



---

# Curevia: An Integrated AI Framework for Early Detection of Neurological Disorders, Brain Tumours, and Ocular Diseases Using EEG, MRI, and Retinal Imaging

Lakshay<sup>1</sup>, Ritika<sup>2</sup>, Nancy<sup>3</sup>, Param<sup>4</sup>, Rajni<sup>5</sup>

AI and DS Department, Panipat Institute of Engineering and Technology, Samalakha, Haryana, India<sup>1,2,3,4,5</sup>

[batralakshay2@gmail.com](mailto:batralakshay2@gmail.com)<sup>1</sup>, [ritikataneja.info@gmail.com](mailto:ritikataneja.info@gmail.com)<sup>2</sup>, [nancy2003saini@gmail.com](mailto:nancy2003saini@gmail.com)<sup>3</sup>,  
[paramnamdev131001@gmail.com](mailto:paramnamdev131001@gmail.com)<sup>4</sup>, [rajni.cse-@piet.co.in](mailto:rajni.cse-@piet.co.in)<sup>5</sup>

**Abstract.** *The exponential growth of neurological and ophthalmic disorders worldwide has intensified the demand for diagnostic frameworks that are precise, scalable, and accessible beyond the confines of specialist-intensive clinical settings. This paper presents Curevia, a multi-modal artificial intelligence platform engineered to deliver early, automated identification of brain tumors, ocular pathologies, and neurological anomalies by harnessing three complementary data streams: Magnetic Resonance Imaging (MRI), Optical Coherence Tomography (OCT) retinal scans, and Electroencephalogram (EEG) signals. The architecture employs Convolutional Neural Networks for spatial feature extraction from volumetric brain scans and retinal images, while Recurrent Neural Networks and Long Short-Term Memory networks decode temporal patterns embedded in EEG time-series data. A data fusion mechanism integrates predictions across all three modalities, yielding a holistic diagnostic profile that surpasses the capability of any single-modality approach. Experimental evaluation demonstrates tumor segmentation accuracy reaching 94%, retinal disease classification at 92%, and the system processes each imaging study within six to ten seconds, substantially outperforming conventional baseline CNN models in both accuracy and processing efficiency. The platform further incorporates an interactive, web-based interface built with Streamlit, enabling seamless image and signal uploads, real-time result visualization, and personalized clinical recommendations. This work establishes Curevia as a clinically viable, computationally efficient, and ethically compliant diagnostic assistant capable of bridging the accessibility gap between specialist expertise and patient need.*

**Keywords:** Convolutional Neural Networks, EEG Signal Processing, Brain Tumor Detection, Ocular Disease Classification, Multi-modal Medical AI, Deep Learning Healthcare.

## Introduction

Neurological and visual disorders are all one of the most clinically and socioeconomically burdensome disease classifications of modern medicine. Brain tumors, including malignant gliomas, meningiomas and pituitary adenomas, still have the same high mortality rates in case of diagnosis after the late clinical stage. Ocular conditions they include Choroidal Neovascularization (CNV), Diabetic Macular Edema (DME), and age-related Drusen formation which are top. reasons behind irreversible vision loss, especially among populations with metabolic comorbidities. Simultaneously, neurological disorders that can be identified by EEG-epileptic seizure disorders, sleep disorders, and cortical abnormalities-affect hundreds of millions in

---



---

the world but often fail to be detected in time because neural signal data is so difficult to interpret. Traditional diagnostic processes are strongly dependent on the presence of specialists, organized imaging conditions, and prolonged turnaround times. An ophthalmologist who looks at scans of the retina, a radiologist who looks at MRI volumes, and a neurologist. EEG traces readings are all conducted in local clinical silos. This fragmentation subjects diagnostic delays, increases risks of interpretive inaccuracy, and makes screening at an early stage practically unattainable in a healthcare setting under resource constraints settings. The deficit is especially severe in semi-urban and rural geographies with the density of specialists being a fraction of clinical demand. Deep learning and computer vision have become a groundbreaking technology in medical diagnostics, providing the ability to scale high-dimensional imaging data quickly and accurately that is challenging and in many benchmarks faster than, experienced human interpretation. Convolutional Neural Networks have been shown to achieve state-of-the-art results in tumor delimiting of boundaries, grading of retinal lesions, and localization of anomalies. Nevertheless, most of the current AI diagnostic tools focus on one disease domain at a time, not taking advantage of the existing clinically meaningful correlations between domains neurological and eye symptoms. Curevia came up as a solution to this architectural drawback. It started as a minor-project platform in which the focus was on MRI-based tumor diagnosis and eye disease classification via OCT, the system has been extended over its major-project phase to incorporate. Third diagnostic modality EEG signal analysis. This development allows a truly multi-modal evaluation process: MRI characterizes pathology of the structural brain, OCT measures retinal layering and vascular abnormality, and EEG records the electrophysiological state of neural networks. By combining these three streams of information in one platform, Curevia will be in a position to do so to provide diagnostic information of a richness and a richness that is impossible to achieve through individual modalities.

