NeuroCare: Multimodal Machine Learning and Clinical Decision Support for Migraine Prediction

Mohini Lawande[1]

[1]Department of Technology, Savitribai Phule Pune University, Pune, India

Publication Date: 2026/05/30

Abstract: Migraine is a disabling neurological condition whose symptoms, intensity, triggers, and progression vary across patients, making timely prediction difficult in routine care. This paper presents NeuroCare, a hybrid artificial intelligence framework that combines multiclass symptom-based migraine subtype classification with short-term sequential migraine risk prediction using sleep and heart rate variability data. The study draws on a structured migraine symptom dataset with 400 records and seven diagnostic categories, a sleep diary dataset containing 1,372 daily observations from 49 users, and a filtered HRV dataset with 38,913 physiological windows collected from 49 devices or users. A classical machine learning pipeline using Logistic Regression, Random Forest, and XGBoost was developed for subtype classification after preprocessing, standardization, and class balancing with SMOTE. In parallel, a 7-day LSTM framework was implemented to model temporal patterns from aligned sleep and HRV signals for next-day migraine risk estimation. The work also includes exploratory data analysis, dataset harmonization, artifact saving for deployment, and a hospital-style Tkinter-based GUI that supports patient intake, symptom capture, sequential entry, prediction review, and recommendation display. The proposed framework is intended as a practical research prototype rather than a notebook-only experiment, and it demonstrates how predictive analytics, temporal modeling, and user-centered interface design can be integrated into one clinically meaningful workflow.

Keywords: *Migraine Prediction, Heart Rate Variability, Sleep Diary, LSTM, XGBoost, Clinical Decision Support, GUI, Wearable Sensors.*

How to Cite: Mohini Lawande (2026) NeuroCare: Multimodal Machine Learning and Clinical Decision Support for Migraine Prediction. *International Journal of Innovative Science and Research Technology,* 11(5), 2465-2475. https://doi.org/10.38124/ijisrt/26may1446

I. INTRODUCTION

Migraine is a chronic neurological condition characterised by recurrent attacks of headache that may be accompanied by nausea, photophobia, phonophobia, visual aura and multiple sensory symptoms. In many hospitals and clinics, assessment still depends on a brief consultation and the patient's retrospective recall of episodes over the preceding weeks, which can lead to incomplete capture of triggers, sleep patterns and day-to-day variability in symptoms. At the same time, consumer wearables and mobile applications are making it easier to collect time-stamped sleep and physiological data outside the clinic, but these data streams rarely reach the consulting neurologist in an organised form that can be used during decision making.

Machine learning methods have shown promise for both static and temporal migraine prediction tasks. Classical classifiers, such as gradient-boosting models, can map structured symptom features to migraine subtypes or attack likelihood, while deep sequence models such as LSTMs can exploit daily time series of lifestyle or physiological signals to anticipate next-day attacks. However, many published models remain limited to research notebooks, with little emphasis on translation into the hospital workflow. There is a clear need for prototypes that connect curated datasets, trained models and a usable clinical interface in a single application so that the feasibility and acceptability of such tools can be studied in realistic settings.

This paper addresses that gap by designing and implementing NeuroCare, a hospital-style migraine screening suite that combines symptom-based classification, sequence-based next-day risk prediction and a recommendation view inside one desktop application. The work draws on multiple datasets symptom classification, sleep diary, wearable HRV and survey responses and aims to demonstrate how these can be fused into a coherent prediction and advice pipeline suitable for outpatient neurology clinics

II. LITERATURE REVIEW

Research on automated migraine prediction can broadly be divided into three strands: symptom-based diagnosis support, trigger and lifestyle modelling, and physiological signal analysis. Symptom-based decision support tools typically encode diagnostic criteria from guidelines and combine them with supervised learning to assist clinicians in

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1446

differentiating migraine with aura, migraine without aura and other primary headache types. Sequence-based approaches, on the other hand, treat daily measurements of sleep duration, stress scores, hydration or step counts as time series and use recurrent neural networks to forecast the likelihood of a headache on the following day.

A growing body of work focuses specifically on heart-rate variability as an indirect measure of autonomic nervous system balance in migraine. HRV indices such as mean heart rate, SDNN, RMSSD and frequency-domain metrics have been investigated as potential biomarkers for both interictal and peri-ictal periods. Similarly, sleep fragmentation and reduced sleep efficiency have been associated with increased migraine frequency and poorer quality of life. These findings support the clinical intuition that combining sleep and HRV features could yield an informative temporal view of each patient's risk profile.

Despite these advances, relatively few studies translate predictive models into clinician-facing applications with structured intake, dashboards and recommendation panels. Many prototypes are restricted to research notebooks or small web demos aimed at patients rather than hospitals. This project therefore positions NeuroCare as a proof-of-concept for bridging that translational gap: it uses standard machine learning methods, open-source tools and simple desktop technologies to create an integrated console that can be run in a neurology department on ordinary hardware.

III. MATERIALS AND DATASETS

➤ *Symptom Classification Dataset*

The migraine symptom classification dataset contains structured records where each row corresponds to a single consultation or episode and each column represents a clinically meaningful symptom or attribute. Features capture headache frequency, laterality, character (for example, throbbing or pressing), intensity scores, associated symptoms such as nausea, vomiting, phonophobia, photophobia, visual or sensory disturbances, and simple demographic fields such as age and sex. The target label is a multiclass migraine category, typically mapping to commonly recognised clinical subtypes.

Before model training, textual symptom descriptions are encoded into numerical categories, and consistency checks are applied to ensure that the coded values match the drop-down options presented to clinicians in the NeuroCare interface. Missing values are imputed using domain-appropriate strategies such as mode imputation for categorical fields and median imputation for ordinal severity scores. Because some migraine subtypes are less common than others, the dataset exhibits class imbalance that must be addressed during training. The resulting class distribution across the seven migraine subtypes is illustrated in Fig. 1.

