

Early Stroke and Panic Detection using AI in Neurology

Sachin Phuldev Yadav, Prof. Bhosale S.B, Prof. Bangar A.P, Prof. A.A. Khatri

Department of Computer Engineering, Jaihind College of Engineering, Pune, Maharashtra, India

Abstract- Neurological emergencies span a wide spectrum, from sudden vascular events such as ischemic stroke to acute psychiatric-neurological crises such as panic attacks. Both categories are time-sensitive and both currently rely heavily on manual clinical assessment, which limits the speed and consistency of detection. This paper presents a detailed review of six recent peer-reviewed sources (2024–2025) on artificial intelligence (AI) in neurology, and uses their combined findings to build a case for, and a design of, an integrated two-module early-detection framework: an imaging-based deep-learning pipeline for stroke triage, and a wearable-sensor-based sequence-modelling pipeline for panic/anxiety detection. The review covers foundational AI concepts (machine learning, deep learning, neural network architectures), the current clinical evidence base for stroke-imaging AI, the comparatively nascent evidence for physiological/behavioural AI in affective and psychiatric-neurological states, and the shared ethical, technical, and regulatory challenges — interpretability, bias, data privacy, and validation — that recur across every source reviewed. The paper then details a proposed system architecture, dataset and preprocessing requirements, candidate model architectures, an evaluation plan with specific metrics, expected outcomes, and a discussion of limitations and future work. The overall conclusion is that stroke-detection AI is a mature and clinically validated field that can serve as a design template for a still-emerging class of wearable, sensor-driven AI tools for detecting acute panic and anxiety states.

Keywords— Artificial Intelligence, Machine Learning, Deep Learning, Convolutional Neural Networks, Recurrent Neural Networks, Stroke Detection, Panic Disorder, Anxiety, Wearable Sensors, Neurology, Clinical Decision Support.

I. INTRODUCTION AND RESEARCH OBJECTIVES

1. Background

Artificial intelligence has moved from a peripheral research interest to a central pillar of modern neurological practice. Rizzo and Dawson (2025) trace this evolution from early cybernetics and perceptron models in the 1940s and 1950s, through the neural-network 'winters' of the 1970s–80s, to the deep-learning resurgence following the 2012 ImageNet competition, and finally to the transformer-based large language models of the 2020s. They argue that neurology and neuroscience have not merely been passive beneficiaries of this evolution but have actively catalysed it, since brain-inspired computation (McCulloch–Pitts neurons, Hebbian learning) supplied many of AI's founding ideas. Within this broad and rapidly evolving field, two categories of neurological emergency stand out for

their time-criticality and their contrasting level of AI maturity: acute ischemic stroke, where treatment efficacy is known to fall sharply with every additional minute of delay, and acute panic or anxiety episodes, which — although rarely immediately fatal in isolation — can precipitate serious cardiovascular strain, escalate into psychiatric crises, and are almost never monitored objectively outside of controlled clinical settings.

2. Problem Statement

Despite decades of clinical protocol refinement, both stroke and panic/anxiety detection still depend on a human observer being present at the right moment: a radiologist reading a scan, a triage nurse applying a stroke scale, or a patient recognising and reporting their own symptoms. This creates three structural weaknesses. First, diagnostic delay: specialist availability and manual image interpretation take time that is critical in stroke, where the phrase 'time is brain' reflects the roughly two-million-neuron loss

estimated per minute of untreated large-vessel occlusion. Second, subjectivity: clinical scoring scales and self-reported symptom checklists vary between observers and between patients, reducing consistency. Third, absence of continuous monitoring: neither stroke risk nor panic/anxiety symptoms are tracked outside of clinical encounters, so the earliest, most treatable phase of an episode is frequently missed entirely.

3. Motivation for an AI-Based Approach

AI offers a technically plausible way to close this gap because it can continuously and tirelessly analyse large volumes of imaging, physiological, and behavioural data at a speed and consistency no human observer can match. As Kalani and Anjankar (2024) summarise, AI systems can now perform many of the 'sophisticated perceptual and cognitive capacities of biological systems, such as object identification and decision making,' which is precisely the capability needed to flag an evolving stroke on a CT scan or an evolving panic episode on a wearable sensor stream before the patient or clinician might otherwise notice.