## **II Related Work**

The use of deep learning in medical imaging has a rich and fast-growing literature. Foundational contributions determined the dominance of CNN architectures over conventional computer vision algorithms thresholding, edge detection, and support vector machines- to classify brain MRI datasets [1]. Later Kermany et al. proved that massive OCT datasets, combined with CNN-based architectures trained using transfer learning, were able to be used to achieve expert results level relative accuracy in the grading of retinal diseases such as age-related macular degeneration and diabetic retinopathy [2]. The encoder-decoder architectures have developed segmentation of tumor regions on brain MRI, especially U-Net and its volumetric extensions, as they allow accurate visualization of tumor boundaries without the need to use manual region-of-interest annotation [3]. Hybrid networks that use CNNs to extract spatial features and Recurrent Neural Networks to sequential context have also enhanced the prediction of tumor growth and longitudinal disease tracking [4]. Islam and Rastegarnia studied the developmental trends of non-invasive Brain-Computer Interface applications in the EEG field. and emphasized the deep learning transformative nature in decoding intricate temporal neural patterns that conventional spectral. There are no methods which can reliably characterize [5]. Long Short-Term Memory networks have demonstrated specific potential with regard to seizure. continuous EEG prediction, providing enough lead time to take clinical action. Although these rich modality-specific studies exist, combined systems that simultaneously record MRI, OCT, and EEG signals in a single diagnostic system are rare. The multi-modal systems that are currently in use normally combine two modalities only. typically MRI including clinical metadata, and



---

uncommonly to electrophysiological signal combination. Curevia fills in this gap by integrating all three streams of data and generating composite diagnostic results using a modality fusion layer.

### **III System Architecture and Methodology**

#### **A. Architectural Overview**

The Curevia system is structured into four main processing steps; ingestion and preprocessing, feature extraction and. Detecting region-of-interest, multi-modal classification, and creating reports with secured data storage. Each stage is as a software module to allow upgrading of modules separately without interfering with the broader pipeline.

#### **B. Data Acquisition and Preprocessing**

Users submit medical data through a Streamlit-based web interface, which accepts MRI scans, OCT retinal images, and EEG recordings in standard formats (JPEG, PNG, DICOM for imaging; EDF, BDF for EEG). Before any model inference occurs, all inputs pass through a multi-step preprocessing pipeline.

For imaging data, preprocessing encompasses resolution normalization to a consistent spatial dimension, Gaussian and bilateral filtering for noise suppression, Contrast-Limited Adaptive Histogram Equalization (CLAHE) to amplify clinically relevant tissue contrast, and orientation alignment to ensure anatomical consistency across heterogeneous imaging sources. MRI volumes additionally undergo skull-stripping to isolate parenchymal structures.

EEG signals are preprocessed through a dedicated pipeline built on the MNE-Python library. This includes bandpass filtering to retain clinically informative frequency bands (delta: 0.5–4 Hz, theta: 4–8 Hz, alpha: 8–13 Hz, beta: 13–30 Hz, gamma: >30 Hz), Independent Component Analysis (ICA) for ocular and muscular artifact removal, epoch segmentation, and baseline normalization. The resulting cleaned signal segments are converted to time-frequency representations using Short-Time Fourier Transform or continuous wavelet decomposition, generating spectrotemporal feature maps suitable for input to convolutional layers.

#### **C. Feature Extraction and Region-of-Interest Detection**

Spatial feature extraction from imaging data is performed by CNN architectures that progressively encode local low-level patterns—edges, textures, intensity gradients—into high-level semantic representations corresponding to disease-specific tissue signatures. U-Net segmentation models delineate tumor boundaries in MRI scans and isolate retinal layers in OCT volumes, producing binary segmentation masks that define the Region of Interest (ROI) for downstream classification. Mask R-CNN is employed for instance-level lesion localization, enabling the system to enumerate and independently characterize multiple distinct abnormalities within a single scan.

For EEG data, two-dimensional spectrotemporal maps are processed by a CNN encoder, which extracts spatial features from frequency-time representations. These spatial embeddings are subsequently passed to an LSTM layer that models inter-epoch temporal dependencies, capturing the dynamic evolution of neural activity over extended recording windows—a critical capability for seizure pattern recognition and cortical arousal assessment.