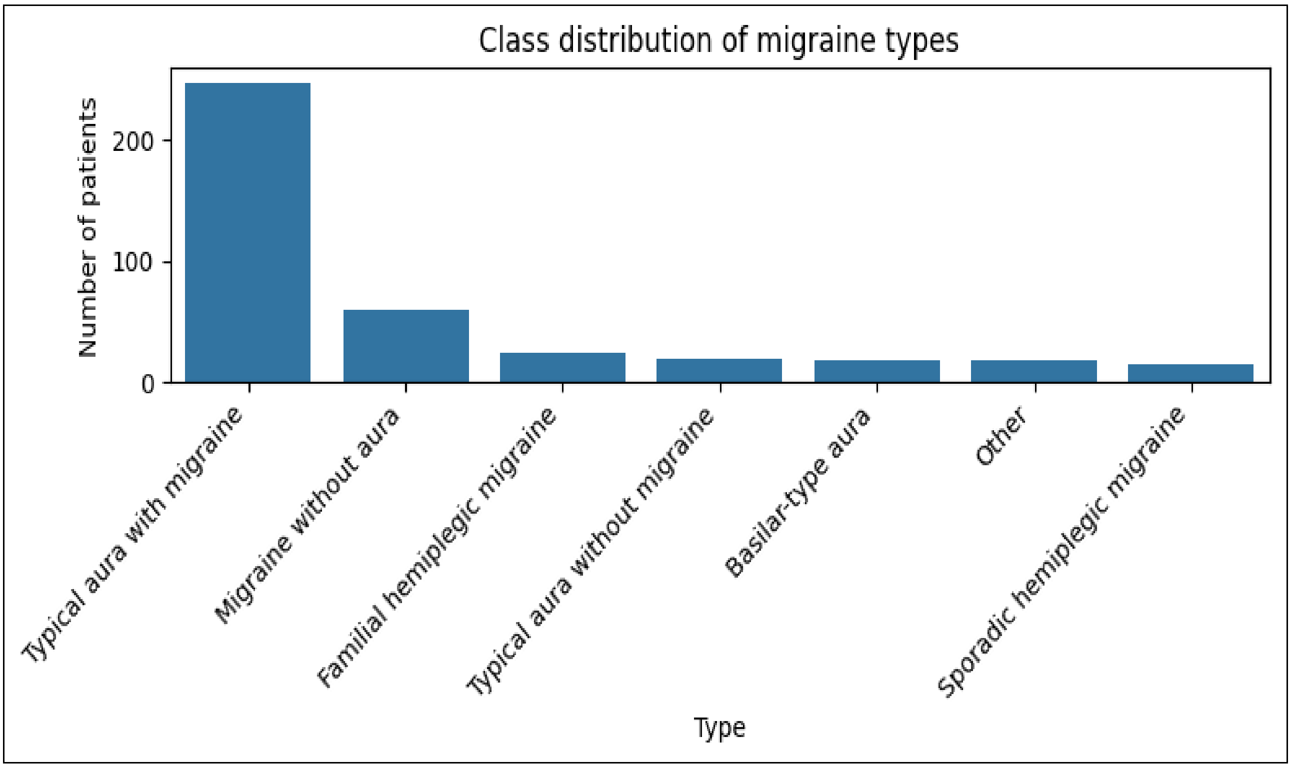

Fig 1 Class Distribution of Migraine Subtypes in the Symptom Dataset.

To characterise the population, age and pain intensity distributions were examined as shown in Fig. 2.

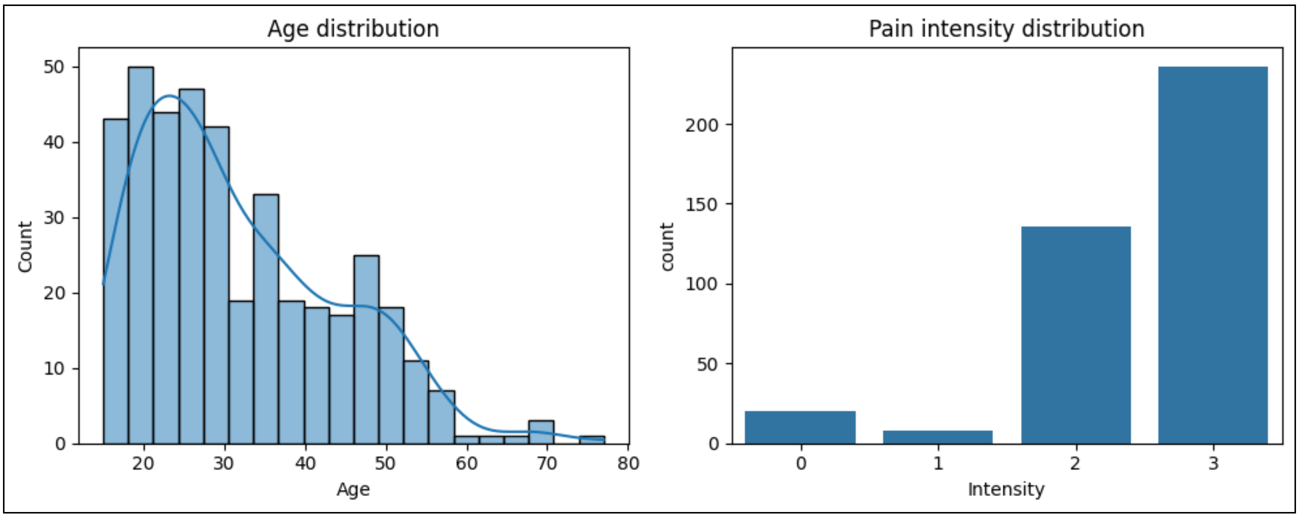

Fig 2 Age Distribution and Pain Intensity Levels for Patients in the Migraine Symptom Dataset.

The overall prevalence of individual symptoms across all records is summarised in Fig. 3, highlighting the dominance of visual aura, nausea, photophobia and phonophobia compared with rarer neurological signs.

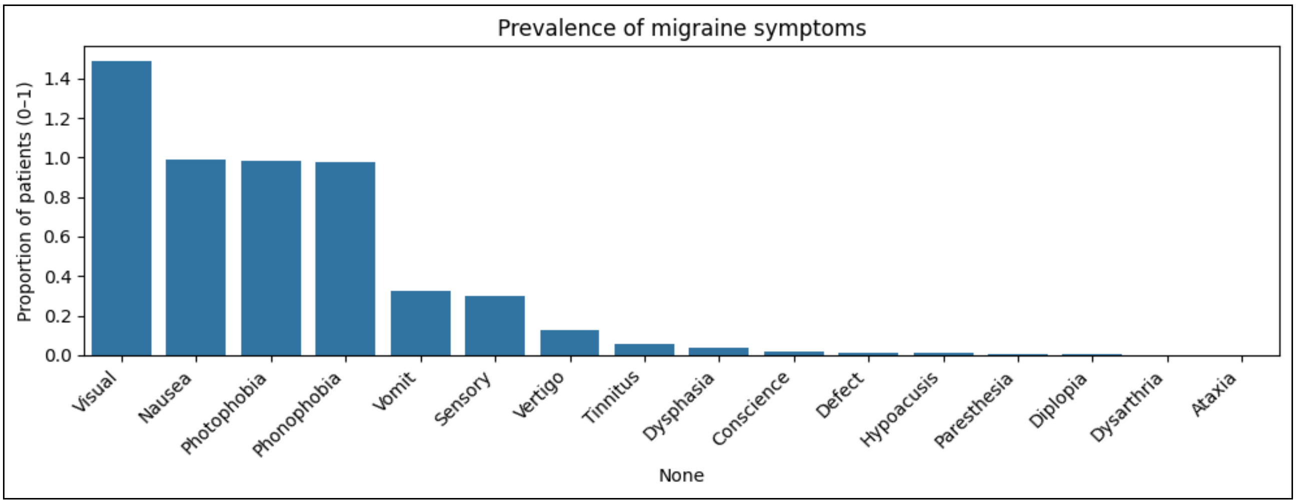

Fig 3 Prevalence of Selected Migraine-Related Symptoms Across the Dataset.

➢ *Sleep Diary Dataset*

The sleep diary dataset provides a seven-day moving window of daily entries per subject. For each day, the table includes fields such as sleep duration, time spent in bed, sleep latency, wake-after-sleep onset, calculated sleep efficiency and self-reported sleep quality. In addition, the diary may capture related lifestyle factors such as caffeine intake, hydration and screen exposure close to bedtime, depending on the specific questionnaire design.