4. Scope and Objectives of This Paper

- Review foundational AI/ML/DL concepts as they specifically apply to neurology, drawing on the taxonomy proposed by Rizzo and Dawson (2025).
- Synthesise current clinical evidence for AI-assisted stroke detection and triage from four independent reviews (Bösel et al., 2025; Gutman et al., 2024; AbuAlrob & Mesraoua, 2024; Kalani & Anjankar, 2024).
- Examine the comparatively limited but growing evidence base for AI-assisted detection of panic, anxiety, and related affective-neurological states, extrapolating from adjacent validated use cases such as seizure forecasting and depression-treatment prediction.
- Propose a concrete, two-module system architecture combining both detection pipelines under a single clinician-supervised dashboard.
- Detail a feasible methodology — dataset requirements, preprocessing, model selection, and evaluation metrics — for implementing and validating such a system.

- Critically discuss the ethical, technical, and regulatory barriers that recur across all six reviewed sources, and outline future research directions.

II. RELATED WORK / LITERATURE REVIEW

1. Foundational AI Concepts in Neurology

Rizzo and Dawson (2025) provide the most comprehensive taxonomy among the reviewed sources, situating AI as the broadest category, with machine learning (ML) as a subset that 'develops algorithms that help computers learn from experience,' and deep learning (DL) as a further subset built on multi-layered neural networks capable of automatically extracting hierarchical features from raw data rather than relying on hand-engineered features. They classify ML into five methodological 'tribes' — symbolic, connectionist, evolutionary, Bayesian, and analogical — each mapped to specific neurology applications: symbolic systems for rule-based diagnostic decision support (e.g., distinguishing multiple sclerosis from other demyelinating conditions), connectionist (neural-network) systems for image-based diagnosis such as tumour or stroke detection, evolutionary algorithms for optimising treatment parameters such as spinal-cord-stimulation settings, Bayesian methods for probabilistic diagnosis and disease-progression prediction, and analogical methods (e.g., support vector machines, k-nearest neighbours) for classifying conditions by similarity to previously diagnosed cases. Namusisi (2025) and Kalani and Anjankar (2024) reach a compatible but simpler conclusion, framing ML and DL as the two dominant paradigms driving neurological AI today, with DL's capacity to learn directly from raw imaging or sensor data being the key enabler of its strong performance in tasks such as stroke-lesion segmentation and Alzheimer's disease detection from MRI.

2. AI for Stroke Detection and Triage

All four clinically-oriented sources converge on the finding that AI has already achieved measurable, quantified impact in stroke care, making this the most mature use case reviewed.

Imaging-based detection

Bösel et al. (2025) report that deep learning 'can considerably enhance the detection of strokes or vessel occlusions,' and that AI-based vascular-imaging tools have moved beyond research settings into FDA-approved clinical products — specifically naming RapidASPECTS for detecting early and advanced infarcts and FastStroke for perfusion imaging. They further note that ML/DL approaches are increasingly combining imaging data with history, laboratory tests, and chart records to support triage, treatment decisions, and admission pathways, and that outcome prediction after stroke is shifting from traditional rigid scoring systems toward AI-driven 'precision' or 'personalised' approaches.

Quantified workflow impact

Gutman et al. (2024) provide the most concrete performance figures among the six sources: a convolutional-neural-network-based algorithm that automatically detects large-vessel occlusions (LVOs) on CT and autonomously alerts the stroke treatment team via a mobile app was associated with an average reduction of 52 minutes in time to LVO-treatment initiation. They also cite AI-based automated Alberta Stroke Program Early CT Score (ASPECTS) calculation achieving performance comparable to expert neuroradiologists, and note that AI-based infarct-volume and treatment-outcome prediction models have shown accuracy improvements over traditional scoring, including a reported 7% improvement in forecasting symptomatic intracerebral haemorrhage risk in one cited study.

Generalisability and resource-setting considerations

AbuAlrob and Mesraoua (2024) add a cautionary but still positive perspective, stating that AI-driven analysis of CT/MRI can reduce stroke diagnosis time 'from hours to minutes' and can detect subtle occlusions that might be missed on an initial manual read — a benefit they argue is particularly valuable in low-resource settings where neuroradiology expertise is scarce. However, they simultaneously highlight that most existing stroke-AI evidence originates from high-resource environments with

advanced imaging infrastructure, and call for dedicated research into lower-computational-cost AI models suitable for low- and middle-income country deployment.

3. AI for Panic, Anxiety, and Affective-State Detection

Direct literature on AI-based panic-attack detection is essentially absent from the six sources reviewed, which is itself an important finding: it indicates a genuine research gap rather than an oversight in this paper's search. However, three adjacent findings suggest strong technical feasibility for extending existing neurological AI methods to this use case.

