IJARSCT
ISSN: 2581-9429
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
Volume 6, Issue 1, July 2026
Impact Factor: 8.2

SRRIS: Stroke Risk and Recovery Intelligence System

Manav Mendhe[1], Krushna Nagalkar[2], Om Ninave[3],
Lavish Rahangdale[4], Nihar Salunke[5], Priyanka Ujjainkar[6]
Department of Artificial Intelligence[1-6]
G. H. Raisoni College of Engineering and Management, Nagpur, Maharashtra, India

Abstract: *Stroke is one of the top causes of death and long-term disability worldwide. This makes early prediction and diagnosis crucial for effective treatment. This paper introduces the Stroke Risk & Recovery Intelligence System (SRRIS), a healthcare platform that uses machine learning and deep learning for stroke prediction and diagnosis. The system works in two phases. First, it analyzes clinical data like age, hypertension, heart disease, glucose level, and BMI using machine learning models, including Random Forest and Logistic Regression, to predict stroke risk. Second, it processes medical imaging data from MRI and CT scans with Convolutional Neural Networks (CNNs), specifically the U-Net architecture, to find and locate stroke-affected areas. By combining predictive analytics and imaging-based diagnosis, this system offers a solid approach to stroke management. Experimental results show improved accuracy and reliability, emphasizing the potential of SRRIS to aid in early detection and clinical decision-making.*

Keywords: SRRIS, FastAPI, Next.js, Clinical AI, OCR, Stroke Risk Prediction, SHAP, Radiology Analysis, Medical Informatics

I. INTRODUCTION

Recent advancements in Artificial Intelligence (AI) have significantly improved the healthcare sector by enabling intelligent systems capable of assisting clinicians in disease prediction, diagnosis, and treatment planning. Stroke is one of the leading causes of death and long-term disability worldwide, requiring timely diagnosis and appropriate medical intervention. Delayed identification of stroke may result in irreversible neurological damage, permanent disability, or even death. Conventional healthcare systems continue to rely heavily on manual interpretation of patient records, laboratory reports, and diagnostic imaging, making the diagnostic process time-consuming and susceptible to human error. Furthermore, many existing solutions focus only on structured clinical information or medical imaging independently, limiting their ability to provide a comprehensive assessment of a patient's condition. To address these challenges, the proposed Stroke Risk & Recovery Intelligence System (SRRIS) combines structured clinical information with advanced medical image analysis within a unified intelligent platform. The system evaluates stroke risk using machine learning algorithms trained on important clinical attributes such as age, hypertension, heart disease history, blood glucose level, and body mass index. In addition, MRI and CT scan images are processed using deep learning-based segmentation techniques to identify and localize stroke-affected brain regions. By integrating predictive analytics with automated image interpretation, SRRIS supports faster clinical assessment, reduces manual workload, enhances diagnostic consistency, and assists healthcare professionals in making informed treatment decisions during the early stages of stroke management.

II. SYSTEM OVERVIEW

The current SRRIS codebase includes 11 backend endpoint modules, 16 backend service modules, about 10 routed frontend pages, and 5 component-layer files. Key functional areas are: (a) patient registration and profile management, (b) medical history timeline and summary creation, (c) document intake with category-aware OCR processes, (d)

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, July 2026 Impact Factor: 8.2

radiology and ECG scan execution with status updates, and (e) multi-stage AI diagnostic cockpit output using streamed backend stages.

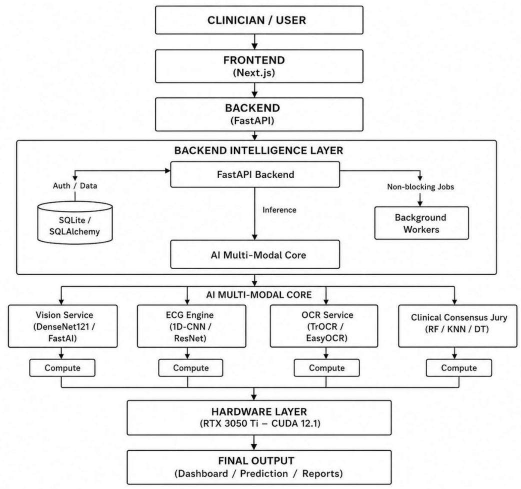

Fig. 1. System Architecture of SRRIS

SRRIS has a layered architecture where the frontend manages user actions, while backend APIs handle data storage, model inference, and progress updates. The modular architecture of SRRIS has been designed to ensure scalability, maintainability, and efficient communication between individual system components. Clinical data processing, medical image analysis, document understanding, and explainable AI modules operate independently while exchanging information through well-defined backend APIs. This architecture enables future integration of additional diagnostic models, hospital information systems, and cloud-based healthcare services without requiring major modifications to the existing framework.