#### **D. Multi-modal Classification and Fusion**

Classification is performed by dedicated model branches for each modality. Brain tumor classification employs VGG16 fine-tuned on the Kaggle Brain Tumor MRI Dataset, yielding a probability distribution over the classes Glioma, Meningioma, Pituitary Tumor, and No Tumor. Retinal disease classification is handled by MobileNetV3, pre-trained on the Mendeley OCT Dataset and fine-tuned to classify CNV,

---



DME, Drusen, and Normal retinal states. EEG-based neurological classification is performed by the CNN-LSTM model described above, predicting the likelihood of abnormal neural patterns indicative of epileptiform activity, sleep disorder, or cortical dysfunction.

A modality fusion layer aggregates the probabilistic outputs of all three branches using a learned weighted averaging scheme, in which weights are conditioned on input quality scores assigned during preprocessing. This approach downweights predictions from low-quality inputs—such as motion-corrupted MRI volumes or EEG recordings with high residual artifact—while amplifying contributions from high-confidence data streams. The fused output represents the system's composite diagnostic assessment.

Confidence calibration is applied to all outputs through temperature scaling, ensuring that reported probability scores reliably reflect empirical accuracy. Cases in which the fused confidence score falls below a predefined threshold are automatically flagged for manual review by a qualified clinician, maintaining a safety buffer against uncertain AI-generated diagnoses.

#### E. Technology Stack

The implementation stack comprises Python 3.9 as the primary development language, TensorFlow and Keras for CNN and segmentation model development, PyTorch for RNN and LSTM implementations, MNE-Python for EEG signal processing, OpenCV and Scikit-image for image preprocessing, Streamlit for the user-facing web interface, FastAPI for backend service orchestration, and MongoDB for encrypted patient data persistence. Cloud deployment is provisioned on AWS/GCP infrastructure, with AES-256 encryption applied to all data in transit and at rest, ensuring compliance with HIPAA-equivalent data protection frameworks.

## IV Experimental Results

### A. Hardware and Evaluation Environment

Experimental evaluation was conducted on a system equipped with an Intel Core i7 (13th Gen) processor, an NVIDIA RTX 4060 GPU, 32 GB DDR4 RAM, and a 1 TB NVMe SSD. Model training for MRI classification and OCT disease detection was performed using the GPU-accelerated TensorFlow backend, while EEG model training employed PyTorch with CUDA support. Operating systems used were Ubuntu 20.04 LTS for training and deployment, and Windows 11 for interface testing.




Figure 1: Curevia dual-panel interface — Brain MRI Analysis Platform (left) and OCT Retinal Analysis Platform (right), each offering drag-and-drop upload, detection categories, and confidence level indicators.



Figure 2: Curevia main dashboard presenting the unified AI-Powered Medical Disease Detection Platform, with navigation to Eye Disease Detection and Brain Tumour Detection modules.



Figure 3: File browser for selecting MRI scan images (meningioma test set shown). Users select JPEG/PNG brain scan images for AI-powered tumour classification.



Figure 4: Live diagnostic output from the OCT Eye Disease Detection module. The uploaded retinal scan is classified as CNV (Choroidal Neovascularization) with Analysis Complete status displayed in real time.



### B. Accuracy Performance

Table 1 presents the accuracy comparison between the Curevia multi-modal system and a baseline single-modality CNN across the three diagnostic tasks. Curevia consistently outperforms the baseline across all evaluated domains.

**Table 1:** Diagnostic Accuracy Comparison — Curevia vs. Baseline CNN

Diagnostic Module	Curevia (%)	Baseline CNN (%)
Brain Tumor Segmentation	94	88
Retinal Disease Classification	92	85
Preprocessing Quality Score	96	90

### C. Processing Time

End-to-end processing time for a single MRI or OCT image—encompassing preprocessing, model inference, and result generation—ranges from six to ten seconds. The breakdown is approximately three seconds for image preprocessing, four seconds for model prediction, and one to two seconds for database logging and report assembly. This performance profile is clinically significant, as it enables real-time diagnostic support during patient consultations.

**Table 2:** Processing Time Comparison (Seconds per Study)

Processing Stage	Curevia (s)	Baseline CNN (s)
MRI / OCT Preprocessing	3	4
Tumor / Eye Prediction	4	5
Database Logging	1–2	4

### D. Error Reduction Analysis

A comparative analysis of diagnostic error rates between manual clinical assessment and Curevia's automated pipeline reveals a substantial reduction in both false positive and false negative outcomes. Table 3 summarizes the error counts recorded during evaluation studies.

