IAES International Journal of Artificial Intelligence (IJ-AI)
Vol. 15, No. 6, December 2026, pp. 5788∼5798
ISSN: 2252-8938, DOI: 10.11591/ijai.v15.i6.pp5788-5798

Magnetic resonance imaging recommendation using electromyography and medical text on edge devices

Prashant Ronad[1], Eshaan Mathur[1], Anirudh Suresh[1], Abhishek Manja Naik[2], Sudarshan Mahantesh Ullegaddi[3], Jayati Vachhani[4], Jyoti Shetty[1]
[1]Department of Computer Science and Engineering, RV College of Engineering, Bengaluru, India
[2]Department of Electronics and Communication Engineering, RV College of Engineering, Bengaluru, India
[3]Department of Electrical and Electronics Engineering, RV College of Engineering, Bengaluru, India
[4]Department of Aerospace Engineering, RV College of Engineering, Bengaluru, India

Article Info

Article history:

Received Jul 18, 2025
Revised Aug 8, 2026
Accepted Aug 31, 2026

Keywords:

Clinical decision support systems
Edge computing
Electromyography
Graph neural networks
Magnetic resonance imaging
Medical named entity recognition
Multimodal learning

ABSTRACT

Healthcare faces persistent challenges in managing unstructured clinical information, interpreting physiological signals, and ensuring timely diagnostic decisions, compounded by inconsistent medical terminology and limited diagnostic capabilities. This paper introduces a multimodal diagnostic system combining unstructured text and electromyography (EMG) signals to improve neurological condition detection, using two parallel inference models: a natural language processing (NLP) pipeline that extracts key symptoms via a bidirectional long short-term memory-conditional random field (BiLSTM-CRF) model, and a real-time biosignal analyzer. Experiments show the system's robustness, achieving an area under the curve (AUC) of 0.99 for the EMG classifier, while the BiLSTM-CRF entity extractor combined with a graph-based SortPooling model achieved 76.35% classification accuracy and 76.10% F1-score. By leveraging textual and biosignal data within a unified, lightweight architecture, the system shows strong potential for scalable, edge-deployable, context-aware healthcare solutions, particularly where specialist access is limited. Unlike prior works focused mainly on disease classification, this study addresses the novel problem of automatically predicting the need for magnetic resonance imaging (MRI). Simultaneously identifying semantic symptom correlations and physiological patterns allows the framework to overcome the limitations of traditional unimodal systems while remaining deployable at edge.

This is an open access article under the CC BY-SA license.

Corresponding Author:

Jyoti Shetty
Department of Computer Science and Engineering, RV College of Engineering
R.V Vidyanikethan Post, Mysuru Road, Bengaluru, Karnataka 560059, India
Email: jyothis@rvce.edu.in

1. INTRODUCTION

The healthcare industry faces challenges in managing heterogeneous patient data, as unstructured clinical text comprises roughly 80% of healthcare information [1]. Natural language processing (NLP) has improved clinical entity extraction through transformer-based and hybrid architectures, and a prior NLP-and-graph pipeline introduced by this group predicts the magnetic resonance imaging (MRI) requirement using a graph neural network [2]; the present study extends that pipeline with a physiological modality. Textual data alone, however, lacks physiological evidence of neuromuscular dysfunction. Biosignals such

as electromyography (EMG) provide direct, quantitative insight into neuromuscular activity and are important in diagnosing neuropathies, myopathies, motor neuron diseases, and radiculopathies, with both surface and intramuscular EMG used to assess activation patterns, conduction velocity, recruitment, and fatigue.

When electronic health records (EHRs) are combined with medical imaging using a multimodal learning approach, it has shown promising results in improving clinical decision-making [3], motivating the combination of textual and biosignal data. Time-, frequency-, and time–frequency-domain features characterize neuromuscular behavior [4], EMG serves as a biomarker for conditions such as Parkinson's disease [5], and deep architectures (convolutional neural networks (CNNs), recurrent neural networks (RNNs), and attention models) have improved fatigue detection, movement classification, and pathology identification [6], [7]. Importantly, quantitative MRI distinguishes EMG severity grades in neuropathy [8] and, in combination with EMG, predicts surgery outcomes [9]; however, cost remains a barrier, with brain-MRI prices ranging from 1,024 to 3,197 (median of 1,900) [10].

In spite of these advances, MRI referrals have become inefficient processes, dependent on arbitrary judgments made by clinicians, which can vary from one to another that is not consistent and resource-intensive in under-resourced environments. There is a clinical gap that exists in modality integration which is that clinical texts hold important historical information but lack systematic analysis, while EMG offers objective evidence yet is rarely interpreted alongside textual data, such as EHRs, and current implementations treat these modalities in isolation without a combined implementation for real-time MRI requirement decisions at the appropriate point of care.

This work is focused on that underexplored issue by proposing a lightweight multimodal system that can process clinical text with a deep learning sequential model to extract symptoms and organs [11], construct a knowledge graph, and predict the outcome with a graph neural network [12], while a one-dimensional convolutional neural network (1D-CNN) analyzes the real-time EMG signals. This work's primary contribution lies in integrating these modalities into a resource-constrained edge platform. Predictions from the text and EMG models are combined using a linear opinion pool to generate a single recommendation, reducing the limitations of relying on either modality alone. This enables an affordable, real-time neurological screening approach that supports clinical assessment with physiological evidence and is well suited for both hospitals and remote healthcare settings.

