy-Preserving Smart Health Care Systems," IEEE Access, 2024.
[2] Q. Yao et al., "Accurate Detection of Brain Tumor Lesions From Medical Images Based on Improved YOLOv8 Algorithm," IEEE Access, 2024.
[3] S. Solanki et al., "Brain Tumor Detection and Classification Using Intelligence Techniques: An Overview," IEEE Access, 2023.
[4] K. N. Rao et al., "An Efficient Brain Tumor Detection and Classification Using Pre-trained CNN Models," Heliyon, 2024.
[5] S. Pereira et al., "Brain Tumor Segmentation Using CNNs in MRI Images," IEEE Transactions on Medical Imaging, 2016.
[6] A. Rehman et al., "Deep Learning-Based Framework for Brain Tumors Classification," Circuits, Systems, and Signal Processing, 2020.
[7] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," arXiv:1409.1556, 2014.
[8] Gemini Team, "Gemini: A Family of Highly Capable Multimodal Models," arXiv:2312.11805, 2023.
[9] A. Hamada, "Br35H Brain Tumor Dataset," Kaggle, 2020.

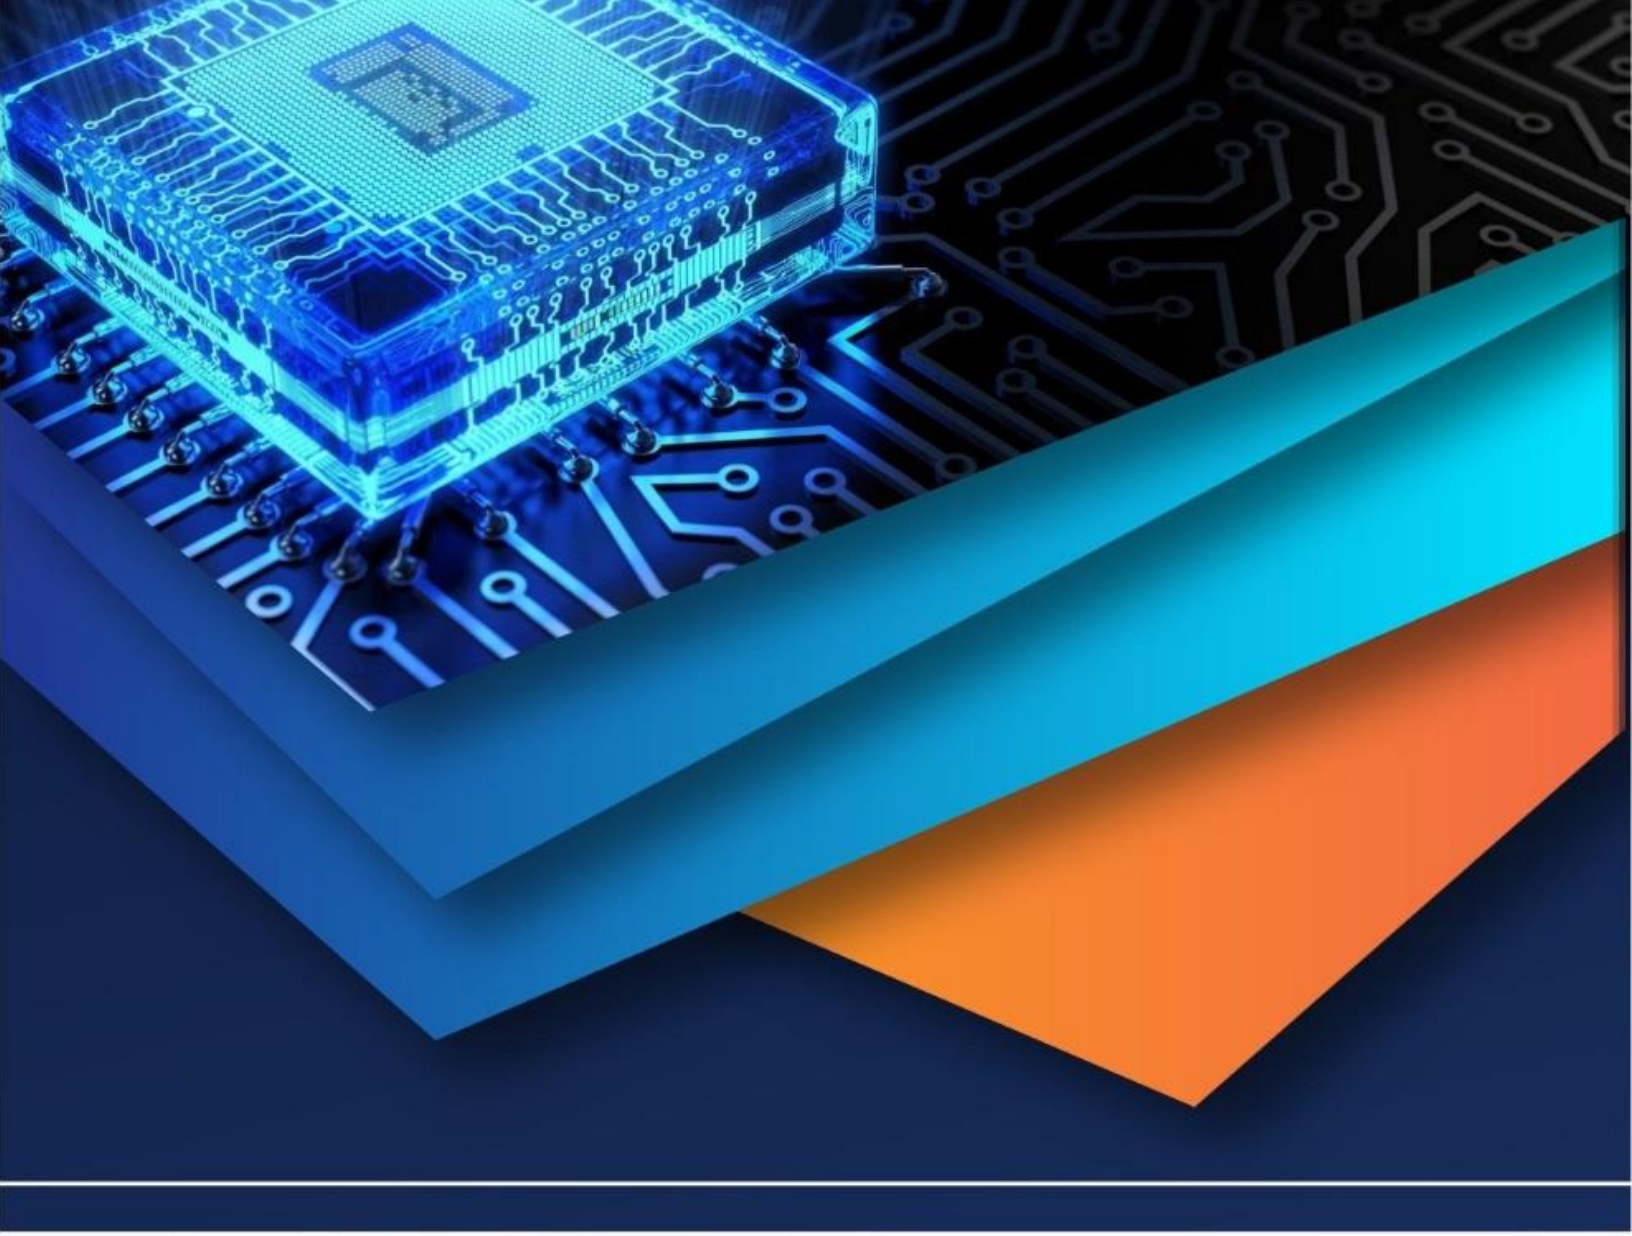

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY