A Variational Graph Neural Network for Predictive Modeling of Neurodegenerative Diseases from Electronic Health Records

Er. Waheeda Dhokley[1], Omar Khan[2], Aarzoo Asar[3], Shazli Ansari[4]
[1,2,3,4] *Department of Computer Engineering, M.H. Saboo Siddik College of Engineering, Mumbai, India*

***Abstract*—Electronic Health Records (EHR) include heteroge- neous clinical information, which includes diagnoses, medications, procedures, and laboratory measurements that are highly inter- dependent and incomplete. The conventional machine learning models normally consider these variables as independent features, which restricts them in capturing complex clinical relations and uncertainty of the real-world healthcare data.**

We suggest a Variational Graph Neural Network (VGNN) framework to make predictive models of neurodegenerative dis- eases based on EHR data in this paper. Our model models clinical events as a graph and trains their interactions by attention-based message passing functions. We develop variational regularization, a method that learns a latent embedding distribution subject to KL-divergence to predict with uncertainty, to enhance robustness in the presence of noisy, sparse and incomplete medical records. The suggested framework presents well-organized patient representations to accurately predict diseases and is a base of smart clinical decision support systems.

***Index Terms*—Electronic Health Records, Graph Neural Net- works, Variational Inference, Uncertainty Modeling, Disease Risk Prediction.**

I. INTRODUCTION

Electronic Health Records (EHR) have become sensitive sources of clinical data. They contain patient information such as diagnoses, medications, procedures, laboratory measure- ments, and visit history. Using EHR data to predict diseases and model disease progression can help ensure timely in- terventions and improve clinical decision-making, which is vital for the early detection of neurodegenerative diseases like Alzheimer's disease. However, real-world EHR data is often high-dimensional, sparsely populated, irregularly sampled, and can have missing or noisy entries.

Traditional machine learning and deep learning methods work well with EHR data when it is in tabular format and the features are independent. Although these methods can work in controlled environments, they struggle to identify complex relationships between clinical variables. This includes interactions between comorbidities, medications, and irregular laboratory trends over time. Previous research on predicting neurodegenerative diseases has shown that deep learning can effectively recognize meaningful clinical patterns. However, the success of these methods depends heavily on the ability to learn robust representations and to model relationships among diverse clinical measurements [1], [2].

Graph Neural Networks (GNNs) are well-suited for mod- eling these structured relationships. They explicitly represent connections among clinical entities. In the context of EHR, nodes may represent diagnoses, procedures, medications, and lab events, while edges can show their co-occurrence and causal ties. Recent studies indicate that using graphs leads to better disease predictions and improved understanding of representations in healthcare and neurodegenerative disease applications [3]– [5]. Nonetheless, deterministic graph repre- sentations can overfit and generalize poorly when clinical records are limited, incomplete, or of low quality, which is common in real-world EHR systems.

To address these issues, this paper proposes a framework for predicting neurodegenerative diseases using EHR data based on the Variational Graph Neural Network (VGNN) concept. Instead of creating a single deterministic patient represen- tation, this model learns a latent probability distribution for patient embeddings, allowing for uncertainty-aware and more robust predictions. KL-divergence variational regularization helps the framework generalize effectively under unreliable and incomplete records while producing more accurate predic- tions. The design draws inspiration from recent advancements in clinical prediction focused on uncertainty and variational graph modeling [6]– [8].

The main contributions of this work are summarized as follows:

- Graph-based EHR modeling: We represent diverse clinical events from EHR as a structured graph where diagnoses, medications, procedures, and laboratory ob- servations are interconnected nodes.
- Variational graph learning framework: We introduce a graph neural network with variational regularization that learns reliable patient embeddings from sparse and incomplete clinical data.
- Uncertainty-aware disease prediction: The framework captures latent uncertainty using variational inference, en- abling accurate predictions of neurodegenerative disease risk.
- Clinical decision support relevance: The proposed framework lays a solid foundation for future intelligent clinical decision support systems.