2. METHOD

The methodology links the informal medical narratives to structured clinical reasoning. The system is then expanded to incorporate a multimodal framework, which involves physiological signal processing using EMG and fusion of multiple inference paths via linear opinion pool. Two separate models in the system produce an independent inference. The first one makes use of medical text and the second one makes use of EMG signals.

For the purpose of obtaining a real-time outcome in a resource-lacking environment, EMG sensors attached to an ESP32 functions as edge devices. They collect EMG signals from the individual and conduct minimal pre-processing and transfer these signals over a network to a Raspberry Pi. The Raspberry Pi here executes the EMG model, thereby acting as a intermediary for inferencing the data and also functions as a central server to execute the named entity recognition (NER)-based model. The fusion of the models uses a linear opinion pool as shown in (1).

$$P_{final}(c) = w_1 P_1(c) + w_2 P_2(c) \tag{1}$$

Where $P_1(c)$ and $P_2(c)$ are the class probabilities predicted by the EMG model and the NLP-imaging model respectively, and w_1, w_2 are weights such that $w_1 + w_2 = 1$ and $w_1, w_2 \in [0, 1]$. These weights can be assigned in a static manner using the cross-validation process or can be dynamically modified according to the confidence metrics. This multimodal approach enhances the accuracy of diagnosis by combining physiological data with textual and EMG information, leading to a more thorough patient assessment procedure.

2.1. Data collection and preparation

This research employs two modalities, which are clinical text and EMG signals. Clinical data were extracted from multiple healthcare institutions in Bengaluru, India along with the PMC-patients dataset, including discharge summary information for around 167 thousand patients [2], [13]. With respect to the

preceding text pipeline [2], sentence segmentation and tokenization along with NER were performed on each narrative and contextual embeddings obtained using biomedical bidirectional encoder representations from transformers (BioBERT). A total of 3,322 5-second signal samples were then collected with BioAmp EMG candy in regulated environment. Preprocessing consisted only of basic noise filtering and the use of normalization preventing changes to data that would reduce clinically meaningful variability, after which all signals were extracted in fixed length segments for time-series analysis with respect to a 1D-CNN run on a Raspberry Pi for real-time inference. Figure 1 shows the architecture of the system, that is composed of the following major components: extraction of clinical entities using NER, EMG signal acquisition, and outcome generation with a 1D-CNN model.

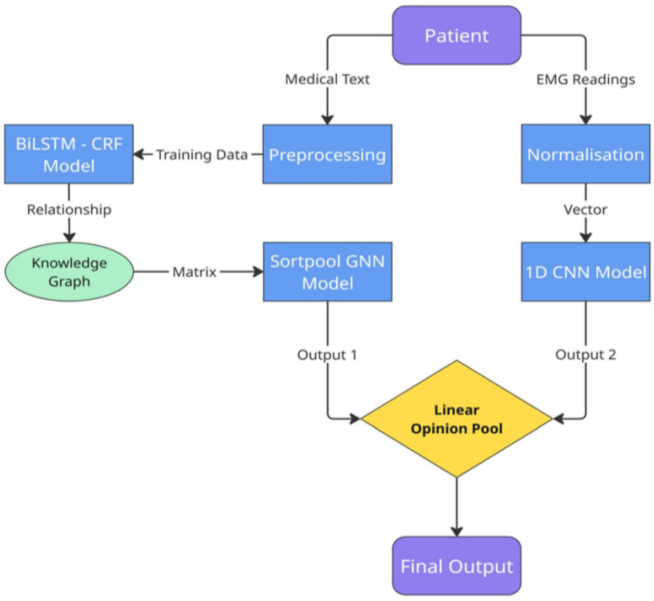

Figure 1. Proposed system workflow

2.2. Medical entity recognition and extraction

The initial step in the extraction of clinical features is to build a dataset which is annotated and could assist to train an NER-based model. Medical NER has been widely studied and various approaches have been proposed including dictionary-based approaches, transfer learning, and conditional random fields models [14]–[16], BERT-based and hybrid deep learning models [17]–[21], and hybrid architectures combining transformer-based models with bidirectional long short-term memory - conditional random field (BiLSTM-CRF) and graph-based neural networks [22]. Consistent with the prior pipeline [2], first, the dataset was annotated automatically with a pre-trained model, biomedical NER. Since automatic annotation is prone to errors, the first annotations were manually checked and corrected to improve the annotation accuracy. For this purpose, a six-category tagging scheme was developed where non-medical tokens that have no clinical significance were tagged as label 0, label 1 and label 2 for initiating and continuing tokens of a symptom range, label 3 for temporal duration of a symptom, label 4 for patient gender, and label 5 for the involved anatomical location.

In the last phase of the extraction, the BiLSTM-CRF architecture was selected over the alternatives consisting of only a graph convolutional network (GCN) or a conditional random field (CRF), since it had shown better performance and is generally well suited to sequence labeling tasks. The network consists of an embedding layer based on BioBERT which produces 768-dimensional representations of the tokens. These embeddings are then fed into a two-layer bidirectional long short-term memory (BiLSTM) having a hidden state dimension, which is of 256, applying 30% dropout, after which a linear layer reduces the 256-dimensional output to six classes. A final CRF layer is used to ensure valid transitions between consecutive tags, thus keeping the predicted sequence of labels structurally consistent.