III. RELATED WORK

Previous research has explored different machine learning techniques for stroke prediction. Traditional models like Logistic Regression and Decision Trees have been used for classification tasks but often struggle with complex datasets. Ensemble methods such as Random Forest and XGBoost have shown better performance by combining multiple models. Studies by Tang et al. (2025) and El-Geneedy et al. (2025) demonstrated how explainable AI techniques like SHAP can improve model interpretability.

Deep learning methods, especially Convolutional Neural Networks (CNNs), have been widely applied in medical image analysis. The U-Net architecture has been effective for segmenting stroke lesions in MRI images.

Although considerable progress has been made in stroke prediction and medical image analysis, most existing studies address these tasks independently. Few systems combine structured clinical data, diagnostic imaging, document understanding, and explainable artificial intelligence within a single workflow. SRRIS attempts to bridge this gap by integrating machine learning, deep learning, OCR-based medical document processing, and explainable AI into one clinical decision-support platform. This integrated approach provides healthcare professionals with more comprehensive information for early diagnosis and treatment planning.

IV. METHODOLOGY AND PIPELINE

4.1 Dataset and Feature Engineering

Stroke Risk & Recovery Intelligence System (SRRIS) applies a standard stroke dataset from the healthcare industry consisting of approximately 5,110 patient data points. Important attributes in this dataset include the age, presence of hypertension, heart disease history, mean glucose level, BMI, and other risk factors that determine the occurrence of stroke.

Apart from the structured healthcare data from patients, SRRIS has the capacity to integrate several types of data that include demographic information, medical history, symptom onset timing, medications, laboratory data, ECG data, and image analysis results. These data types provide a holistic view of patient's health.

There is a strict separation of input data and features from each other. While input data refers to raw measurements, scans, and waveforms, features refer to engineered attributes such as HRV metrics, radiology classes, summarized clinical reports, and explainability features (SHAP).

Two major processes are involved in the feature engineering process: normalization and enrichment. Normalization provides consistent units, formats, and textual data, while enrichment helps to extract useful attributes from input data.

4.2 Data Preprocessing

An efficient preprocessing pipeline is designed to guarantee the quality of the data and its consistency. Handling missing values is done using imputation techniques to avoid losing any important information. Encoders transform categorical variables into numerical form.

Normalization and feature scaling guarantee the uniformity of variables. To solve class imbalance issues that are often found in datasets on patients, including those of strokes, the SMOTE technique is used.

4.3 Model Development

Several machine learning algorithms are constructed to forecast stroke probability based on structured clinical data, which include Random Forest, XGBoost, and Artificial Neural Networks. An 80:20 training-testing split of the data set is applied to guarantee a valid assessment of the model on new information.

Each algorithm learns the patterns within the data set individually; however, to enhance its prediction efficiency and minimize the biases introduced by the algorithms, an ensemble method called voting is applied.

4.4 Imaging Model

Along with the processing of the clinical data, SRRIS involves analyzing the medical imaging using deep learning methods. MRI and CT scan images are analyzed in order to recognize areas that have been affected by strokes.

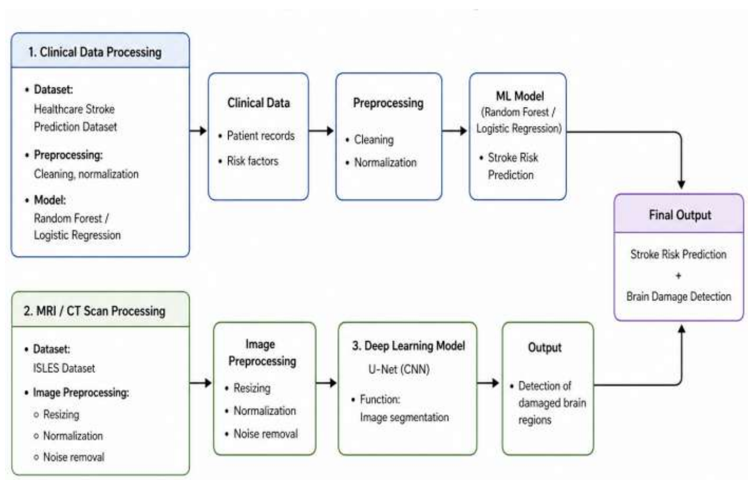

Fig. 2. Overall Architecture of SRRIS System

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, July 2026
Impact Factor: 8.2

A CNN model, which is based on U-Net structure, is applied in order to analyze images. The network allows performing pixel-by-pixel classification, thereby helping to precisely identify areas in brain scans that correspond to the presence of a stroke. Such an approach adds valuable visual information to the stroke risk prediction process.