Behavioural and speech-based affective analysis

AbuAlrob and Mesraoua (2024) note that 'AI-driven models have been developed to assist in the diagnosis of psychiatric and cognitive disorders like depression, schizophrenia, and mild cognitive impairment,' and that such models 'can analyse speech patterns, facial expressions, and behavioural data to detect early signs of cognitive decline.' These are precisely the signal modalities — voice, facial affect, movement — that would plausibly carry the behavioural signature of an acute panic or anxiety episode.

EEG and connectivity-based mood-state prediction

Gutman et al. (2024) describe validated examples of neural-signal-based classifiers for affective states, including a resting-state-connectivity biomarker study that identified neurophysiological subtypes of depression, and a study using an electroencephalographic signature to predict antidepressant treatment response. These examples demonstrate that EEG and connectivity-derived features are already being used in machine-learning pipelines to characterise mood and affective states, a methodological precedent directly transferable to detecting the acute physiological signature of panic or anxiety.

Wearable-sensor precedent from seizure forecasting

The strongest and most directly transferable precedent comes from epilepsy care. Bösel et al.

(2025) describe FDA-approved wearable devices — specifically the Empatica wristband — that use electrodermal-activity and accelerometer signals to detect seizures in real time and notify caregivers or contacts, and note that similar wearable- and AI-based seizure forecasting methods analyse wristband sensor data to predict seizures minutes to hours in advance. Since panic/anxiety episodes, like seizures, are transient, physiologically distinct events superimposed on a baseline signal, the same class of architecture — continuous wearable monitoring feeding a sequence model trained to recognise a

characteristic onset pattern — is a natural candidate for adaptation.

4. Comparative Summary of Reviewed Sources

Table 1 summarises the AI techniques, primary application area, and single most relevant finding from each of the six reviewed sources, illustrating both the concentration of evidence around stroke/imaging AI and the comparative scarcity of panic/anxiety-specific evidence.

Table 1. Comparative summary of AI techniques and findings across the six reviewed sources.

Source	Primary AI Technique(s)	Application Area	Key Finding Relevant to This Paper
Bösel et al. (2025)	CNN (vascular imaging), wearable sensor + RNN-type analytics	Stroke, epilepsy, movement disorders	FDA-approved LVO/ASPECTS tools in clinical routine; Empatica wearable detects seizures via electrodermal activity — a direct template for panic detection.
Gutman et al. (2024)	CNN (LVO detection), EEG/connectivity ML classifiers	Stroke workflow, depression prediction	Automated LVO alerting reduced treatment time by ~52 minutes; EEG-based classifiers already predict antidepressant response.
AbuAlrob & Mesraoua (2024)	CNN/DL for imaging, NLP/speech/facial-expression AI	Stroke, TBI, spinal injury, psychiatric-cognitive disorders	Stroke diagnosis time cut from hours to minutes; speech/facial-expression AI used for depression and cognitive-decline detection.
Kalani & Anjankar (2024)	ML and DL (general)	Neuroimaging, BCIs, treatment optimisation	AI performs pattern recognition and decision-making comparable to biological systems across neurology tasks.
Rizzo & Dawson (2025)	Full AI/ML/DL taxonomy, CNN, RNN/LSTM, transformers	Cross-cutting neurology and neuroscience	Provides architecture-level detail (CNNs for imaging, RNNs/LSTMs for sequential physiological data) directly usable for system design.
Namusisi (2025)	ML and DL (general)	Prediction and management of neurological disorders	Confirms ML/DL as dominant paradigms; highlights ethical and data-privacy concerns common to all AI-neurology applications.

5. Synthesis: Literature Gaps

Three consistent gaps emerge from this synthesis. First, panic/anxiety-specific AI detection models are

essentially absent from the current neurology-AI literature reviewed here, which instead concentrates on stroke, epilepsy, depression, and cognitive

decline — indicating an open research opportunity rather than an established sub-field. Second, the stroke-AI evidence that does exist is described by its own authors as predominantly retrospective and often single-centre (Bösel et al., 2025), raising generalisability concerns that would apply with at least equal force to any newly developed panic-detection model trained on a necessarily smaller initial dataset. Third, every one of the six sources — regardless of its specific clinical focus — independently raises the 'black-box' interpretability problem and the risk of algorithmic bias as unresolved barriers to clinical trust (AbuAlrob & Mesraoua, 2024; Rizzo & Dawson, 2025; Namusisi, 2025), confirming that these are field-wide concerns rather than issues specific to any one application.