2.3. Graph-based processing framework

2.3.1. Relationship mapping between symptoms and anatomical structures

Although anatomical structures (or organs) and symptoms were separately obtained, modeling them in a graphical structure required relationship establishment between these entities. To accomplish this, an algorithm based on proximity was implemented for identifying the potential organ matching every symptom. A probabilistic value of an anatomical structure connected to a definite symptom was determined after that with the help of naive Bayes. The probabilistic values were calculated from co-occurrences of anatomical structures and a well validated symptom-to-structure mapping, that were derived while training the model.

More precisely, $P(S_1 \mid O_1)$, where O_1 represents an anatomical structure and S_1 denotes a symptom, was calculated. This process generated a definite symptom-to-structure mapping which could help in graph construction. Given that a single symptom can be linked to multiple anatomical structures, the anatomical structure associated with the highest probability under the naive Bayes model is chosen as the predicted selection. The Llama 8b 8192 model was used to generate a set of validated symptom-to-structure mappings. Frequency counts drawn from these mappings, together with information contained in the patient's current medical description, formed the basis for these probability values.

2.3.2. Knowledge graph construction

Once the important medical features pertaining to each individual have been identified, namely symptoms, how long they last (duration), and the organ/anatomical structure that is affected, these medical features are still interrelated and not independent. For instance, each of the symptoms can arise within a specified period of time and have an effect on one or more organs. Moreover, several such entity collections can be linked to every single patient, which can represent complicated relationships. A knowledge graph is therefore constructed to model this more efficiently (Figure 2). This graph-oriented method focuses on an area in which the creation of a graph for each patient has received little attention in recent literature. This knowledge graph facilitates the proper handling of complex relationships between the identified medical concepts and avoids redundancy with the help of denoised graphs, which is consistent with the fundamental theoretical models of automatic creation of health knowledge graphs from electronic medical records [23].

As an illustrative case, patient P_1 produces the following set of extracted relationships: i) S_1, which is a symptom, lasting for D_1 duration, linked to anatomical structure O_1; ii) S_2, which is a symptom, lasting for D_2 duration, linked to anatomical structure O_2; iii) S_3, which is a symptom, lasting for D_3 duration, linked to anatomical structure O_1; and iv) S_4, which is a symptom, lasting for D_2 duration, linked to anatomical structure O_2. Analysis of these relationships reveals several important patterns that must be accounted for during graph construction. First, a single anatomical structure (O_1 or O_2) can be implicated by more than one symptom (for example, (S_1) and (S_3) both link to (O_1), while (S_2) and (S_4) both link to (O_2), indicating that organs frequently serve as shared points of convergence across multiple symptoms. Second, distinct symptoms may share an identical duration value (as seen with (D_2) for both (S_2) and (S_4), showing that temporal attributes are not necessarily unique identifiers of a single symptom. Finally, not every symptom mention in the extracted data necessarily comes with all three associated entities intact, since cases may arise where either the duration or the anatomical structure is missing from the source text.

Based on these observations, a patient-level graph is constructed in which the nodes stand for the extracted elements– symptoms, duration, and anatomical structures. In this graph, the edges signify how these elements are related, such as "Symptom affecting an anatomical element" and "Symptom existing for a duration". Each patient is therefore represented by an individual graph whose topology directly reflects the symptom-duration-organ associations extracted from that patient's clinical narrative, allowing shared organs and durations across multiple symptoms to be captured naturally through overlapping edges.

Incomplete information in unstructured text is an important consideration when building this representation, since clinical narratives do not always report every attribute of a symptom. Symptoms do not have to indicate a period of time or an organ, and in many cases one or both of these attributes may simply be absent from the source text. To accommodate this, the graph structure allows for incompleteness by omitting the corresponding node while maintaining all other valid connections, ensuring that partially specified symptoms still contribute usable structural information to the overall graph.

2.3.3. Graph machine learning

The architecture begins with two GCNConv layers, each followed by rectified linear unit (ReLU), progressively narrowing the feature space from an initial size of 768 down to 32. Because graphs in the dataset

vary in size, a SortPooling operation (with value of k as 5) standardizes the output into a fixed 5 × 32 matrix before further processing. A 1D convolution (kernel size 2, 32 filters) follows, feeding into a final dense layer through a sigmoid nonlinearity that outputs the classification outcome. For optimization, Adam was chosen where the learning rate was set to 1e-3, and 250 training cycles (or epochs) were used for training while loss was tracked after each one.