For the LSTM model, the dataset is reorganised into fixed-length sequences of seven consecutive days, with each sequence labelled according to whether a migraine attack occurs on the following day. This formulation allows the model to learn temporal patterns, such as progressive sleep restriction or irregular bedtimes, that precede an attack.

➢ *HRV Sensor Dataset*

The sensor HRV and filtered HRV files contain wearable-derived heart-rate time series and derived variability features. Raw RR-interval streams are first cleaned to remove artefacts, outliers and ectopic beats. From the filtered signals, summary features are computed for each night or time block, such as mean heart rate, standard deviation of NN intervals (SDNN), root mean square of successive differences (RMSSD), and simple frequency-domain or non-linear descriptors where available. These nightly HRV features are then aligned with the sleep diary records to form multivariate daily feature vectors.

➢ *Survey Dataset*

A small survey dataset captures subjective information from participants or pilot users. It includes questions about migraine frequency, medication usage, perceived usefulness of prediction tools, comfort with wearable devices and preferences regarding the presentation of advice and alerts. Although the survey is not used directly for prediction, it informs the design choices of the GUI, especially the language used in the recommendation tab and the balance between clinical detail and simple explanations.

IV. METHODOLOGY

➢ *Overall System Architecture*

NeuroCare is implemented as a desktop application that wraps trained machine learning models and dataset-matched

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1446

input forms inside a hospital-style console. The workflow can be summarised in four stages:

- Data acquisition and preprocessing for the symptom, sleep diary and HRV datasets.
- Model training and evaluation for the symptom classifier and LSTM sequence predictor.
- Export of trained models and artefacts such as feature lists and encoders into files that can be loaded by the GUI.
- Graphical user interface integration, where clinicians can register patients, enter data through structured forms, invoke predictions and review advice.

The models are developed in a notebook environment using Python machine learning libraries, and the final artefacts (for example, baselinexgb.pkl and lstmdailymigraine.keras) are loaded in the GUI for real-time inference.

➢ *Symptom-Based Multiclass Classification*

For symptom classification, gradient-boosting methods are used because they handle heterogeneous tabular data, capture non-linear feature interactions and often perform well on moderate-sized clinical datasets without requiring extensive feature engineering. The preprocessing pipeline includes:

- Label encoding or one-hot encoding of categorical symptom features.
- Scaling or normalisation where required by specific models.
- Application of Synthetic Minority Oversampling Technique (SMOTE) or similar strategies to balance under-represented migraine classes.

Hyperparameters such as learning rate, maximum tree depth, number of estimators and subsampling ratios are tuned using grid search or random search with cross-validation. Model performance is evaluated using stratified k-fold cross-validation to ensure robustness across different splits. Accuracy, precision, recall, F1-score and macro-averaged metrics are reported to reflect performance on both common and rare migraine types. Confusion matrices are inspected to understand systematic misclassifications, which can later inform how the GUI presents confidence levels to clinicians.

To explore how symptom profiles differ between subtypes, symptom prevalence was computed separately for each migraine class, as shown in Fig. 4.

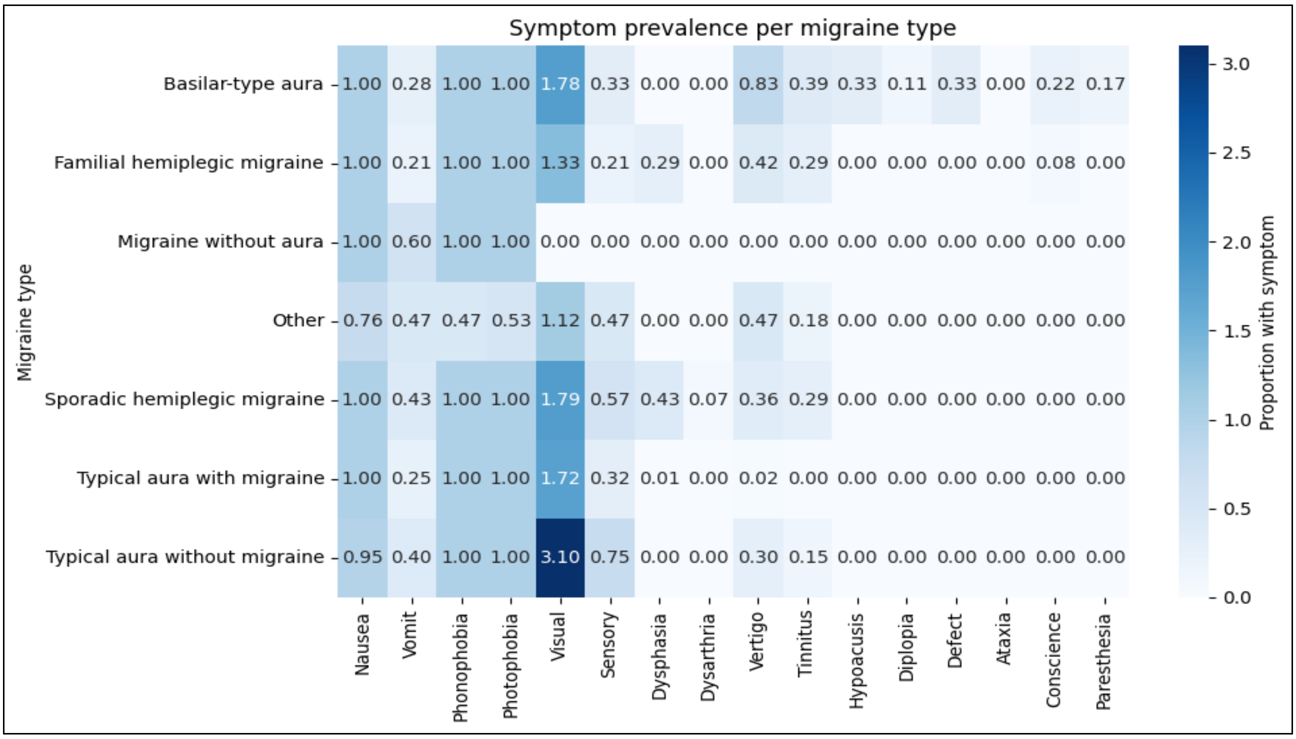

Fig 4 Symptom Prevalence for Each Migraine Subtype.