**Table 3:** Manual Diagnosis vs. Curevia Error Counts

Process Type	Manual Errors	Curevia Errors
Tumor Detection	15	3
Eye Disease Detection	12	4

The 80% reduction in tumor detection errors and the 67% reduction in ocular diagnosis errors underscore the clinical value of removing subjective interpretive variability from diagnostic workflows. These gains are attributable to the combination of rigorous preprocessing—which eliminates artifact-induced



misclassification—and the depth of feature representations learned by large-capacity pre-trained architectures fine-tuned on domain-specific medical datasets.

### **V EEG-Based Neurological Disorder Detection — A Novel Extension**

The integration of EEG signal analysis represents the most clinically novel contribution of Curevia's major-project phase. While MRI and OCT capture static morphological and structural information, EEG encodes the dynamic functional state of the brain across time, offering a complementary and in many respects irreplaceable perspective on neurological health.

#### **A. Clinical Motivation**

Only epilepsy is seen to affect around 50 million individuals worldwide, but most of the patients in low- and middle-income environments. have never been subject to EEG testing because of the cost of the equipment, the special interpretation needs, and the laboratory. infrastructure typically required. The EEG module developed by Curevia is aimed at reducing this barrier, by automating the identification of. ictal, interictal and normal EEG patterns, indicating recordings which should be examined by a specialist and give a preliminary probabilistic description of anomalies detected. In addition to epilepsy, EEG abnormalities can reveal cortical slowing of dementia, metabolic encephalopathy, and disrupted sleep architecture, circumstances in which EEG offers information that cannot be obtained by imaging studies alone. The combination of structural MRI and electrophysiological EEG data in the same diagnostic encounter thus provide a much more enriched clinical picture that each modality alone would give.

#### **B. EEG Processing Pipeline**

Raw EEG data is consumed by clinical acquisition systems or wearable EEG. After the rejection of the artifacts by means of. The cleaned signal is divided into fixed-length epochs, which is followed by ICA and bandpass filtering. Every epoch is changed into a two. dimensional spectrotemporal map short-time fourier analysis to create a representation in which the x-axis encodes time, y-axis represents frequency, color strength represents spectral power. This transformation makes EEG data. compatible with convolutional features extraction. The spectrotemporal maps are encoded by a CNN and yield fixed-dimensional embedding vectors. These embeddings are inputted into a stacked LSTM model which learns time-dependent interactions across epochs, and allows the model to identify. dynamic patterns the gradual accumulation of ictal activity in pre-seizure epochs would be unobservable to models which process each epoch separately. The LSTM output is sent to a softmax depth of classification that assigns. likelihoods between categories of neurological states.

#### **C. Multi-modal Fusion with EEG**

Introduction of the EEG branch to the existing Curevia MRI and OCT pipeline meant the need to develop a cross-modal fusion. mechanism that can be used to reconcile predictions based on fundamentally different data types. The approach adopted is concatenation the feature vectors of the penultimate layer of both modality branches and feeds the concatenated representation into a shared fully connected network which learns cross-modal correlations. In a case of a patient who comes with both imaging and EEG data, the fused representation records interactions like co-occurrence of tumor-induced mass effect on MRI with the corresponding. cortical slowing of EEG, which gives finer estimates of probability than any branch alone can produce. When EEG data are not available--as is true in patients who may present only with imaging studies--the system gracefully. no longer requires the two-modality or single-modality inference, but leaves the pipeline intact.



## **VI Discussion**

### **A. Comparative Advantages**

Curevia has a number of substantive benefits compared to the current single-modality medical AI tools. First, its multi-modal architecture embodies manifestations of disease on structural, functional, and metabolic levels in simultaneity, which is reflective of the comprehensive evaluation which skilled clinicians make through the course of various specialist visits. Second, the confidence calibration mechanism gives a clear estimate of uncertainty as well as each diagnosis, which make human-AI responsible cooperation whereby the system amplifies problematic cases instead of making a commitment to the possibility of making incorrect decisions. Third the built-in Streamlit interface removes the need of users to have technical knowledge to use it, making it the platform available to general practitioners, paramedical employees and possibly patients who are pursuing self-managed health monitoring.