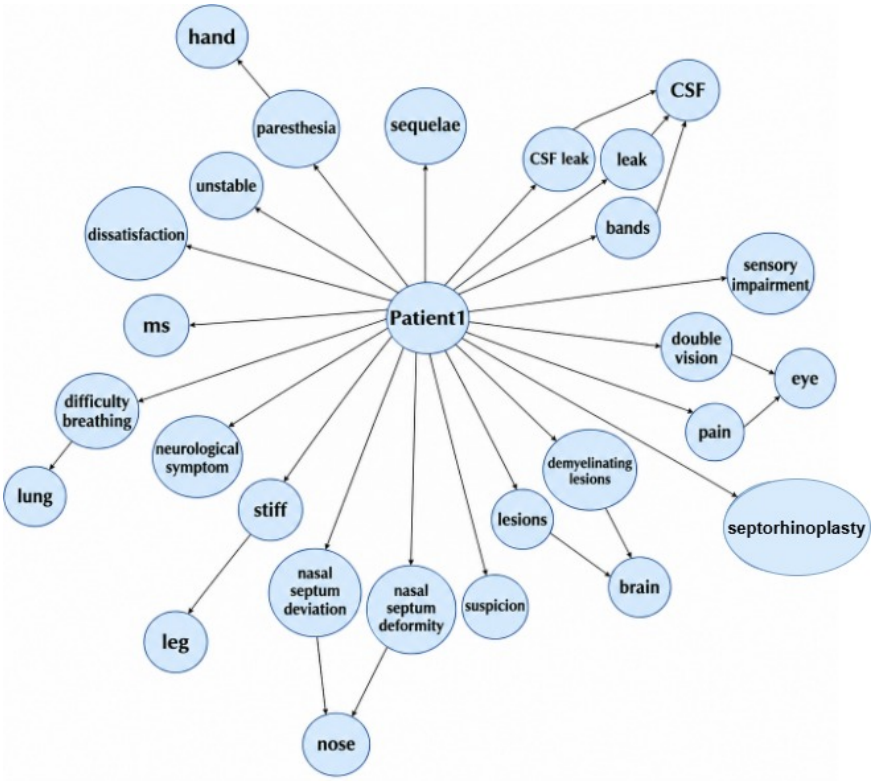

Figure 2. Example graph for symptoms and organs network

2.4. Electromyography classification

This module executes instantaneous identification of neuromuscular irregularities through a streamlined, edge-centric framework that combines economical sensing devices with integrated processing, encompassing EMG signal collection, wireless communication, preprocessing, deep learning-driven categorization, and result distribution. The hardware is composed of three main parts, a BioAmp EMG candy sensor, an ESP32 microcontroller and a Raspberry Pi 5. The BioAmp EMG Candy serves as the analog front-end for surface EMG acquisition, providing high input impedance, low noise amplification, and basic analog filtering to minimize motion artifacts and power-line interference. The surface electrodes are placed over a given muscle group and pick up the electrical impulses which are produced during muscle contraction. The continuous analog EMG signal is directly fed into the ESP32's integrated analog to digital converter (ADC), which converts the biosignal into discrete digital samples; the ESP32 operates as a low-power, Wi-Fi enabled acquisition node that supports local data buffering and wireless transmission. In order to allow near real-time analysis with less overhead, the ESP32 partitions the continuous data flow into fixed 5-second intervals, striking a balance between temporal resolution and classification reliability. Each interval is serialized into an HTTP POST request sent over a local wireless network to an HTTP server running on the Raspberry Pi 5, offloading all intensive processing and inference tasks to the Raspberry Pi, thus preserving the ESP32's computational efficiency.

Upon receiving the data, the Raspberry Pi 5 reshapes the signal into a fixed-length one-dimensional array matching the model's expected input dimensions, since data arrives in batches; rather than relying on handcrafted features, the system employs an end-to-end deep learning approach that automatically captures discriminative temporal patterns such as abnormal firing rates, irregular burst activity, or altered amplitude

distributions indicative of neuromuscular dysfunction. This module's primary element is a previously trained 1D-CNN which is capable of processing time-series sEMG data as convolutional filters effectively capture localized temporal relationships and frequency characteristics from minimally processed data. It comprises convolutional layers which have small kernel sizes, followed by pooling layers to diminish the dimensionality and is catered for resource-limited edge devices through a limited number of parameters, similar to existing edge-based approaches which are used for real-time neurological detection [24], with a final sigmoid layer used to classify each segment as normal or abnormal. Inference is performed locally on the Raspberry Pi 5, rendering the system independent of the cloud and agnostic to cloud services, while the classification outcomes are integrated into the dashboard operating on the Raspberry Pi, where the linear opinion pool amalgamates them with the graph model's results to yield the ultimate decision. The main advantages of the system are portability and cost effectiveness due to easily available cheap components, low latency for inference on the device and scalability, that is, the possibility to replace or update individual components like sensor, microcontroller or classification model separately.

3. RESULTS AND DISCUSSION

3.1. Natural language processing based classification results

This module incorporates three core models required for the NLP component: BiLSTM-CRF which extracts features, naive Bayes probabilistic distribution which establishes relationships between medical features, and the SortPool algorithm, which performs a classification task on the constructed graph. Evaluation of these models was performed on 20 per cent of the data. The BiLSTM-CRF model outputs were passed onto the probabilistic relationship establishment and then a graph of those relationships was passed onto a neural network operating over graph-structured data (GNN). Final results obtained from the NLP pipeline are the results of the GNN model depicted in Table 1. All these models are a part of the pipeline of prediction alongside the EMG classification model. Regarding the classification task on the graph, evaluation was performed on two primary models: sort pool and global mean pool [25]. The models were trained using the collection of graphs of the patients produced in the earlier stages.