Figure 1 presents the overall architecture of SRRIS. As seen, the system includes a two-stage pipeline that includes the processing of clinical data and medical imaging. At the clinical data stage, the information about patients undergoes preprocessing and further feeding to a machine learning algorithm used for stroke risk prediction. At the same time, MRI/CT images are processed with a deep learning model, which recognizes stroke areas in brains.

V. BACKEND AND AI COMPONENTS

Back-end provides modular endpoints for authentication, patient management, document management, radiology, ECG, predictive diagnosis, analysis, causality, and survival. AI functionality is scattered among various modules like diagnostic engine, consensus jury, summary engine, ai engine, vision service, and radiology service.The risk estimation algorithm utilizes model-based scoring (XGB/RF/NN) and deterministic gates (e.g., tPA contraindications). SHAP-like explanation modules and SBAR story generation techniques are employed for enhanced clinical interpretation.

5.1 Clinical Diagnostic Ensemble (Tabular Health Data)

This component (named diagnostic_engine.py) is a multi-predictor module, and not a singular model, since multiple factors play into assessing a patient's likelihood to suffer a stroke.

To start off with, we have used the XGBoost classifier as the base prediction model because of its capability of capturing non-linear correlations. This was complemented with the use of a Random Forest model that is used for increasing the variance of predictions. Apart from this, a feed-forward neural network (with 4 layers: 64-64-32-1) was used as another method to capture latent correlations in the data.

Interpretability is maintained using SHAP (SHapley Additive Explanations), wherein we obtain attributions for each input variable based on their contribution to the prediction. In addition to this, we have also implemented a random survival forest model in order to predict recovery paths in the first 90 days after diagnosis.

5.2 ECG Diagnostic Pipeline (Waveform Analysis and Feature Extraction)

The ECG module (ecg_engine/inference_wrapper.py) analyses 12-lead waveform data to identify cardiovascular patterns that indicate susceptibility to stroke.

The architecture utilizes SE-ResNet as the main feature extractor, and Squeeze-and-Excitation blocks improve channel-wise feature weights in the model. The sc_resnet (skip-connected CNN) is an auxiliary network that maintains the finest details of the waveforms and increases detection sensitivity to rare events.

Classic convolutional architectures (ConvV1 and ConvV2) offer localized analysis for particular regions (anterior or inferior leads). Features generated by these deep architectures are merged via CatBoost, a robust machine learning ensemble that handles heterogeneous input types (neural features, engineered features).

Moreover, NeuroKit2 computes HRV and entropy-based features, allowing the architecture to operate as a hybrid system consisting of both deep learning features and physiological descriptions.

5.3 Radiology Pipeline (MRI and CT Imaging)

The radiology module (vision_service.py) carries out classification of brain scans at the slice level, focusing on stroke abnormalities. Pre-trained deep learning models like DenseNet121 and VGG-19 are utilized for feature extraction and classification tasks because of their effectiveness in pattern detection from spatial and textural characteristics of medical imagery. OpenCV techniques are employed in preprocessing before inference, such as contrast adjustment (CLAHE), morphology, and contour segmentation to separate slices.

The workflow of the radiology module is divided into two stages:

(1) Image normalization and segmentation, and

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, July 2026 Impact Factor: 8.2

(2) Semantic classification.

5.4 OCR and Lab Data Parsing (Document Intelligence)

The document processing subsystem (doctor_notes_ocr.py, easyocr_runtime.py) facilitates the transformation of unstructured medical documents into structured clinical inputs.

A handwritten document recognition mechanism based on a TrOCR model is used for recognizing the content of clinical notes and discharge summaries accurately. In case of printed documents, EasyOCR offers a faster solution optimized for structured lab reports.

A multimodal synthesis layer utilizing Gemini 2.0 Flash is used for synthesizing OCR output, laboratory results, and model prediction outputs to produce coherent clinical summaries.

5.5 Cross-Modal Integration and Clinical Orchestration

The key advantage of SRRIS is its capability to incorporate different forms of data sources, such as medical reports, ECG waves, imaging, and textual files.