III. PROPOSED FRAMEWORK

Building on the gaps and precedents identified above, this paper proposes an integrated two-module AI framework for early detection of stroke and panic/anxiety episodes. The framework is deliberately designed to be clinician-supervised at every stage rather than autonomous, in direct response to the interpretability and trust concerns raised throughout the literature.

1. System Overview

The system consists of two independent but architecturally parallel detection pipelines (Module A for stroke, Module B for panic/anxiety) that both report into a shared Integration and Alerting Layer. This design allows each module to be developed, validated, and — critically — regulatorily approved on its own timeline, since Module A relies on mature, already-validated techniques while Module B is exploratory.

Table 2. Side-by-side pipeline design for Module A (stroke) and Module B (panic/anxiety).

Pipeline Stage	Module A — Stroke Detection	Module B — Panic/Anxiety Detection
Input data	Non-contrast CT / CT-angiography	Continuous HRV, EDA, respiration,

Pipeline Stage	Module A — Stroke Detection	Module B — Panic/Anxiety Detection
	images; vitals; history	optional single-channel EEG
Preprocessing	Normalisation, skull-stripping, artefact removal	Denoising, missing-data imputation, feature extraction (e.g., HRV frequency bands)
Core model	CNN (ResNet/U-Net-style) for classification + segmentation	RNN / LSTM or GRU for sequence classification
Output	LVO flag, infarct segmentation, automated ASPECTS score	Panic/anxiety-onset probability score with time-window
Action on positive detection	Automated alert to stroke response team	Alert to patient/caregiver + optional clinician dashboard flag

2. Module A — Imaging-Based Stroke Detection

Module A follows the pattern already validated in commercial stroke-AI tools described by Bösel et al. (2025) and Gutman et al. (2024). A convolutional neural network — for example a U-Net-style architecture for pixel-level segmentation combined with a ResNet-style classifier for occlusion detection, consistent with the network types catalogued by Rizzo and Dawson (2025) — is trained on labelled non-contrast CT and CT-angiography images to (i) classify the presence or absence of a large-vessel occlusion, (ii) segment the infarct core and surrounding penumbra, and (iii) compute an automated ASPECTS score. A positive detection at

any stage triggers an automated, time-stamped notification to the stroke response team, directly mirroring the workflow-acceleration effect quantified by Gutman et al. (2024).

3. Module B — Wearable-Sensor-Based Panic/Anxiety Detection

Module B adapts the wearable-sensor architecture used for seizure forecasting (Bösel et al., 2025) to the panic/anxiety use case. Continuous streams of heart-rate variability, electrodermal activity, respiration rate, and — where available — single-channel EEG are fed into a recurrent neural network, specifically an LSTM or GRU, following the sequential-data architectures catalogued by Rizzo and Dawson (2025) as well-suited to time-series with long-range temporal dependencies. The network is trained to recognise the characteristic physiological signature preceding or accompanying an acute panic/anxiety episode. As a secondary, corroborating signal available during clinical or telehealth encounters, optional NLP-based speech analysis and computer-vision-based facial-expression analysis — the same modalities

AbuAlrob and Mesraoua (2024) describe for depression and cognitive-decline detection — can be layered on top to reduce false-positive alerts from physiological noise alone (e.g., exercise-induced heart-rate elevation).

3. Module B — Wearable-Sensor-Based Panic/Anxiety Detection

Module B adapts the wearable-sensor architecture used for seizure forecasting (Bösel et al., 2025) to the panic/anxiety use case. Continuous streams of heart-rate variability, electrodermal activity, respiration rate, and — where available — single-channel EEG are fed into a recurrent neural network, specifically an LSTM or GRU, following the sequential-data architectures catalogued by Rizzo and Dawson (2025) as well-suited to time-series with long-range temporal dependencies. The network is trained to recognise the characteristic physiological signature preceding or accompanying an acute panic/anxiety episode. As a secondary, corroborating signal available during clinical or telehealth encounters, optional NLP-based speech analysis and computer-

vision-based facial-expression analysis — the same modalities

4. Integration and Alerting Layer

Both modules report to a shared clinical dashboard, which a neurologist, emergency physician, or (for Module B) the patient's care team supervises. Rather than presenting a bare binary verdict, the dashboard displays a confidence score alongside the specific evidence driving it — the flagged imaging region and its segmentation overlay for Module A, or the physiological trend segment and its key contributing features for Module B. This design choice is a direct, deliberate response to the explainability concerns raised by every source in the review (Rizzo & Dawson, 2025; AbuAlrob & Mesraoua, 2024; Namusisi, 2025).