Table 1. Graph classification performance comparison across pooling strategies

Metric	Sort pooling	Global mean pooling
Total accuracy	0.7635	0.7260
Recall	0.7410	0.7080
Precision	0.7820	0.7530
F1 score	0.7610	0.7300

By concentrating on the most informative nodes, the sort pooling algorithm as in Table 1 improved classification performance on graphs characterized by sparsity and structural heterogeneity. Selective focus enabled the model to distinguish meaningful nodes from background noise with greater precision. However, a trade-off was observed as recall decreased. This was most likely due to informative nodes having lower degrees that were filtered out in top-k selection. Table 2 shows the outcomes of the naive Bayes implementation for linking the symptoms with the specific organ, derived from the computed probabilities. Every symptom node is initialized with a frequency value of 1 and connected to its corresponding organ (for example, O_1, O_2, O_3) each carrying its associated naive Bayes probability factor. For instance, S_1 shows a strong likelihood of being associated with O_1, with a probability value of 0.93714, whereas symptoms such as S_2 and S_3 display more distributed probabilities spanning several anatomical structures. This probabilistic framework enables inference over the most plausible organ corresponding to a designated symptom, ultimately assisting in the overall diagnosis.

3.2. Electromyography-based classification results

To help identify neuromuscular disorders through EMG signals, a 1D-CNN model was designed in the sEMG classification part of the framework. A careful evaluation was performed on the test dataset resulting in an area under the curve (AUC) of 0.99. The indication about the model is that its almost perfect in classifying normality and abnormality in the signal patterns. The substantial AUC means that the tool has an excellent ability when it comes to classifying the correct class label. To understand the performance in detail, a confusion matrix was constructed using the test data, as presented in Figure 3. The model was able to correctly identified

109 abnormal cases (true positives) and 136 normal cases (true negatives). Only 3 cases were identified as abnormal when they were originally marked as normal (false positives), while 2 cases with abnormalities were incorrectly flagged as normal cases (false negatives). This error rate, shown in Figure 3, reflects a classifier, having large values of precision and sensitivity. High precision indicates that the model performs fewer false alarms which can help reduce patient anxiety. The model also shows a high recall value which indicates that critical conditions are unlikely to go undetected, a balance that matters greatly in clinical settings.

To optimize the decision boundary for improved user experience, the F1-score was assessed at multiple classification thresholds. The F1-score is valuable for class imbalance and when the costs of a false positive and a false negative diverge. As shown in Figure 4, the threshold of 0.5 is traditionally used as a cutoff point for classification models. However, the F1-score curve suggests that a classification model should be more optimal at a threshold of 0.407. The 0.407 threshold demonstrates the most equitable balance of sensitivity (or recall) and specificity (or precision). A threshold of 0.407 should increase the clinical application of this testing model.

Table 2. Illustrative naive Bayes association outputs

Symptom	Number of occurances	Anatomical structure	Probability factor
S_1	1	O_1	0.937140
S_2	1	O_2	0.50
S_3	1	O_3	0.40
S_4	1	O_3	0.40
S_5	1	O_2	0.166670
S_7	1	O_2	0.166670
S_4	1	O_2	0.166670
S_8	1	O_3	0.10
S_2	1	O_3	0.10
S_9	1	O_1	0.042860
S_2	1	O_1	0.017140

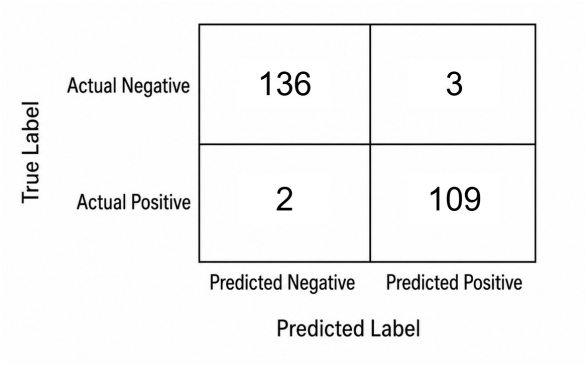

Figure 3. Confusion matrix showing the test results of the EMG classifier

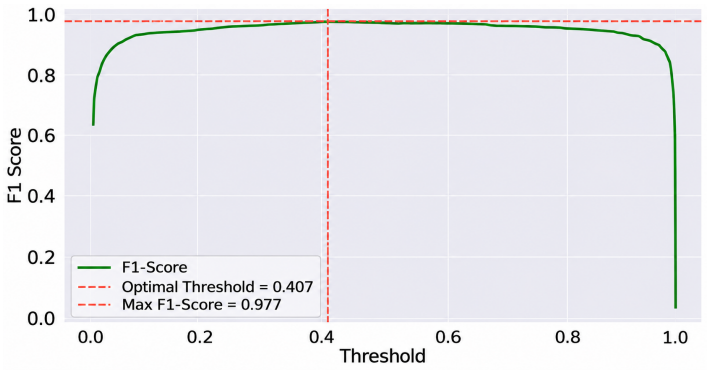

Figure 4. F1-score vs. classification threshold showing optimal threshold at 0.407

Threshold tuning was validated by comparing the performance metrics at the default threshold of 0.5 with those at the optimal threshold of 0.407. The analysis showed that the optimal threshold increases the true positive rate (TPR) from 0.968 to 0.974, while the false positive rate (FPR) only slightly increases from 0.014 to 0.017. This means that the system has become more sensitive to detecting true positives for neuromuscular abnormality cases, and has incurred only a negligible loss of specificity. Figure 5 illustrates the receiver operating characteristic (ROC) curve analysis to show how performance is improved at the optimal threshold.