### **B. Limitations**

The present implementation has a number of constraints. The quality of input images can significantly reduce model performance, particularly when it falls below the minimum standards required by the preprocessing validation module, which may be difficult to satisfy in practice. This limitation is challenging to eliminate due to the heterogeneity of imaging equipment used in clinical settings. Similarly, EEG data obtained through wearable consumer-grade devices exhibit substantially higher noise levels compared to hospital-grade acquisition systems. As a result, the current artifact rejection pipeline may not perform adequately across all scenarios. Another important limitation is the computational infrastructure required for real-time multimodal inference. GPU-based servers with sufficient memory bandwidth are necessary, which can pose a barrier to adoption in resource-constrained environments. However, this challenge can be partially mitigated through cloud-based deployment strategies. Additionally, the models have primarily been assessed using publicly available benchmark datasets. Evaluation on real-world clinical data, encompassing diverse patient populations and imaging centers, remains an essential step prior to clinical deployment. Algorithmic bias is another critical concern. Medical AI systems are often trained on skewed datasets with imbalances in demographic representation, imaging procedures, and disease prevalence. These discrepancies may limit the generalizability of the model to the intended target population. Continuous performance evaluation across demographic subgroups, along with the inclusion of underrepresented populations in training datasets, will be crucial to ensure equitable diagnostic performance.

## **VII Future Scope**

The strategy of further development of Curevia covers a variety of opportunities. Wearable EEG continuous, longitudinal neural health monitoring devices would expand the clinical applicability of the platform beyond episodic diagnostic experiences in the direction of proactive health, allowing the provision of seizure forecasting models providing advance warning ictal events hrs before ictal. Scaling to new disease areas cardiac arrhythmia detection using ECG waveforms, classification of dermatological lesions based on dermoscopic images, and pulmonary pathology based on chest radiographs-- they may be accomplished by adding new specialized branches to the existing modular architecture. Federated learning is a technically advanced and yet clinically significant direction of the future. By training model updates feeding gradient information at a central server and only aggregating it and not sending any gradient information to the participating healthcare facilities, federated learning enables the process of improving models using a variety of institutional datasets with no effort to get the patient data out of the institution of



---

origin- a competitive privacy and regulatory benefit. HIPAA Compliance Certification, the EU Medical Device. A requirement to commercial clinical deployment is regulation, and India has a Digital Personal Data Protection Act that will need special validation research and regulatory involvement. Lastly, the creation of mobile and offline-enabled. The rural healthcare workers and community health programs would be able to access healthcare services dramatically through interface versions operating in low-connectivity environments.

### **VIII Conclusion**

Curevia demonstrates that a thoughtfully engineered multi-modal AI framework can substantially improve upon both the accuracy and the efficiency of traditional diagnostic pathways for neurological and ophthalmic disorders. By combining Convolutional Neural Networks for MRI and OCT analysis with LSTM-based EEG signal interpretation and fusing the outputs of all three modalities, the system achieves 94% tumor segmentation accuracy and 92% retinal disease classification accuracy while processing each clinical study in under ten seconds. Error rates relative to manual diagnosis are reduced by 67–80%, underscoring the practical clinical value of automated, artifact-resistant interpretation. The EEG integration introduced in Curevia's major-project phase represents a meaningful step toward genuinely holistic AI- assisted diagnosis—one that encompasses not only the structural anatomy captured by imaging but also the functional electrophysiology that reveals the living brain's operational state. Together, these capabilities position Curevia as a clinically viable, ethically responsible, and technically scalable contribution to the growing field of AI-enabled healthcare, with the potential to improve diagnostic outcomes for underserved populations who currently lack reliable access to specialist medical expertise.

### **References:**

- [1] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- [2] Kermany, D. S., et al. (2018). Deep learning models for OCT-based retinal disease classification. *Cell*, 172(5), 1122–1131.
- [3] Wang, Y., Liu, L., & Song, W. W. (2022). Deep learning based brain tumor segmentation: A survey. *Journal of Ambient Intelligence and Humanized Computing*.
- [4] Al Showarah, S., Bibi, R., Muthanna, M. S. A., & Ayoubi, T. M. A. (2023). Detection and classification of brain tumor using hybrid deep learning based fine-tuned EfficientNets. *Scientific Reports*.
- [5] Islam, M. K., & Rastegarnia, A. (2023). Recent advances in EEG (non-invasive) based BCI applications. *Frontiers in Computational Neuroscience*, 17, 1151852.
- [6] Géron, A. (2019). *Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow*. O'Reilly Media.
- [7] Szeliski, R. (2020). *Computer Vision: Algorithms and Applications*. Springer.
- [8] Gururaj, H. L., Flammini, F., Kumar, V. R., & Prema, N. S. (2023). Recent trends in healthcare using AI-driven medical imaging. *Frontiers in Medicine*.