IV. METHODOLOGY

1. Data Requirements

Module A requires labelled CT/CT-angiography datasets with confirmed LVO status and expert-annotated ASPECTS scores, ideally aggregated across multiple institutions to mitigate the single-centre bias explicitly flagged by Bösel et al. (2025) as a limitation of current neurocritical-care AI studies. Module B requires continuous physiological time-series data labelled with clinician- or self-reported panic/anxiety episode onset and offset times — a dataset type that does not yet exist at scale in the public literature reviewed here and would need to be prospectively collected, following a data-collection model similar to the seizure-diary and wearable datasets already established for epilepsy research (Bösel et al., 2025).

2. Preprocessing Pipeline


Imaging data for Module A would undergo standard normalisation, skull-stripping, and artefact-removal steps, consistent with the preprocessing pipeline Rizzo and Dawson (2025) describe for neuroimaging AI (MRI/CT normalisation, feature extraction from EEG/MEG signals, and statistical handling of missing data). Physiological signal data for Module B would require denoising, missing-data imputation, and feature engineering — for example extracting frequency-domain HRV features (low-

frequency/high-frequency power ratio) known to correlate with autonomic arousal — prior to sequence modelling.

3. Model Selection and Training

Both modules would use a standard 70/15/15 train/validation/test split, with k-fold cross-validation applied during training to mitigate overfitting risk, as recommended by the general AI-in-neurology methodology guidance in Rizzo and Dawson (2025). For Module A, transfer learning from pre-trained medical-imaging backbones is recommended to reduce the data volume required for acceptable performance, given the relatively limited size of labelled stroke-imaging datasets compared to general computer-vision corpora. For Module B, given the anticipated scarcity of labelled panic/anxiety physiological data in an initial pilot phase, a semi-supervised or few-shot learning approach may be necessary, potentially bootstrapped from the larger existing seizure-detection wearable datasets through domain adaptation.

4. Architecture



5. Evaluation Metrics

Both modules would be evaluated using sensitivity, specificity, precision, recall, F1-score, and area under the receiver-operating-characteristic curve (AUC), consistent with standard practice for diagnostic AI models in the reviewed literature. In addition, Module A would be evaluated on time-to-alert as a workflow metric, directly comparable to the 52-minute reduction reported by Gutman et al. (2024),

while Module B would be evaluated on false-alarm rate per patient-day, a usability metric critical to sustained patient adherence with a wearable device.

6. Clinical Validation Plan

- Retrospective validation: test both modules against existing labelled datasets (stroke-imaging archives for Module A; any available wearable panic/anxiety pilot data for Module B) before any prospective deployment.
- Prospective single-centre pilot: deploy Module A in an emergency-department workflow and Module B on a small volunteer cohort with diagnosed panic/anxiety disorder, comparing AI-flagged events against clinician-confirmed ground truth.
- Multicentre validation: extend to multiple institutions to test generalisability across patient populations and equipment types, directly addressing the single-centre bias concern raised in the literature.
- Regulatory review: pursue approval pathways comparable to those already achieved by commercial stroke-imaging AI tools (e.g., FDA clearance processes referenced for RapidASPECTS and the Empatica device).

V. EXPECTED RESULTS AND DISCUSSION

1. Expected Outcomes

Based on the performance figures already reported for comparable, already-deployed systems, Module A would be expected to achieve sensitivity and specificity for LVO detection in the range reported for existing commercial tools (approaching or exceeding neuroradiologist-level accuracy per Gutman et al., 2024), with a proportional reduction in time-to-treatment. Module B, being exploratory, would more realistically be expected to achieve moderate initial sensitivity with a non-trivial false-alarm rate in its first pilot phase, improving through iterative retraining as more labelled panic/anxiety physiological data becomes available — a trajectory consistent with how wearable seizure-forecasting accuracy has improved over successive device generations (Bösel et al., 2025).

2. Discussion of Feasibility

The literature reviewed here supports the technical feasibility of both modules individually, since each reuses architectures and signal types already validated in closely related neurological applications — CNN-based imaging analysis on one side, and RNN/LSTM-based wearable-sensor monitoring (as used for seizures) on the other. The integration of the two into a single system is, however, a novel proposal rather than an established practice, and several open questions would need to be resolved before deployment.

