International Journal of Web of Multidisciplinary Studies
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: https://ijwos.com OPEN ACCESS E-ISSN : 3049-2424
Volume 3 Issue 9 , September 2026 DOI: https://doi.org/10.71366/ijwos

NeuroScan AI: An Explainable Deep Learning Framework for Brain Hemorrhage Detection and Clinical Decision Support Using CT Scan Images

Krupa raani S[1], Varshitha N[2]
[1] *Assistant Professor, Department of Computer Science, GM University Davangere,India*
[2] *Student, Department of Computer Science, GM University Davangere, India.*

Article Info

Article History:

Published: 02 Sept 2026

Publication Issue:
Volume 3, Issue 9
September-2026

Page Number:
14-20

Corresponding Author:
S,Varshitha N

Abstract:
Brain hemorrhage is a life-threatening neurological emergency in which rapid diagnosis directly affects patient survival and long-term recovery. In current clinical practice, CT scans are read manually by radiologists, a process that is accurate but slow and heavily dependent on the availability of experienced specialists, particularly in busy emergency departments. This paper introduces NeuroScan AI, an explainable deep learning framework built to detect brain hemorrhage from CT scans and support clinicians with automated, easy-to-understand diagnostic reasoning. At its core, the system uses a Convolutional Neural Network (CNN) to classify CT images as hemorrhage or non-hemorrhage. What sets NeuroScan AI apart is that its predictions are not left as raw labels — they are passed through OpenAI's language model to produce clear, clinically relevant explanations that help physicians understand why a given classification was made. The framework is delivered as a secure, web-based platform that also stores patient records and diagnostic reports for later reference. In testing, the system produced fast, accurate, and interpretable predictions, shortening the time needed to reach a preliminary diagnosis while giving clinicians more confidence in the AI's output. These results suggest that NeuroScan AI could meaningfully improve the efficiency, reliability, and accessibility of brain hemorrhage diagnosis in real-world healthcare settings.

Keywords: Brain Hemorrhage Detection, Deep Learning, Convolutional Neural Network (CNN), Explainable AI (XAI), CT Scan Analysis, OpenAI, Clinical Decision Support, Medical Image Analysis

1. INTRODUCTION

Bleeding inside or around the brain — commonly referred to as brain hemorrhage — is one of the most urgent conditions a hospital can encounter. It can be triggered by head trauma, stroke, or the rupture of a blood vessel, and because brain tissue is so sensitive to pressure and oxygen deprivation, even a short delay in treatment can leave a patient with permanent neurological damage or worse. Computed Tomography (CT) remains the imaging method of choice for spotting this kind of bleeding, mainly because it is fast and widely available. The catch is that someone still has to read the scan, and that job usually falls to a radiologist. In a quiet department that works fine, but during a surge of emergency admissions, waiting for an available specialist can cost precious minutes.

Over the past several years, artificial intelligence — and deep learning in particular — has started to close that gap. Convolutional Neural Networks (CNNs) now perform impressively well at recognizing patterns in medical images, hemorrhage detection included. The problem is that most of these models operate as black boxes: they output a label, hemorrhage or no hemorrhage, without explaining why. For clinicians who are ultimately responsible for the patient, a prediction with no reasoning attached is hard to trust, and that lack of trust is a real barrier to getting AI tools adopted in practice.

This paper responds to that gap with NeuroScan AI, a framework designed to detect brain hemorrhage from CT scans and, just as importantly, explain its own reasoning. The detection side of the system is handled by a CNN trained to separate hemorrhage cases from normal scans. What makes the framework different is what happens next: each prediction is passed to OpenAI's language model, which converts the raw classification into a clinical explanation a physician can actually read and evaluate. The goal is not to replace the radiologist's judgment but to give it better support.

Beyond detection and explanation, NeuroScan AI is built as a complete web platform. Clinicians can upload scans, view results, and keep a running history of patient records and diagnostic reports in one place. Taken together — accurate detection, AI-generated explanation, and organized record-keeping — the framework is aimed at shortening the time to diagnosis and giving healthcare teams more confidence in using AI-assisted tools during time-critical situations.

2. STATEMENT OF THE PROBLEM

Brain hemorrhage diagnosis today still depends almost entirely on a radiologist manually reading a CT scan. That approach works, but it is slow by nature and tied directly to how much clinical expertise happens to be available at the moment a patient arrives. When emergency departments are busy, that dependency becomes a bottleneck, one that can directly affect whether a patient is treated in time.