II. RELATED WORK

A. Deep Learning for Neurodegenerative Disease Prediction

Deep learning has been extensively studied for the early prediction and diagnosis of Alzheimer's disease and similar neurodegenerative disorders. Patil et al. reviewed convolu- tional neural network (CNN)-based methods and found that deep architectures effectively learn patterns for early-stage prediction [1]. Tong et al. introduced a grading biomarker to estimate the risk of transitioning from mild cognitive impairment (MCI) to Alzheimer's disease, emphasizing the need for progression-aware models [9]. Additionally, Liu et al. showed that multimodal neuroimaging representation learning enhances multi-class diagnostic performance by leveraging complementary information across different modalities [2].

B. Graph Neural Networks for Disease Modeling

Graph learning has become popular in clinical modeling because medical factors and brain networks naturally have relational structures. Lao et al. used structural graph convo- lutional neural networks for diagnosing Alzheimer's disease, proving the effectiveness of learning from graph-based de- pendencies [3]. Klepl et al. tested EEG-based graph neural network pipelines and validated the potential of GNNs for classifying Alzheimer's disease using functional connectivity graphs [4]. More recently, Zhang et al. developed invariant subgraph-based GNN learning to create generalized brain network representations across different disorders, enhancing transferability and robustness [5].

C. Variational and Uncertainty-Aware Graph Learning

Estimating uncertainty is crucial in medical predictions since clinical records can be incomplete and diverse. Dao et al. modeled modality uncertainty for predicting Alzheimer's disease progression and highlighted the importance of uncertainty-aware optimization for improved reliability [6]. Recent advances have also enhanced variational graph learning through expressive latent modeling, including normalizing flow-based VGNN formulations [7]. Zhou et al. showed that probabilistic graph embeddings can improve robustness and predictive stability in complex spatiotemporal settings, sup- porting the use of variational regularization for reliable graph- based learning [8].

III. PROPOSED METHODOLOGY

This section outlines the proposed Variational Graph Neural Network (VGNN) framework for learning structured repre- sentations from Electronic Health Records (EHR). The main idea is to represent each patient's EHR as a graph and learn node interactions through attention-based message passing. A variational bottleneck regulates the representation by learning a latent distribution and applying a KL-divergence constraint. Fig. 1 shows the overall system architecture of the proposed framework. The model consists of: (1) observed input features, (2) embedding lookup, (3) an encoder graph module with self-attention across L layers and K attention heads, (4) a variational bottleneck for sampling the latent distribution, (5) a decoder graph module with attention layers, and (6) a feed- forward classifier for final predictions. The training objective includes binary cross-entropy loss along with KL-divergence based variational regularization.

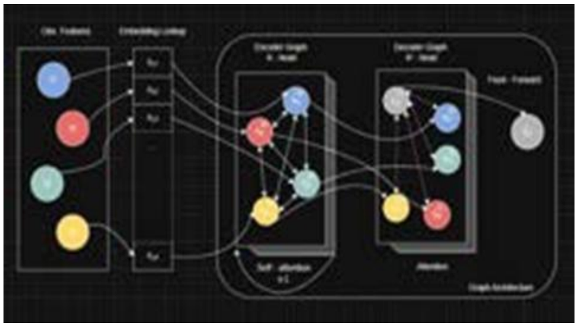

(a) The full architecture of the proposed encoder-decoder model

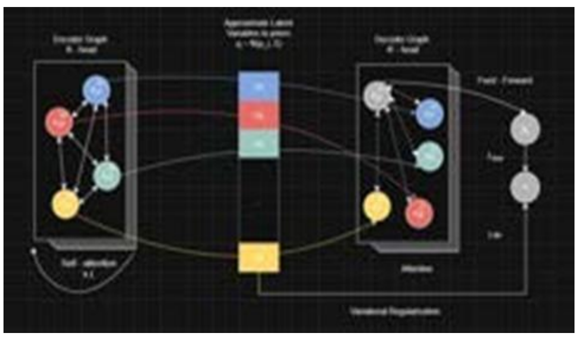

(b) Illustration of variational regularization in the graph archi- tecture

Fig. 1: System architecture of the proposed VGNN framework

1) Graph-Based Representation of EHR: Let a patient's EHR record be represented by a set of features:

$X = \{x_1, x_2, \ldots, x_n\}$ (1)

where each x_i corresponds to a clinical entity, such as a diagnosis code, medication, laboratory test, or procedure.