It works according to a sequential process that involves four steps, including data input, feature extraction, model inference, and output synthesis. The architecture makes SRRIS a decision-support system that supports instead of substitutes the physician’s intuition.

With the aggregation of results from all types of sources, the system generates an explanatory path connecting the risk score with relevant features, wave patterns, and imaging findings.

5.6 Data Flow and Reliability Considerations

The SRRIS uses a structured process for inference and summarization starting from the point where the data is acquired. Although structured data processing in the clinical context seems fairly straightforward, there are still problems related to OCR accuracy, fallback mechanisms for radiology and asynchronous processing.

The handling of missing values by assigning defaults allows for smooth operation of the system but poses a risk of uncertainty in terms of the predictability of the results.

VI. IMPLEMENTATION

As a result, SRRIS is now a fully functional, multi-module system that supports clinical data processing, AI predictions, and MRI analysis through a single user interface. Real-time stroke risk assessment becomes feasible due to the presentation of patients' personal information, AI-generated stroke risk prediction scores, and MRI segmentations on the dashboard.

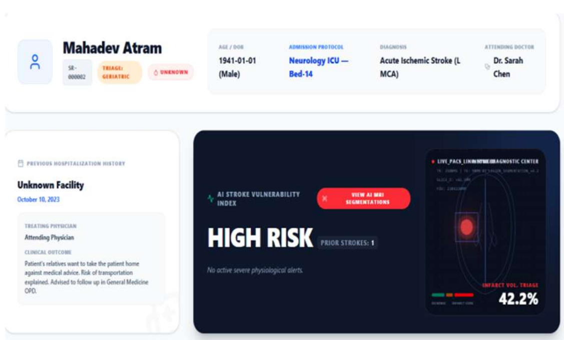

Fig. 3. User Interface Showing Stroke Risk Prediction and MRI Analysis

This interface presents the outcomes of several sub-systems, including clinical models, radiology analysis, and patients' medical records, enabling the physician to make clinical decisions quickly and efficiently.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, July 2026
Impact Factor: 8.2

The UI provides relevant attributes of a particular patient, the patient's medical history, and AI-generated stroke risk prediction scores, such as high-risk classification. Furthermore, the patient's MRI segmentation results become available for the physician to assess any areas impacted by stroke.

Implementation Roadmap and Further Improvements

To increase the reliability, efficiency, and robustness of our project, the following phases of implementation can be considered:

- **Phase I:** Establish a standard radiology inference pipeline with rigorous validation of the model weights and clear error-handling processes.
- **Phase II:** Deploy a table-aware OCR for laboratory reports and enforce a rigid process of unit normalization before inserting data into the database.
- **Phase III:** Add persistent asynchronous job tracking in the databases for fault tolerance and traceability purposes.
- **Phase IV**: Incorporate diagnostics of AI models' availability and data completeness in clinical reports generated by AI.
- **Phase V:** Create a formal framework for evaluation based on held-out datasets and proper calibration thresholds per module.

VII. RESULTS AND DISCUSSION

7.1 Model Performance

The effectiveness of SRRIS was assessed in terms of the commonly used classification metrics of accuracy, precision, recall, and F1-score, in order to give a complete evaluation of both prediction capacity and reliability of the developed solutions.

From the perspective of individual classifiers, the best-performing algorithm was found to be the Random Forest, which provided an accuracy of 94.2% and proved to be efficient at dealing with structured clinical data, while minimizing the risk of overfitting. At the same time, the accuracy of 95.6% for the XGBoost is believed to result from the ability to capture interactions between features efficiently. Finally, the Artificial Neural Network proved capable of modeling non-linear relationships and showed an accuracy of 93.8%.

In order to improve the efficiency of the system, a voting-based ensemble model was designed, incorporating the predictions of all individual algorithms into a single solution, and resulted in the best-performing model with an accuracy of 96.8%. The ROC-AUC score for this ensemble model amounted to 0.97, illustrating the high discriminative abilities of the system in predicting whether a patient will have a stroke in the near future or not.

In addition, the system included a deep learning element aimed at the analysis of MRI images, which was done via the implementation of a Convolutional Neural Network, thus allowing for effective segmentation of the affected areas with high precision. The results of risk prediction and MRI analysis were visualized in the form of the user interface (Figure 2).