An upper-triangle correlation heatmap of the numeric features is presented in Fig. 5 to visualise linear associations and potential groups of related variables.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1446

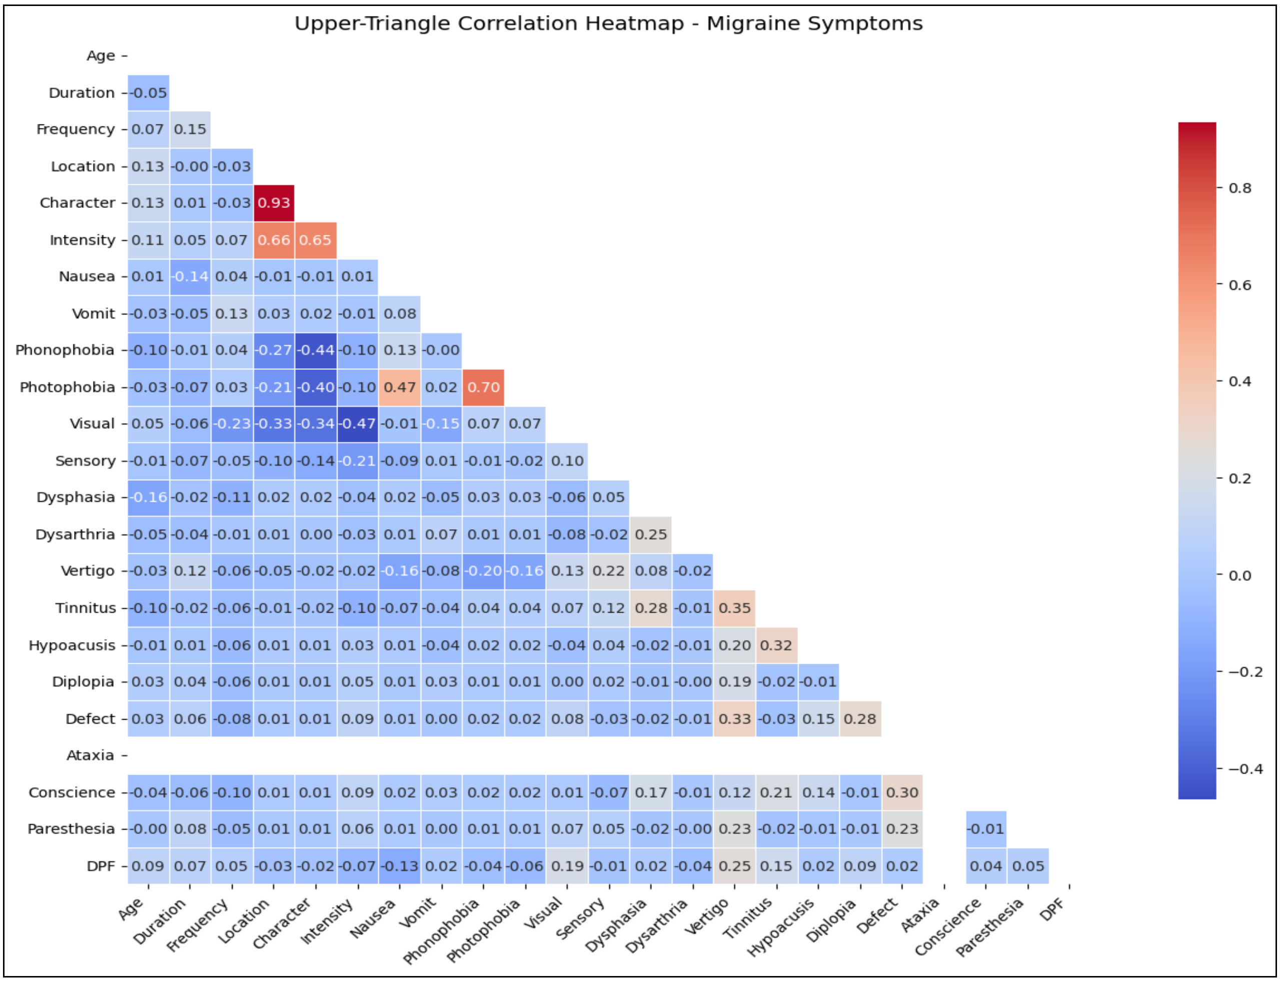

Fig 5 Upper-Triangle Correlation Heatmap for Numeric Migraine-Related Features.

➢ *LSTM-Based Next-Day Risk Prediction*

The second model is a binary sequence classifier that predicts whether a migraine attack will occur on the day following a seven-day window of sleep and HRV features. The input to the LSTM is a three-dimensional tensor with shape (number of sequences, 7, number of daily features), where each sequence contains seven days of sleep and HRV measurements. Each sequence is normalised feature-wise to stabilise training.

The network architecture typically includes:

- One or two stacked LSTM layers with an appropriate number of hidden units.
- Dropout layers to reduce overfitting.
- A dense output layer with sigmoid activation that yields a probability estimate for next-day migraine.

The model is trained using a binary cross-entropy loss function and optimised with the Adam optimiser. Early stopping on a validation set is employed to avoid overfitting, and class weights or resampling are used if the "migraine tomorrow" class is relatively rare. Evaluation uses metrics such as area under the ROC curve (AUC), precision-recall curves, sensitivity at fixed specificity, and calibration plots to assess how well the predicted probabilities align with observed frequencies.

➢ *Integration and Clinical Logic*

To make predictions clinically interpretable, outputs from both models are combined in the Prediction & Clinical Summary tab of the GUI. For a given patient and visit:

- The symptom classifier returns a ranked list of migraine subtypes with associated probabilities and a narrative summary explaining which symptom patterns led to the top prediction.
- The LSTM model outputs a risk score for next-day migraine based on the most recent seven-day window of sleep and HRV data.

Simple rule-based logic is then applied to translate numeric outputs into coloured risk labels (for example, low, moderate, high) and to trigger appropriate advice blocks in the Recommendation tab. This makes the system's behaviour transparent to clinicians while still benefiting from data-driven predictions.

V. SYSTEM IMPLEMENTATION AND HOSPITAL-STYLE GUI

➢ *User Interface Design*

The NeuroCare interface adopts a console layout familiar to hospital staff. A persistent left-hand navigation panel provides access to the main modules:

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1446

- *Dashboard*

Displays the configured number of symptom fields, the LSTM window length and the count of sequence features. It also summarises the recommended prediction workflow so that new users can follow a consistent sequence of steps. An overview of the NeuroCare dashboard layout is shown in Fig. 6.