A threshold of 0.407 was preferred instead of the standard 0.5 due to the former being nearer the optimum threshold value with high TPR and low FPR, giving improved sensitivity that would aid early detection and prompt medical treatment, as well as improved accuracy of diagnosis. The trade-off in the latter is clinically permissible due to the greater danger of missing an abnormal case. The comprehensive evaluation metrics for the proposed system are detailed in Table 3. These findings point to a high level of classification reliability, with all the metrics reflecting balance in the performance of the model. This strong statistical validity ensures the effectiveness of the model in assessing the need for MRIs.

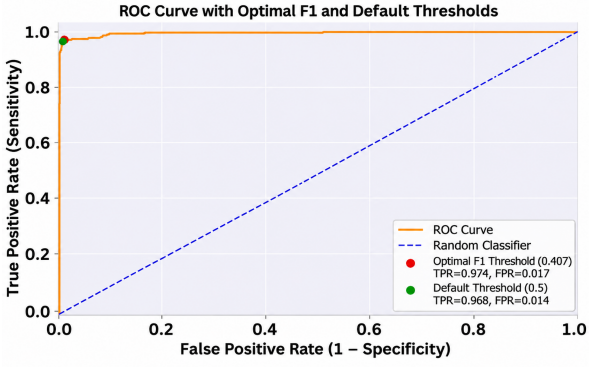

Figure 5. Threshold tuning comparison in the ROC curve: performance at optimal F1 threshold (0.407) vs. default threshold (0.5)

Table 3. Evaluation metrics of the system

Metric	Value
Total Accuracy	0.9800
Recall	0.9820
Precision	0.9732
F1 score	0.9776

4. CONCLUSION

This paper puts forward a multimodal diagnostic approach that uses the combination of clinical text mining along with the analysis of EMG signals to achieve high precision and effectiveness in neurological diagnosis. A BioBERT fine-tuned BiLSTM-CRF model extracts clinical features from the input text, which are then structured into knowledge graphs. The SortPooling technique is applied to GNN in order to predict the need for an MRI scan. The system utilizes real-time data acquisition by means of the EMG signals using ESP32 sensors, with the edge-based computation being done through Raspberry Pi 5, allowing for instant neuromuscular evaluation at the point of treatment. Combining these multimodal predictions with the linear opinion pool method uses both types of predictions. It is a cost-effective and scalable step toward smart health care systems.

5. LIMITATIONS AND FUTURE SCOPE

In spite of the encouraging results which were produced by the multimodal approach, there still exists a bunch of drawbacks in the system. First, the EMG model was built based on limited training data from a primary dataset gathered through the single-channel BioAmp EMG sensor. In turn, this can prevent us

from taking into account possible neuromuscular disease features. Second, the system's generalizability is restricted to the biceps brachii muscle and the clinical text formats used during training. Third, in order to perform a real-time prediction in resource-limited devices (ESP32 and Raspberry Pi), lightweight models were prioritized over computationally intensive transformer- or large language model (LLM)-based architectures, potentially limiting maximum accuracy. Future work will expand the dataset with pathological clinical samples, investigate model quantization for deploying more advanced models on edge hardware, and validate the system through prospective clinical trials by comparing its recommendations with clinician referrals using a structured clinical evaluation framework.

FUNDING INFORMATION

No funding involved in the project.

AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Prashant Ronad	✓	✓	✓		✓			✓		✓			✓	
Eshaan Mathur	✓	✓	✓		✓			✓		✓			✓	
Anirudh Suresh	✓	✓		✓	✓				✓		✓			
Abhishek Manja Naik	✓			✓		✓		✓	✓		✓			
Sudarshan Mahantesh Ullegaddi	✓	✓		✓		✓			✓		✓			
Jayati Vachhani	✓					✓			✓	✓	✓			
Jyoti Shetty	✓						✓			✓		✓	✓	

C : **C**onceptualization
M : **M**ethodology
So : **So**ftware
Va : **Va**lidation
Fo : **Fo**rmal Analysis
I : **I**nvestigation
R : **R**esources
D : **D**ata Curation
O : Writing - **O**riginal Draft
E : Writing - Review & **E**diting
Vi : **Vi**sualization
Su : **Su**pervision
P : **P**roject Administration
Fu : **Fu**nding Acquisition

CONFLICT OF INTEREST STATEMENT

No conflicts of interest are declared by the authors in relation to the work presented here.

DATA AVAILABILITY

Data supporting the results reported in this study are available from the corresponding author, [JS], on request.

REFERENCES

[1] Markets and Markets, "NLP in healthcare & life sciences market," marketsandmarkets.com. Accessed: Nov. 30, 2024. [Online]. Available: https://www.marketsandmarkets.com/Market-Reports/healthcare-lifesciences-nlp-market-131821021.html

[2] V. Prabhu, S. G, H. Watanuki, E. Mathur, P. Ronad, and N. Vasu, "NLP pipeline for electronic health records and predicting the necessity of MRI scan requirement using graph neural network," in *2025 5th International Conference on Robotics, Automation, and Artificial Intelligence (RAAI)*, Dec. 2025, pp. 1–7, doi: 10.1109/RAAI67517.2025.11423345.

[3] J. N. Acosta, G. J. Falcone, P. Rajpurkar, and E. J. Topol, "Multimodal biomedical AI," *Nature Medicine*, vol. 28, no. 9, pp. 1773–1784, Sep. 2022, doi: 10.1038/s41591-022-01981-2.