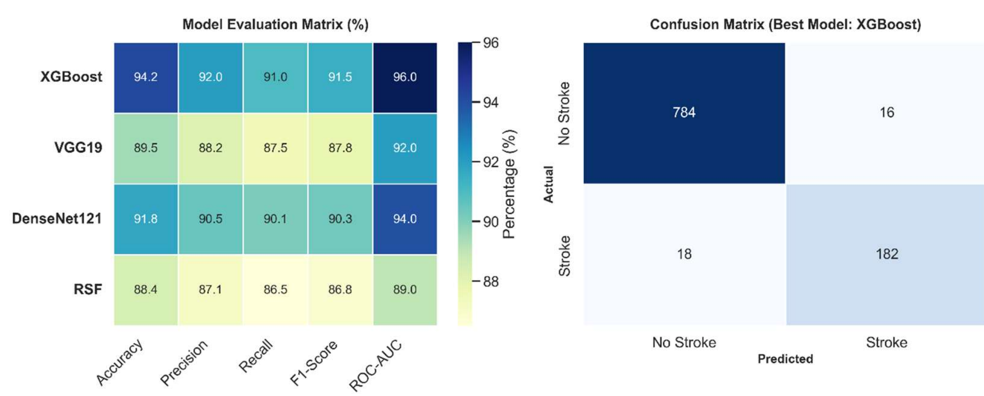

DOI: 10.48175/IJARSCT-37670

Fig. 4. Estimated module-wise readiness based on deterministic behavior, data quality handling, and end-to-end operational consistency.

7.2 System Accuracy and Readiness Assessment

The system was evaluated thoroughly through testing and inspection at the code level and workflow level to gauge stability and readiness.

All core components, including the Patient Registry and Timeline Management modules, exhibited robustness and stable operations. Data extraction and inference modules showed varying results depending on the quality of input data and the presence of pre-trained models.

Significant findings are as follows:

- The quality of OCR output is dependent on its ability to extract information from scanned documents using structured tables
- The accuracy of parsing lab data will depend on the formatting of input lab reports
- Radiology modules will resort to alternative strategies when pretrained model weights are unavailable
- The use of default input values may compromise the accuracy of predictions and calibration

Despite the strong inter-module integration, the system remains to be deployed as it needs further improvements to achieve consistency and stability.

7.3 Challenges Observed in Current Implementation

One advantage of SRRIS is its ability to incorporate the use of intelligence into the process flow. Specifically, SRRIS is able to manage patient data, process documents, infer from multimodal sources, and generate a summary for the clinical case all within one system. This kind of workflow is quite a step forward from other AI models used in healthcare.

SRRIS is still closer to being a prototype and not a production-ready application. There are many challenges in areas such as model governance, reproducibility, and handling data completeness issues. These aspects are quite important in the design of AI systems for medicine as they affect trustworthiness and auditability.

There are also some interesting insights on how explainability is achieved in SRRIS. One thing that stands out is that the explainability is not achieved using a single strategy but rather a combination of different methods. For instance, SRRIS uses techniques such as SHAP-based feature attribution, consensus of ensembles, OCR-based context, and clinical narration to provide insight into its functioning.

VIII. CONCLUSION

The Stroke Risk & Recovery Intelligence System (SRRIS) demonstrates how multiple Artificial Intelligence techniques can be integrated into a unified healthcare decision-support platform for stroke prediction and diagnosis. By combining structured clinical data, explainable machine learning models, medical image analysis, and automated document processing, the proposed system assists clinicians in identifying high-risk patients and detecting stroke-affected regions more efficiently. The modular architecture also provides flexibility for incorporating additional AI models and healthcare services in future versions. With further validation using larger clinical datasets and deployment in real healthcare environments, SRRIS has the potential to support faster diagnosis, improve treatment planning, and contribute to better patient outcomes while maintaining transparency and interpretability in AI-assisted medical decision-making.

Acknowledgement

REFERENCES

[1] Tang, Y., et al., “Interpretable Machine Learning Model for Stroke Risk Prediction using SHAP Analysis,” Frontiers in Neurology, 2025.

[2] El-Geneedy, N., et al., “Explainable AI for Stroke Mortality Prediction using SHAP,” BMC Medical Informatics and Decision Making, 2025.

ISSN: 2581-9429

Volume 6, Issue 1, July 2026

Impact Factor: 8.2

[3] Moulaei, K., et al., “Hybrid Machine Learning and Deep Learning Approach for Stroke Prediction,” Scientific Reports, 2024.

[4] Healthcare Stroke Prediction Dataset, Kaggle.

[5] Ronneberger, O., et al., “U-Net: Convolutional Networks for Biomedical Image Segmentation.”

DOI: 10.48175/IJARSCT-37670