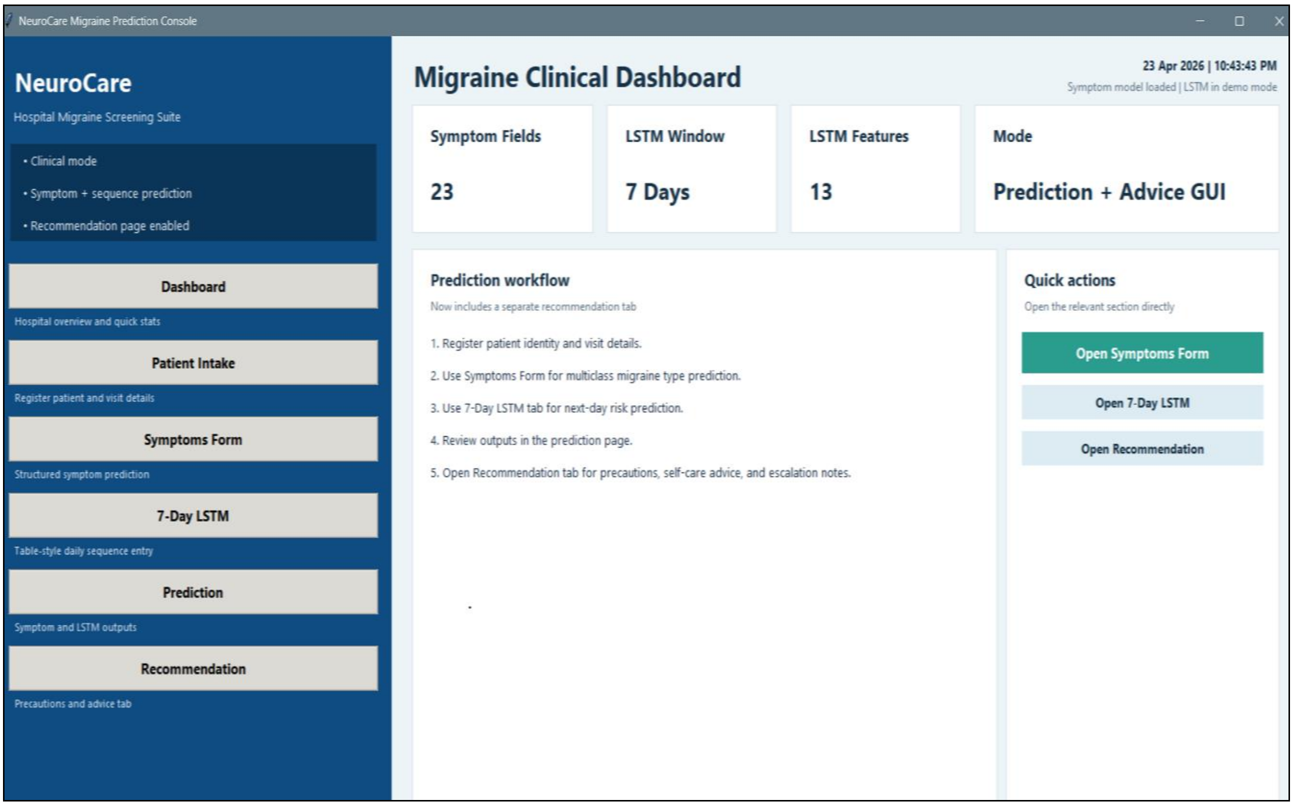

Fig 6 NeuroCare Dashboard Showing System Indicators and Navigation Panel.

- *Patient Intake*

Offers structured fields for patient ID, name, age, gender, contact information, consultant and visit type, mirroring typical neurology outpatient forms. Fig. 7 illustrates the structured Patient Intake screen used for registration."

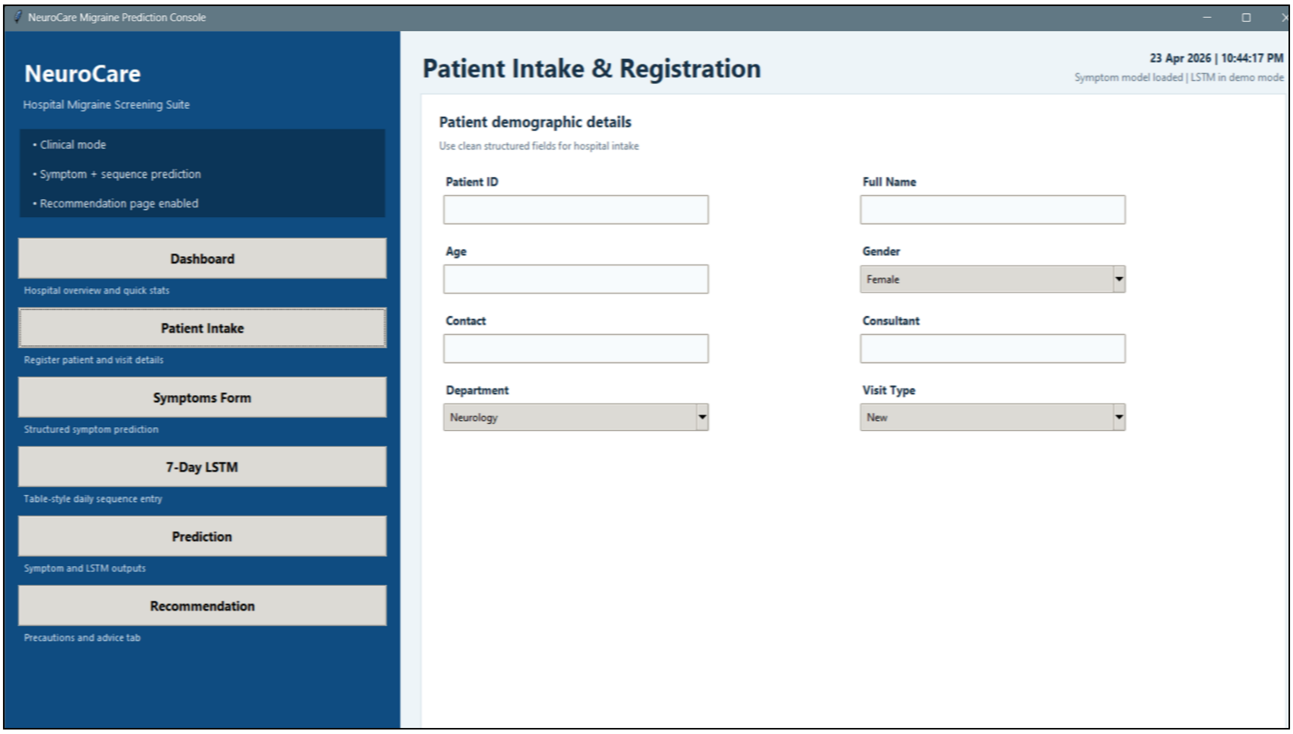

Fig 7 Patient Intake form of Capturing Demographic and Visit Details.

- *Symptoms Form*

Presents a vertically scrollable structured symptom assessment form that is tightly aligned with the coding used in the symptom dataset. Drop-downs show both the encoded value and a human-readable medical description to reduce input errors. The symptom entry interface aligned to the coded dataset is shown in Fig. 8.