[4] V. Gohel and N. Mehendale, "Review on electromyography signal acquisition and processing," *Biophysical Reviews*, vol. 12, no. 6, pp. 1361–1367, Dec. 2020, doi: 10.1007/s12551-020-00770-w.

[5] E. D. Pasmanasari and J. A. Pawitan, "The potential of electromyography signals as markers to detect and monitor Parkinson's disease," *Biomedical and Pharmacology Journal*, vol. 14, no. 1, pp. 373–378, Mar. 2021, doi: 10.13005/bpj/2136.

[6] H. Bengacemi, A. H.-Gharbi, P. Ravier, K. A.-Meraim, and O. Buttelli, "Surface EMG signal classification for Parkinson's disease using WCC descriptor and ANN classifier," in *International Conference on Pattern Recognition Applications and Methods*, 2021, pp. 287–294, doi: 10.5220/0010254402870294.

[7] P. V. Nandankar, A. Kothiyal, K. K. D, A. Patil, H. Patil, and R. Maranan, "Parkinson's disease early detection and classification based on EMG signal using spherical convolutional neural network," in *2024 8th International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC)*, Oct. 2024, pp. 1140–1146, doi: 10.1109/I-SMAC61858.2024.10714636.
[8] E. T. Tan *et al.*, "Quantitative MRI differentiates electromyography severity grades of denervated muscle in neuropathy of the brachial plexus," *Journal of Magnetic Resonance Imaging*, vol. 56, no. 4, pp. 1104–1115, Oct. 2022, doi: 10.1002/jmri.28125.
[9] F.-J. Liu, Y.-P. Sun, Y. Shen, W.-Y. Ding, and L.-F. Wang, "Prognostic value of magnetic resonance imaging combined with electromyography in the surgical management of cervical spondylotic myelopathy," *Experimental and Therapeutic Medicine*, vol. 5, no. 4, pp. 1214–1218, Apr. 2013, doi: 10.3892/etm.2013.934.
[10] J. X. Jiang, A. Malhotra, and G. Bai, "Factors associated with hospital commercial negotiated price for magnetic resonance imaging of brain," *JAMA Network Open*, vol. 6, no. 3, Mar. 2023, doi: 10.1001/jamanetworkopen.2023.3875.
[11] G. Lample, M. Ballesteros, S. Subramanian, K. Kawakami, and C. Dyer, "Neural architectures for named entity recognition," in *2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2016, pp. 260–270, doi: 10.18653/v1/N16-1030.
[12] M. Zhang, Z. Cui, M. Neumann, and Y. Chen, "An end-to-end deep learning architecture for graph classification," in *32nd AAAI Conference on Artificial Intelligence*, 2018, pp. 4438–4445, doi: 10.1609/aaai.v32i1.11782.
[13] Z. Zhao, Q. Jin, F. Chen, T. Peng, and S. Yu, "A large-scale dataset of patient summaries for retrieval-based clinical decision support systems," *Scientific Data*, vol. 10, no. 1, Dec. 2023, doi: 10.1038/s41597-023-02814-8.
[14] A. P. Quimbaya *et al.*, "Named entity recognition over electronic health records through a combined dictionary-based approach," *Procedia Computer Science*, vol. 100, pp. 55–61, 2016, doi: 10.1016/j.procs.2016.09.123.
[15] R. Panchendrarajan and A. Amaresan, "Bidirectional LSTM-CRF for named entity recognition," in *32nd Pacific Asia Conference on Language, Information and Computation*, 2018, pp. 531-540.
[16] Ö. Uzuner, I. Solti, and E. Cadag, "Extracting medication information from clinical text," *Journal of the American Medical Informatics Association*, vol. 17, no. 5, pp. 514–518, Sep. 2010, doi: 10.1136/jamia.2010.003947.
[17] J. Li, R. Liu, C. Chen, S. Zhou, X. Shang, and Y. Wang, "An RG-FLAT-CRF model for named entity recognition of chinese electronic clinical records," *Electronics*, vol. 11, no. 8, Apr. 2022, doi: 10.3390/electronics11081282.
[18] L. Gligic, A. Kormilitzin, P. Goldberg, and A. N.-Holgado, "Named entity recognition in electronic health records using transfer learning bootstrapped neural networks," *Neural Networks*, vol. 121, pp. 132–139, 2020, doi: 10.1016/j.neunet.2019.08.032.
[19] K. Xu, Z. Yang, P. Kang, Q. Wang, and W. Liu, "Document-level attention-based BiLSTM-CRF incorporating disease dictionary for disease named entity recognition," *Computers in Biology and Medicine*, vol. 108, pp. 122–132, May 2019, doi: 10.1016/j.compbiomed.2019.04.002.
[20] F. Ren, X. Dong, X. Xu, J. Chen, and D. Zhang, "BERT-FLAT-CRF: a parallel model for chinese policy named entity recognition," *SSRN*, 2023, doi: 10.2139/ssrn.4442846.
[21] J. Zhong, Z. Xuan, K. Wang, and Z. Cheng, "A BERT-Span model for Chinese named entity recognition in rehabilitation medicine," *PeerJ Computer Science*, vol. 9, Aug. 2023, doi: 10.7717/peerj-cs.1535.
[22] W. Yoon, C. H. So, J. Lee, and J. Kang, "CollaboNet: collaboration of deep neural networks for biomedical named entity recognition," *BMC Bioinformatics*, vol. 20, no. S10, May 2019, doi: 10.1186/s12859-019-2813-6.
[23] M. Rotmensch, Y. Halpern, A. Tlimat, S. Horng, and D. Sontag, "Learning a health knowledge graph from electronic medical records," *Scientific Reports*, vol. 7, no. 1, Jul. 2017, doi: 10.1038/s41598-017-05778-z.
[24] B. Coffen and M. S. Mahmud, "TinyDL: edge computing and deep learning based real-time hand gesture recognition using wearable Sensor," in *2020 IEEE International Conference on E-health Networking, Application & Services (HEALTHCOM)*, Mar. 2021, pp. 1–6, doi: 10.1109/HEALTHCOM49281.2021.9399005.
[25] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," Feb. 22, 2017, *arXiv:1609.02907*.