A number of deep learning models already exist for automated hemorrhage detection, but most stop at classification: they indicate whether hemorrhage is present, and nothing more. Without an explanation attached, it is difficult for clinicians to build confidence in what the model is doing, which in turn slows real-world adoption. On top of that, most of these systems include no way to manage patient records or generate interpretable clinical summaries, features that matter just as much as raw accuracy when a tool is meant to fit into an actual hospital workflow.

What is missing, then, is a framework that brings these pieces together: reliable deep learning-based classification, an explainable layer that turns predictions into language a clinician can reason about, secure handling of patient data, and a platform simple enough to use under pressure. That is the gap this work sets out to address.

3. STUDY METHODOLOGY

A. Description of Study Area

This study centers on building and evaluating an AI-assisted system for brain hemorrhage detection, intended for use in hospitals and diagnostic centers where CT imaging is part of routine care. The framework gives clinicians a secure web interface where they can upload a scan and receive both a classification and an AI-generated explanation of what that classification means.

Under the hood, a CNN handles the actual image classification, separating hemorrhage cases from normal ones. To make those results useful in a clinical setting rather than just accurate on paper, the predictions are routed through a large language model (OpenAI) that turns them into readable diagnostic summaries. The platform also keeps a secure, searchable record of patient information, prior diagnoses, and generated reports.

The scope of the study covers designing, building, and testing this framework, and evaluating it on prediction accuracy, response speed, how understandable its explanations are, and how usable the interface is for clinical staff. It is worth being clear about what this system is not: it is not a replacement for a radiologist's judgment. It is meant to sit alongside that judgment, speeding up the early stages of diagnosis and giving clinicians a second, transparent opinion to work with.

B. Data Collection and Analysis

NeuroScan AI was trained and evaluated on a publicly available brain CT dataset containing both hemorrhage and non-hemorrhage cases. Before training, the images went through standard preprocessing, including resizing, normalization, and augmentation, to help the model generalize better and avoid overfitting to the training set.

At inference time, a clinician uploads a scan through the web application, the trained CNN produces a classification, and that result is passed to OpenAI to generate a clinical explanation written in plain, understandable language. Everything relevant to that encounter, the patient's information, the uploaded image, the prediction, the confidence score, the AI-generated explanation, and the final report, is stored securely in a PostgreSQL database for future access.

To judge how well the system actually performs, standard classification metrics were used: accuracy, precision, recall, F1-score, and the confusion matrix. Looking at deep learning predictions alongside the AI-generated explanations together, rather than either in isolation, gives a fuller picture of how useful the framework would be in an actual clinical decision-support role.

4. RESULTS AND DISCUSSION

NeuroScan AI was implemented as a complete, web-based diagnostic platform rather than a standalone model, bringing together image classification, AI-assisted interpretation, patient record management, and report generation under one system. Where conventional workflows rely entirely on a radiologist's manual read of the scan, this framework is built to produce a fast, explainable prediction that a clinician can use as a starting point for their own assessment.

The frontend was built with React.js and styled with Tailwind CSS, with Express.js running the backend logic. Patient information, scan images, prediction history, and generated reports all reside in a PostgreSQL database. The classification itself is handled by a CNN trained on brain CT images, and each of its predictions is sent to OpenAI, which turns the raw output into a clinical explanation clinicians can reason about.

A. Brain Hemorrhage Detection System

The core goal here was straightforward: give clinicians a tool that can look at a CT scan and flag hemorrhage quickly. After logging in, a clinician lands on a dashboard showing patient records, past diagnoses, and an option to upload a new scan. Once an image is uploaded, the CNN analyzes it and returns a hemorrhage or non-hemorrhage classification.

That whole process, from upload to result, takes only a few seconds. Alongside the classification, the system shows a confidence score and the AI-generated explanation, so the result is not just a label sitting on a screen; it comes with enough context for a clinician to make sense of it.

B. CT Scan Image Analysis

The trained model consistently classified uploaded CT scans into hemorrhage and non-hemorrhage categories. Each image was resized, normalized, and standardized before being fed to the CNN, which kept the input consistent regardless of how the original scan was captured.

Across testing, the model picked up on hemorrhage patterns reliably and held steady performance across different scans. Compared to a manual read, the automated step reduced significantly the time needed to reach a preliminary result.

C. Prediction Results

Every prediction generated by the system came bundled with the following information:

• Hemorrhage or non-hemorrhage classification.

• Prediction confidence score.

• AI-generated clinical interpretation.

• Associated patient information.

• Date and time of the diagnosis.

• A downloadable diagnostic report.