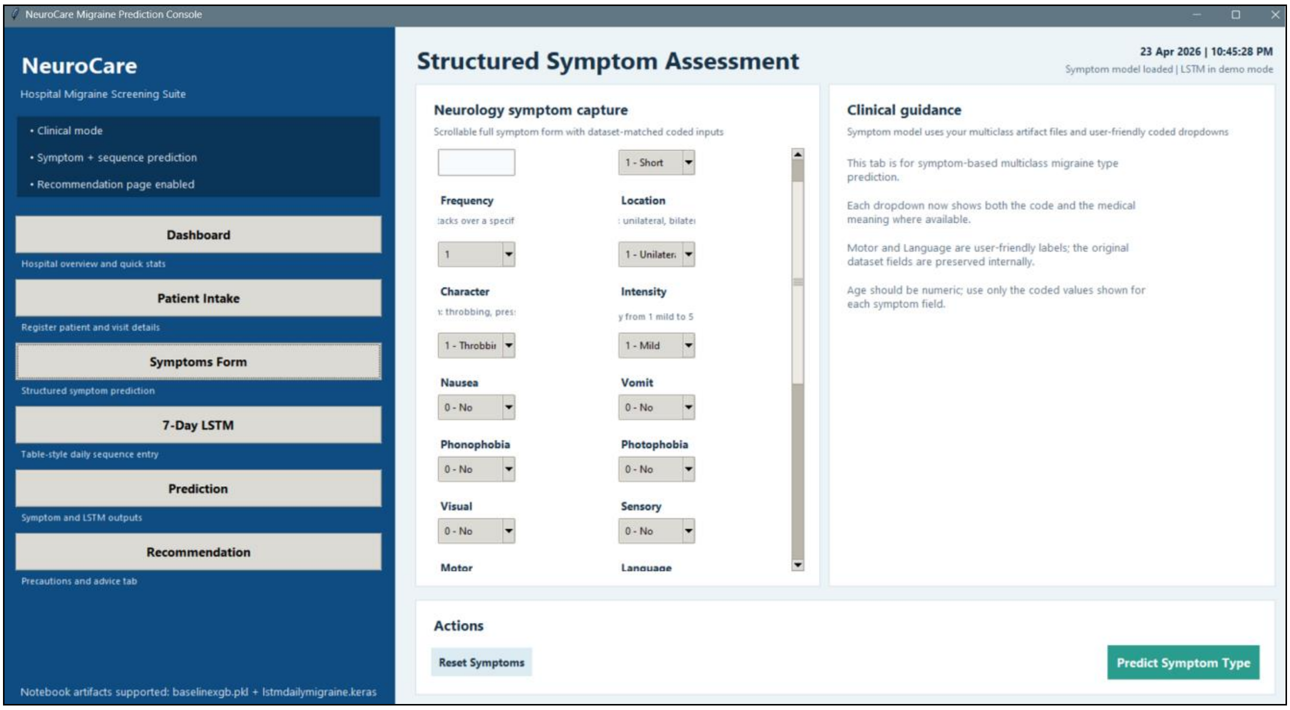

Fig. 8. Structured Symptom Assessment form with Coded Drop-Down Fields.

- *7-Day LSTM*

Provides a table-style grid where each row corresponds to a day and each column to a sleep or HRV feature. Buttons allow the user to clear the table or populate it with demo values for testing. Fig. 9 shows the 7-Day LSTM tab used to enter sleep and HRV features.

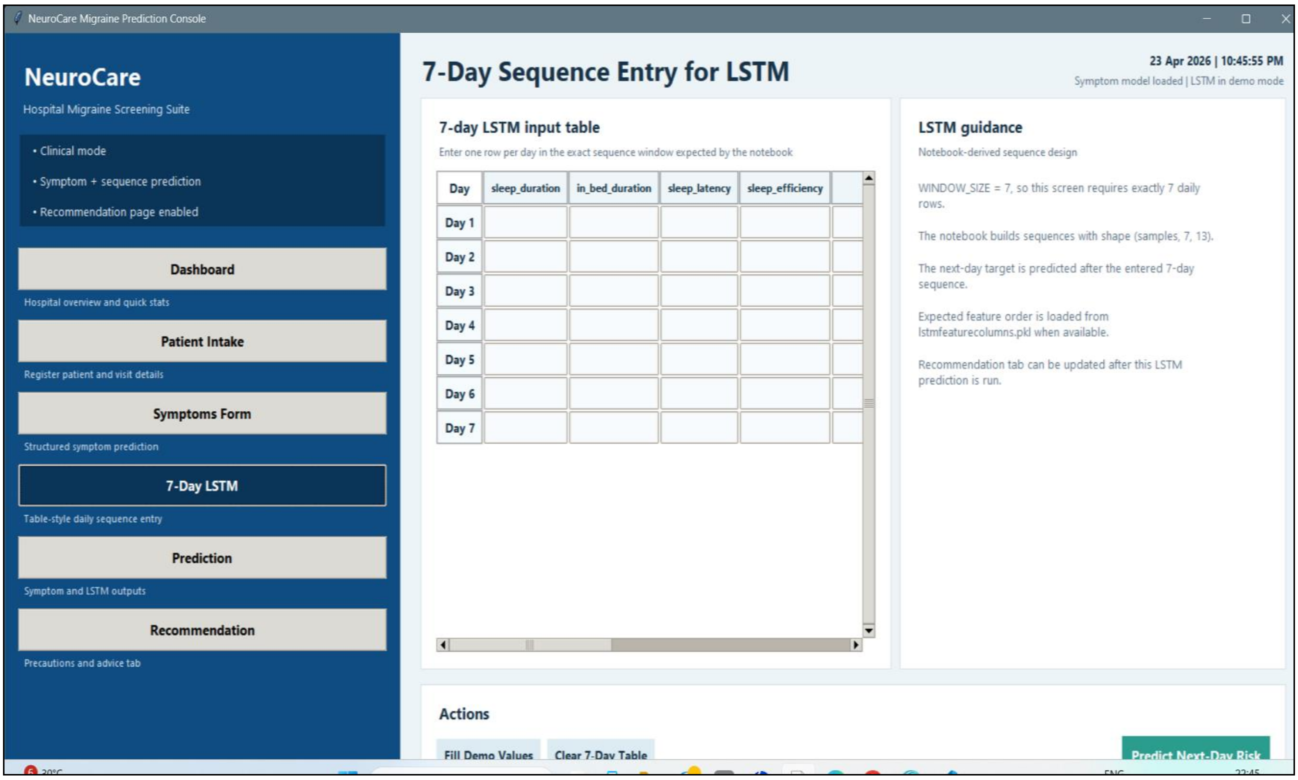

Fig 9 Seven-Day Sequence Entry Tab for Sleep and HRV Features Used by the LSTM Model.

- *Prediction*

Displays current model outputs and probability charts. The Prediction tab, which combines symptom and sequence model outputs, is depicted in Fig. 10.