3. Ethical and Regulatory Considerations

- Clinical validation: panic/anxiety-detection models would need prospective, multicentre trials analogous to those beginning to appear for stroke-AI tools, rather than being deployed on the strength of retrospective or adjacent-domain evidence alone.
- Bias and generalisability: as AbuAlrob and Mesraoua (2024) caution, models trained predominantly on data from specific populations may perform unevenly across ethnic and socioeconomic groups; this risk applies equally to physiological wearable data, which can vary with skin tone (affecting optical heart-rate sensors) and body composition (affecting electrodermal-activity baselines).
- Interpretability: both Rizzo and Dawson (2025) and AbuAlrob and Mesraoua (2024) identify the 'black-box' nature of deep learning as a persistent barrier to clinician trust, reinforcing the design choice in this proposal to surface underlying evidence rather than a bare verdict.
- Privacy and consent: continuous physiological monitoring raises data-governance questions not present in one-off imaging studies, requiring clear informed consent, secure data storage, and compliance with health-data regulations, a concern Namusisi (2025) raises explicitly for AI in neurology generally.
- Human oversight: consistent with every source reviewed, the proposed system is designed to augment — not replace — clinical judgement, with final treatment or intervention decisions always resting with a human clinician or, for Module B, the patient's own care plan.

4. Limitations of This Study

This paper is a literature-based design proposal rather than an empirical study; no models have been trained or tested, and the expected-results figures cited are drawn from performance already reported for related, previously validated systems rather than from original experimentation. The panic/anxiety detection module in particular rests on extrapolation from adjacent use cases (seizure forecasting, depression prediction) rather than on any directly published panic-detection study, since none was identified among the six sources reviewed. Any future implementation would need to treat Module B as a genuinely exploratory research effort requiring its own dedicated data-collection and validation programme before clinical claims could be made.

VI. CONCLUSION AND FUTURE WORK

This paper has presented a detailed review of six recent sources on AI in neurology and used them to construct a proposed integrated framework for early detection of stroke and panic/anxiety episodes. The review confirms that AI-based stroke detection is a mature, clinically validated field — with FDA-approved imaging tools already reducing treatment delays by tens of minutes in real-world deployments — while AI-based panic/anxiety detection remains an open research opportunity that could plausibly borrow architectures from adjacent, already-validated applications such as wearable seizure forecasting and EEG-based mood-state prediction. Future work should prioritise: (i) prospectively collecting labelled panic/anxiety physiological datasets in collaboration with psychiatric and neurology clinics; (ii) conducting a small-scale pilot study of the proposed wearable module using domain-adapted seizure-detection architectures as a starting point; (iii) iteratively improving Module B's sensitivity and false-alarm rate through successive retraining cycles, following the precedent set by wearable seizure-forecasting devices; and (iv) pursuing multicentre validation and, ultimately, regulatory review comparable to that already achieved by commercial stroke-imaging AI tools. Successfully combining both modules could extend the proven benefits of AI-accelerated neurological triage to a broader class of acute neurological and

psychiatric-neurological events, meaningfully improving both physical and mental-health outcomes for patients experiencing time-critical neurological crises.

REFERENCES

1. Bösel, J., Mathur, R., Cheng, L., Varelas, M. S., Hobert, M. A., & Suarez, J. I. (2025). AI and Neurology. *Neurological Research and Practice*, 7(11).
2. Gutman, B., Shmilovitch, A.-H., Aran, D., & Shelly, S. (2024). Twenty-Five Years of AI in Neurology: The Journey of Predictive Medicine and Biological Breakthroughs. *JMIR Neurotechnology*, 3, e59556.
3. AbuAlrob, M. A., & Mesraoua, B. (2024). Harnessing artificial intelligence for the diagnosis and treatment of neurological emergencies: a comprehensive review of recent advances and future directions. *Frontiers in Neurology*, 15:1485799.
4. Kalani, M., & Anjankar, A. (2024). Revolutionizing Neurology: The Role of Artificial Intelligence in Advancing Diagnosis and Treatment. *Cureus*, 16(6): e61706.
5. Rizzo, M., & Dawson, J. D. (2025). AI in Neurology: Everything, Everywhere, All at Once — Part 1: Principles and Practice. *Annals of Neurology*, 98(2), 211–230.
6. Namusisi, H. N. (2025). The Role of Artificial Intelligence in Neurology: Predicting and Managing Disorders. *Eurasian Experiment Journal of Biological Sciences*, 6(1), 18–23.