Having both the numbers, in the form of a confidence score, and the narrative, in the form of an AI explanation, side by side gave clinicians more to work with than a bare label, and that combination is what builds trust in an AI-assisted result.

D. AI-Assisted Clinical Interpretation

Perhaps the most distinctive part of this framework is what happens after classification. Instead of stopping at a label, NeuroScan AI sends the result to OpenAI, which generates a written explanation covering what was detected, why it might matter clinically, and what kind of follow-up attention it may warrant.

That explanation is what turns the tool from a plain classifier into something closer to a decision-support assistant. It does not replace a clinician's judgment, but it gives them something concrete to reason with instead of a single word.

E. Patient Record Management

Every diagnosis run through the system left behind a secure digital record, including the scan, the prediction, the confidence score, the AI explanation, and the report, all stored in PostgreSQL. Clinicians could search past cases, pull up historical diagnoses, and regenerate reports as needed.

Having that history centralized made it easier to track a patient over time and reduced the kind of paperwork that tends to accumulate in manual record-keeping.

F. Initial Performance Observations

Throughout development and testing, the system held up well across every stage. Upload, preprocessing, CNN inference, explanation generation, and report creation all ran without noticeable lag, and the interface stayed responsive throughout.

Getting deep learning, AI-generated explanation, database storage, and a web interface all working together smoothly confirmed that a system like this is genuinely feasible to build and use, not just feasible on paper. The results overall suggest NeuroScan AI addresses several practical shortcomings of existing hemorrhage detection tools by pairing accuracy with actual explainability.

The subsections that follow look more closely at prediction accuracy, the quality of the AI-generated explanations, system performance, and how the framework compares against existing approaches.

G. Prediction Accuracy Analysis

Accuracy was, understandably, the first thing that needed verification. The trained CNN was tested on a held-out set containing both hemorrhage and non-hemorrhage images, and evaluated using accuracy, precision, recall, F1-score, and a confusion matrix.

The model correctly flagged the large majority of hemorrhage cases while keeping both false positives and false negatives low. Because the whole pipeline runs automatically, the time saved compared to a manual read was substantial, evidence that deep learning can genuinely speed up the early stages of diagnosis without sacrificing reliability.

H. Explainable AI Evaluation

Most deep learning models in this space stop at the classification label. NeuroScan AI goes a step further by having OpenAI generate a written explanation for every prediction, summarizing the detected condition in language a clinician does not need a machine learning background to understand.

That extra layer is what moves the needle on transparency; it gives clinicians a real diagnostic description to weigh, not just a yes-or-no output, which in turn builds more confidence in using the tool without asking clinicians to defer to it blindly.

I. System Performance

Performance held up well across the board. Upload, preprocessing, CNN inference, AI explanation generation, and report creation together completed quickly enough to be genuinely useful in emergency scenarios rather than just meeting a benchmark in a lab setting.

The platform also stayed responsive under multiple concurrent requests, which speaks well of how the frontend, backend, model, OpenAI integration, and database were tied together, with no noticeable bottlenecks slowing the system down.

J. Patient Report Generation

After every prediction, the system automatically compiles a diagnostic report containing:

• Patient details.

• The uploaded CT scan image.

• The AI-generated clinical interpretation.

• The date and time of diagnosis.

These reports are stored securely and can be retrieved later, which helps with documentation and makes it easier to track a patient's history across multiple visits.

K. Comparative Analysis with Existing Systems

Table I lays out how NeuroScan AI compares against more traditional hemorrhage diagnosis approaches.

Feature	**Existing Systems**	**Proposed NeuroScan AI**
Manual CT scan analysis	Required	AI-assisted
Automated hemorrhage detection	Limited	Yes
Deep learning classification	Limited	Yes
Explainable AI	No	Yes
AI-generated clinical interpretation	No	Yes
Patient record management	Basic	Comprehensive
Report generation	Basic	Automated
Web-based platform	Limited	Yes
Fast diagnosis support	Limited	Yes

Clinical decision support	No	Yes

TABLE I. COMPARISON OF EXISTING SYSTEMS AND THE PROPOSED NEUROSCAN AI FRAMEWORK

The comparison makes clear that NeuroScan AI brings more to the table than accuracy alone. The combination of explainable AI, automatic report generation, and AI-driven clinical interpretation is what sets it apart from more conventional systems, and what makes it a better fit for real clinical use.

L. System Effectiveness

Overall, the implementation met the goals it was designed around. The CNN reliably picked up on hemorrhage cases, and the OpenAI-generated explanations gave those predictions the kind of clinical context that makes them easier to trust.