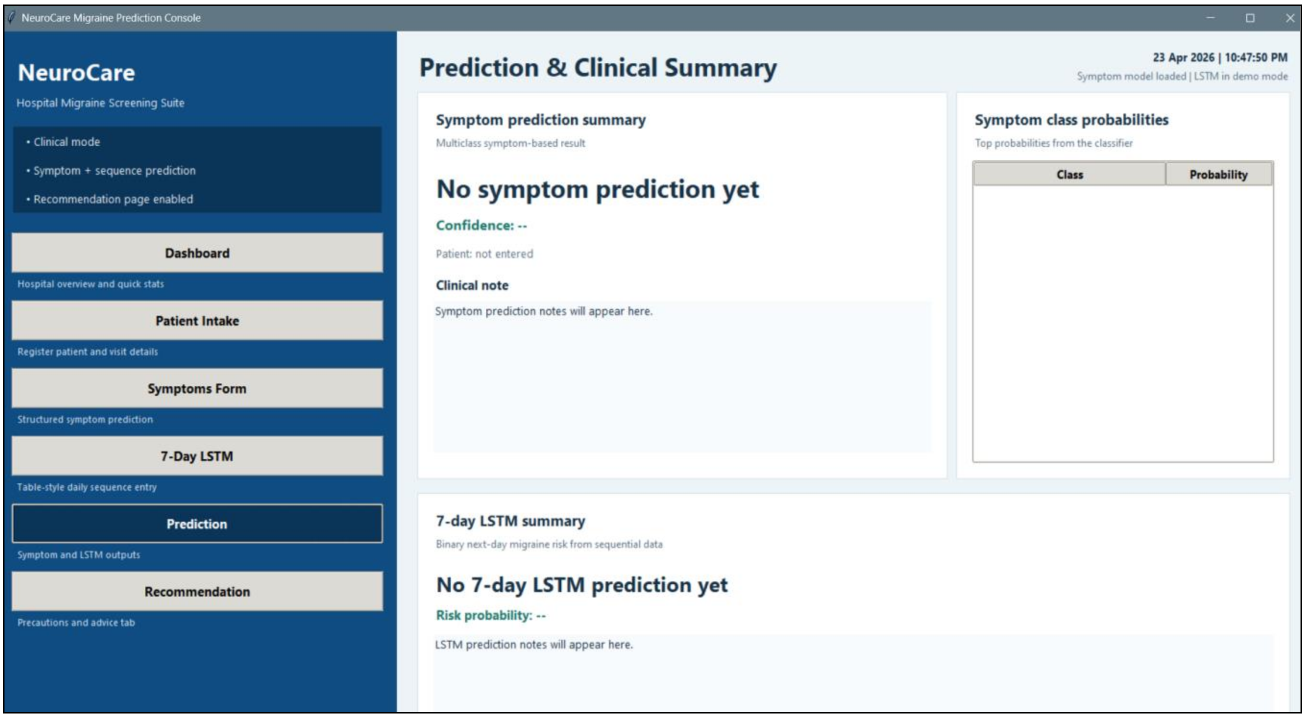

Fig 10 Prediction tab displaying migraine subtype probabilities and next-day risk score.

- *Recommendation*

Shows textual advice, precautions and escalation guidance based on the latest predictions. Fig. 11 illustrates the Recommendation tab where advice blocks are rendered based on risk.

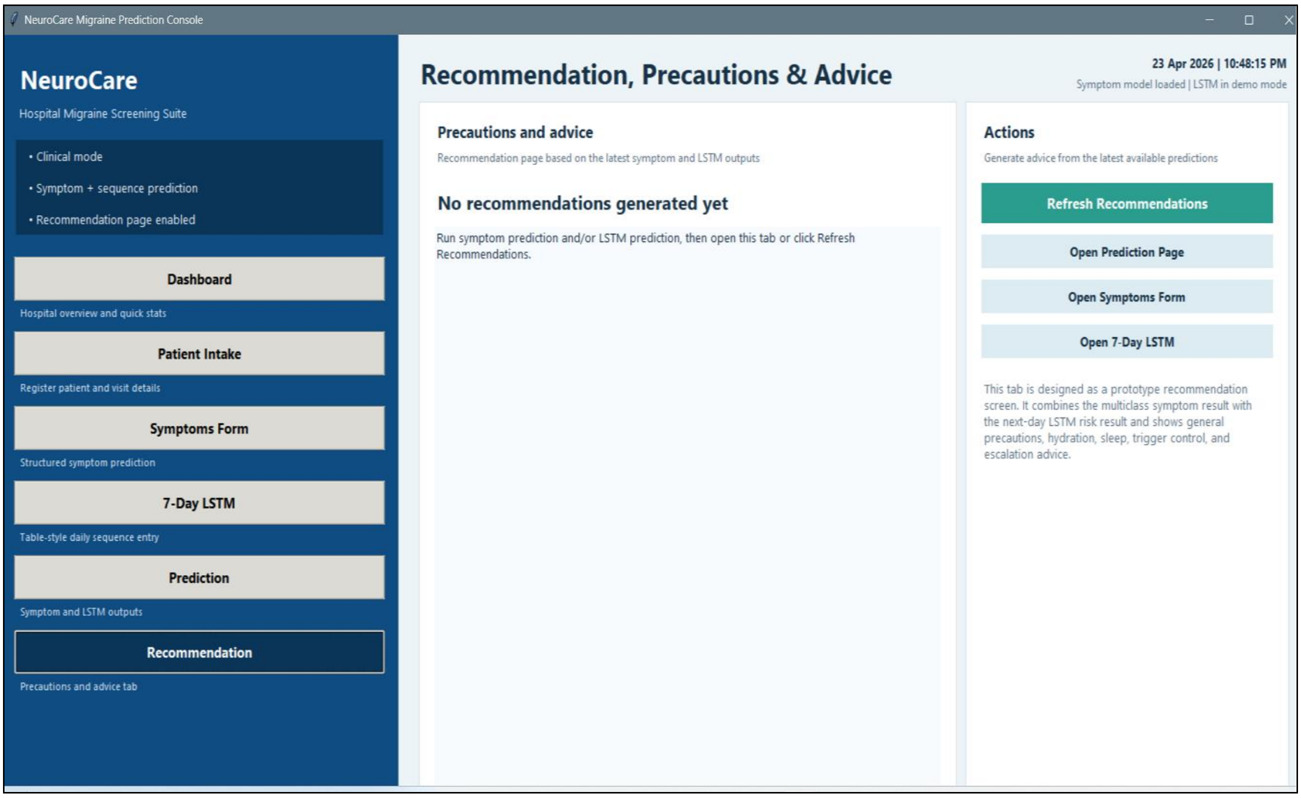

Fig 11 Recommendation tab presenting risk-based advice, precautions and escalation guidance.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1446

The window title and status bar convey that the symptom model and LSTM engine are loaded and indicate whether the system is running in demo or clinical mode. Screenshots of each module illustrate how a neurologist or nurse could use the tool during a visit, from intake to prediction to advice.

➤ *Backend Logic and Code Organisation*

The backend code orchestrates:

- Loading of trained model artefacts and feature encoders at application startup.
- Validation and preprocessing of fields entered through the GUI to match the formats expected by the models.
- Construction of a single row for symptom predictions and a properly shaped tensor for LSTM predictions.
- Error handling and user feedback when required fields are missing or inconsistent.

The code is structured into modules for data utilities, model utilities and GUI components, making it easier to maintain and extend. Logging is used to record prediction requests and allow offline audit or debugging. Where possible, internal representations preserve original dataset field names so that model explanations can reference familiar clinical terms.

VI. RESULTS AND DISCUSSION

➤ *Symptom Classifier Performance*

On the curated symptom dataset, the gradient-boosting classifier achieves strong performance on the main migraine subtypes, with high overall accuracy and competitive macro-averaged F1-scores. Misclassifications often occur between clinically similar categories, such as migraine without aura versus probable migraine, which reflects the inherent overlap in symptom profiles rather than clear model failure. The use of SMOTE or similar oversampling techniques improves recall on under-represented classes, reducing the tendency of the model to default to the majority class.

Feature importance analysis reveals that variables such as headache laterality, pulsating character, attack duration, intensity, presence of nausea or vomiting, photophobia, phonophobia and visual aura contribute most strongly to the decision boundary. These findings are consistent with standard diagnostic criteria and provide reassurance that the model is learning clinically meaningful patterns. The GUI can surface these high-impact features in the narrative summary, helping clinicians judge whether the prediction aligns with their own reasoning.

Table 1 summarises the performance of the XGBoost classifier on the curated symptom dataset. The model achieves an overall accuracy of 0.93 with a weighted F1-score of 0.92, indicating that it performs well not only on the dominant subtypes but also on the less frequent hemiplegic and basilar-type classes.