BIOGRAPHIES OF AUTHORS



Prashant Ronad is a Computer Science and Engineering undergraduate at RV College of Engineering, Bengaluru. His research focuses on generative artificial intelligence (AI), machine learning, and natural language processing. He has developed an image outpainting pipeline using a diffusers model to expand the field of view for XR headsets and designed a workflow to generate high-precision depth maps. His project work includes building a stock analysis chatbot using a retrieval-augmented generation (RAG) pipeline, which increased accuracy to 89.3%, and developing a geospatial ride-hailing optimizer using predictive demand modeling to reduce peak-hour ride denials. He can be contacted at email: prashantronad870@gmail.com.



Eshaan Mathur is an undergraduate student in Computer Science and Engineering at RV College of Engineering, Bengaluru. He recently implemented a push-based, API-driven composable stream processing system with real-time querying, supporting both stateless and stateful queries featuring role-based access control with model context protocol (MCP) for natural language interpretation and adaptive microservices. His project work includes a cassini spacecraft scientific AI assistant using a RAG pipeline over mission data from PDS with LLaMA-3 70B, a DDoS and TLS attacks mitigation app using eBPF, and a real-time fire and smoke detection system using an ESP32-CAM network and Raspberry Pi. He can be contacted at email: eshaanmathur157@gmail.com.



Anirudh Suresh is a Computer Science and Engineering undergraduate student at RV College of Engineering, Bengaluru. With a strong focus on system design, machine learning, and data science, his work spans practical and impactful applications, including the development of a comprehensive bank management tool aimed at streamlining financial operations and a handwriting recognition system powered by deep learning for digitizing written input with high accuracy. He also built an interactive data dashboard to enable real-time insights and visualization for decision-making. His interests lie at the intersection of intelligent systems and scalable architectures, with a commitment to building data-driven solutions that enhance user experience and operational efficiency. He can be contacted at email: anirudh2suresh@gmail.com.



Abhishek Manja Naik is an Electronics and Communication Engineering student at RV College of Engineering, Bengaluru. He has technical expertise in system verilog, verilog, digital communication system, optimization techniques, digital signal processing (DSP), and machine learning. He has worked on UVM-based OSI layers verification, focusing on MAC layer verification. Additionally, he has practical experience in working with analog micro-electronics circuits as well as digital communications systems. His core interests lie in computer network protocols, SysVerilog, and embedded systems. He can be contacted at email: abhishekms0509@gmail.com.



Sudarshan Mahantesh Ullegaddi is an Electrical and Electronics Engineering student at RV College of Engineering, Bengaluru, with technical expertise in verilog, DSP, and machine learning. He has worked on computer vision-based projects integrating STM32 microcontrollers, focusing on real-time image processing and automation tasks. Additionally, he has practical experience in working with GPS and GSM modules for location tracking and wireless communication applications. His core interests lie in embedded systems, internet of things (IoT) solutions, and the fusion of hardware with intelligent algorithms to develop efficient, real-world engineering solutions. He can be contacted at email: sudarshanullegaddi2004@gmail.com.



Jayati Vachhani is currently pursuing a degree in Aerospace Engineering at RV College of Engineering. A passionate multitasker, her core interests lie in structural design and engineering analysis, with strong proficiency in tools such as SolidWorks and ANSYS. She also holds foundational knowledge in programming languages like Python and is keen to explore further in the computing domain. Deeply driven by a commitment to sustainability, she has contributed to projects aimed at reducing carbon emissions in automobiles. Her research interests also extend into quantum computing, where she is presently engaged in solving real-world problems such as gate assignment optimization for airports. She aspires to address practical engineering challenges with innovative and sustainable solutions. She can be contacted at email: jayativachhani2101@gmail.com.



Dr. Jyoti Shetty is an associate professor in the Department of Computer Science and Engineering, RV College of Engineering, Bengaluru, India, and has 16 years of teaching and 2 years of industry experience. Her specialization includes data mining, machine learning, and cloud computing. She has published 50+ research papers in reputed journals and conferences. She has also executed sponsored projects funded by various agencies nationally and internationally. She was the recipient of awards such as the SAP Award of excellence from IIT Bombay for demonstrating ICT in education in 2016, HPCC Systems Mentor Badge Award in 2021, Best Young Teacher Award ISTE, HPCC Systems Community, and HPCC Systems Community Recognition Award. She can be contacted at email: jyothis@rvce.edu.in.