The record-management side of the platform made storing and retrieving diagnostic reports far less of a hassle than paper-based systems, and because it is web-based, clinicians can use it from any authorized device on the network rather than being tied to one workstation.

Put together, the system speeds up diagnosis, adds a layer of transparency that is usually missing from black-box models, and gives healthcare professionals more to work with when making time-sensitive treatment decisions.

5. DISCUSSION

Building NeuroScan AI reinforced something that is becoming clearer across medical AI research: pairing deep learning with explainability is not just a nice-to-have, it changes how usable a system actually is in practice. Traditional hemorrhage diagnosis leans entirely on a radiologist's manual read; this framework instead pairs automated classification with AI-generated reasoning, which makes the whole process more transparent without removing the human from the loop.

The OpenAI integration is really what does the heavy lifting on usability. It is the difference between a system that simply reports "hemorrhage detected" and one that explains what that finding might mean and why. That distinction matters, because the framework is not meant to replace radiologists; it is meant to give them a faster, more transparent starting point during an emergency.

The secure platform, automatic reports, and centralized patient records round out the practical side of things, the parts that determine whether a system like this could actually be adopted in a hospital rather than remaining a research prototype. Taken as a whole, NeuroScan AI addresses the problem it set out to solve: giving clinicians an accurate, explainable, and fast way to support hemorrhage diagnosis.

6. CONCLUSION

This paper presented NeuroScan AI, an explainable deep learning framework built for brain hemorrhage detection and clinical decision support from CT scans. By pairing a CNN with an explainable AI layer, the system addresses a real shortcoming in existing tools, namely accurate classification with essentially no reasoning behind it, and instead offers both a prediction and a clinical explanation clinicians can actually evaluate.

Testing showed the framework reliably analyzing CT scans, producing dependable hemorrhage predictions, and generating readable clinical explanations through OpenAI. The secure record-management and automatic reporting features add real practical value on top of that, making documentation and access to diagnostic history considerably easier. Because it is delivered as a web platform, clinicians can use it efficiently without adding friction to their existing workflow.

None of this replaces a clinician's expertise; the framework is meant to make diagnosis faster, more transparent, and more reliable, not to make the final call. Looking ahead, there is room to extend this work toward multi-class hemorrhage classification, severity estimation, lesion localization through attention mechanisms, integration with hospital information systems, support for other imaging modalities such as MRI, and deployment in real-time clinical settings at a larger scale.

References

[1] G. Litjens, T. Kooi, B. E. Bejnordi, et al., "A survey on deep learning in medical image analysis," Medical Image Analysis, vol. 42, pp. 60–88, 2017.

[2] D. Shen, G. Wu, and H.-I. Suk, "Deep learning in medical image analysis," Annual Review of Biomedical Engineering, vol. 19, pp. 221–248, 2017.

[3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," Advances in Neural Information Processing Systems (NeurIPS), vol. 25, 2012.

[4] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in Medical Image Computing and Computer-Assisted Intervention (MICCAI), 2015.

[5] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016.

[6] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.

[7] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed. Hoboken, NJ, USA: Pearson, 2021.

[8] OpenAI, "GPT-4 technical report," 2023.

[9] TensorFlow, "TensorFlow documentation," Google, 2025.

[10] React Team, "React documentation," Meta Platforms Inc., 2025.

[11] Express.js Team, "Express.js documentation," 2025.

[12] PostgreSQL Global Development Group, "PostgreSQL documentation," version 17, 2025.

[13] Tailwind Labs, "Tailwind CSS documentation," version 4, 2025.

[14] A. Esteva, A. Robicquet, B. Ramsundar, et al., "A guide to deep learning in healthcare," Nature Medicine, vol. 25, pp. 24–29, 2019.

[15] M. T. Islam, S. Aowal, A. Minhaz, and K. Ashraf, "Abnormality detection and localization in chest X-rays using deep convolutional neural networks," IEEE Access, vol. 5, pp. 2210–2218, 2017.

[16] World Health Organization (WHO), "Neurological disorders: Public health challenges," WHO Press.

[17] R. S. Pressman and B. R. Maxim, Software Engineering: A Practitioner's Approach, 9th ed. New York, NY, USA: McGraw-Hill, 2020.

[18] Next.js Team, "Next.js documentation," Vercel Inc., 2025.

[19] Mozilla Developer Network (MDN), "Web APIs documentation," 2025.

[20] T. M. Mitchell, Machine Learning. New York, NY, USA: McGraw-Hill, 1997.