Table 1 Performance of the XGBoost Symptom Classifier on the Multiclass Migraine Dataset.

Class	Precision	Recall	F1-Score	Support
Basilar-type aura	0.75	0.75	0.75	4
Familial hemiplegic migraine	0.67	0.80	0.73	5
Migraine without aura	0.92	1.00	0.96	12
Other	1.00	0.67	0.80	3
Sporadic hemiplegic migraine	1.00	0.67	0.80	3
Typical aura with migraine	0.96	0.96	0.96	49
Typical aura without migraine	1.00	1.00	1.00	4
Overall / averages				
Accuracy	–	–	0.93	80
Macro avg	0.90	0.83	0.86	80
Weighted avg	0.93	0.93	0.92	80

➤ *LSTM Prediction Performance*

The LSTM sequence model demonstrates the ability to discriminate between high-risk and low-risk seven-day windows, capturing temporal patterns that are not evident from single-day summaries. For example, gradual reductions in sleep duration, irregular bedtimes or sustained elevations in night-time heart rate may progressively raise the predicted risk. AUC values above chance level, together with acceptable sensitivity-specificity trade-offs, indicate that the model has learned non-trivial structure from the sleep and HRV sequences.

However, performance varies depending on the amount of training data and the quality of sensor recordings. Artefacts in HRV measurements, missed diary entries and heterogeneous participant behaviour can reduce signal-to-noise ratio. The prototype therefore emphasises careful preprocessing and transparent display of model confidence, encouraging clinicians to treat the risk score as one input among many rather than as a definitive prediction.

Because the current LSTM branch is trained on simulated daily labels, the goal is to evaluate feasibility rather than to claim definitive clinical performance. As reported in Table 2, the model attains an overall accuracy of around 0.70 with an AUC of 0.57, and it shows reasonable discrimination for the majority "no migraine" class while struggling on the much rarer "migraine" days.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1446

Table 2 LSTM Daily Migraine Risk Model Performance on Simulated Seven-Day Sequences.

Class	Precision	Recall	F1-Score	Support
No migraine	0.91	0.73	0.81	207
Migraine	0.14	0.38	0.20	24
Overall / averages				
Accuracy	–	–	0.70	231
Macro avg	0.53	0.55	0.51	231
Weighted avg	0.83	0.70	0.75	231
AUC	–	–	0.57	–
Average precision (PR)	–	–	0.18	–

These values underline that the sequential model can capture some useful structure from the merged sleep and HRV data, but its recall and F1-score for migraine-positive days remain modest. This is expected, given the rarity of positive labels and the fact that the experiment uses simulated outcomes. In later work, the same architecture can be retrained on real daily attack labels to obtain clinically meaningful operating points.

➢ *Usability and Clinical Interpretation*

The hospital-style GUI aims to make advanced models approachable for busy clinicians. By embedding predictions in familiar clinical forms and dashboards, NeuroCare shifts the emphasis from algorithms to workflow. The step-wise navigation (Intake → Symptoms → 7-Day LSTM → Prediction → Recommendation) mirrors how a typical consultation unfolds and can be explained easily to both clinicians and patients.

The Recommendation tab illustrates how numeric outputs can be translated into grouped advice blocks, for example: hydration and regular meals, sleep hygiene, trigger avoidance, early use of prescribed medications and escalation criteria for emergency care. These text blocks are parameterised by risk category and symptom profile, allowing the system to provide personalised but structured guidance. Pilot feedback from the survey data indicates that such integrated advice, presented in clear language, is more acceptable than raw probabilities alone.

At the same time, the current prototype operates in a demonstration mode and is not yet validated in real clinical populations. External validation on larger, diverse cohorts, as well as human-factors studies with neurologists and nurses, would be necessary before deployment in routine care.

VII. LIMITATIONS AND FUTURE WORK

The study has several limitations. First, the datasets used are of modest size and may not fully capture the diversity of migraine presentations across age groups, comorbidities and cultural backgrounds. Second, the LSTM model is trained on aggregated sleep and HRV summaries rather than raw high-resolution signals, which may omit subtle temporal patterns. Third, the current system focuses on prediction and advice, but does not yet incorporate longitudinal tracking of individual patients across multiple visits or automatic integration with electronic medical record systems.

Future work will therefore focus on expanding the datasets, including continuous wearable streams and multi-site clinical data; experimenting with alternative sequence architectures such as temporal convolutional networks or transformers; and adding interpretability modules, for example, attention-based visualisations that highlight influential days within the seven-day window. In parallel, usability studies with clinicians and patients will be conducted to refine the wording of advice, the thresholds for risk categories and the overall user experience.

VIII. CONCLUSION

This paper presents NeuroCare, an integrated migraine prediction and clinical advice system that combines symptom-based gradient-boosting classification with LSTM-based next-day risk prediction from sleep diary and wearable HRV data, all wrapped inside a hospital-style GUI.

By unifying multiple datasets, machine learning models and a clinician-friendly interface, the system demonstrates a practical pathway from raw patient-generated data to actionable bedside insights. Although further validation is required before clinical deployment, the prototype highlights how standard machine learning tools can be configured into a coherent decision support workflow for migraine care and provides a template that can be extended to other neurological conditions.

ACKNOWLEDGMENT

The authors would like to thank the faculty mentors and all who provided guidance on migraine pathophysiology, data interpretation and user-interface feedback during the development of this project. Any datasets or survey responses used in this work were collected and processed in accordance with institutional guidelines and ethical considerations.

REFERENCES

[1]. Quartetti U. et al, ”Forecasting migraine attacks by managing daily lifestyle: a systematic review as a basis to develop predictive algorithms”, Pain Reports,10(2), e1247, 2025.
[2]. Das S. et al., ”Machine Learning-based Migraine Prediction: Analyzing Key Features and Cause-Effect Relationships for Improved Diagnosis and Management”, International Journal of Computer Applications ,Volume 187 - no.11, June 2025.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26may1446

[3]. Amiri P. et al., ”Migraine: A Review on Its History, Global Epidemiology, Risk Factors, and Comorbidities”, Frontiers in neurology, volume 12-2021, February 2022.
[4]. Gulati S. et al. , ”Classification of Migraine Disease using Supervised Machine Learning”, 10th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO)., 2022.
[5]. Jalannavar A. et al., ”Migraine Prediction Using Deep Learning Model”, Fourth International Conference on Emerging Research in Electronics, Computer Science and Technology (ICERECT),2022.
[6]. Srinivasan S. et al. , ”Predictive Analysis of Chronic Migraine Symptoms Detection and Management Using IoT and ML”, 2025 International Conference on Advances in Modern Age Technologies for Health and Engineering Science (AMATHE), 2025.
[7]. Velgina R. et al., "Deciphering Migraine Types: A Machine Learning Odyssey for Precision Prediction", 2024 2nd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT), 2024.
[8]. Chokri Baccouch. et al., "Advanced Machine Learning Approaches for Accurate Migraine Prediction and Classification" (IJACSA) International Journal of Advanced Computer Science and Applications, 2025.
[9]. Khan, Lal, et al. "Migraine headache (MH) classification using machine learning methods with data augmentation." Scientific Reports, volume 14, no. 1, 2 Mar. 2024, p. 5180. Nature.