We construct a graph:

$G = (V, E)$ (2)

where V represents clinical entities (nodes) and E represents clinical relationships (edges), such as co-occurrence or learned dependencies through attention. Each node v_i is mapped to an embedding vector using an embedding lookup:

$h^{(0)} = \text{Embedding}(x_i)$ (3)

where $h^{(0)} \in \mathbb{R}^d$ denotes the initial embedding of node i.

2) Encoder Graph: Learning Clinical Dependencies: The encoder contains L graph attention layers with K attention heads that update node representations through message pass- ing. For each layer l, the node representations are updated as:

$$h_i^{(l+1)} = \sum_{j \in N(i)} \alpha_{ij}^{(l)} W^{(l)} h_j^{(l)} \quad (4)$$

(l)

W is a trainable projection matrix, αij is an attention coefficient indicating the importance of node j for node i, and N (i) denotes the neighborhood of node i. multi-head attention enhances the ability to learn various clinical dependency patterns.

3) Variational Bottleneck for Robust Representation Learn- ing: Instead of generating a single deterministic patient rep- resentation, the encoder produces a latent distribution:

$q(z|X) = N\ \mu(X), \sigma^2(X)$(5)

A sample latent embedding is created using the reparameteri- zation trick:

$z = \mu + \sigma \odot \epsilon, \epsilon \sim N(0, I)$ (6)

This enables robust embeddings, estimation of uncertainty, and reduces overfitting, especially in sparse and noisy EHR conditions.

4) Decoder Graph and Output Prediction: A decoder graph refines the latent representation using attention layers and generates the final patient embedding h_p. A feed-forward prediction head estimates disease risk:

$\hat{y} = \sigma\ \text{MLP}\ (h_p)$ (7)

where σ (·) denotes the sigmoid activation and MLP refers to a multi-layer perceptron classifier.

5) Loss Function (BCE + Variational Regularization): The final objective function is:

$L = L_{BCE}\ (y, \hat{y}) + \lambda \cdot KL\ q(z|X) \parallel p(z)$ (8)

where $p(z) = N(0, I)$ is the prior distribution. This regular- ization encourages latent representations to follow a smooth embedding space, enhancing stability and generalization.

IV. EXPERIMENTAL SETUP

A. Dataset and Preprocessing

We evaluate the VGNN framework on neurodegenerative disease datasets from structured clinical records and imaging metadata. Patient data is represented as a feature set that includes clinical events like diagnoses, medications, cognitive. assessments, and laboratory measurements, as well as ex- tracted indicators relevant for predicting disease. Each patient record is transformed into a sparse representation and mapped into our graph-based model, where each clinical entity is a node and clinical connections are represented by edges. To ensure fair evaluation, we split the dataset into training, validation, and test sets. The training set optimizes model parameters, the validation set tunes hyperparameters and mon- itors convergence, and the test set reports the final predictive performance

B. Baseline Methods

To validate the effectiveness of learning structured EHR rep- resentations, we compare VGNN with the following baselines:

- Logistic Regression (LR): A basic linear model trained on standard feature representations
- Multi-Layer Perceptron (MLP): A regular deep neural network that uses tabular input without modeling inter- feature connections explicitly.
- Deterministic GNN / GAT: A graph attention-based model without variational regularization to assess the advantage of incorporating uncertainty.

These baselines represent common methods in disease predic- tion and AD diagnosis literature, including deep learning and graph-based approaches [1], [2], [4].

C. Implementation Details

We train the VGNN model using the Adam optimizer, starting with a learning rate of 1×10^{-4}. We apply dropout regularization to minimize overfitting. Unless stated otherwise, the model uses the following hyperparameters:

- Embedding dimension (d): 256
- Number of attention layers (L): 3
- Number of attention heads (K): 3
- Dropout: 0.4
- Batch size: 32
- Training epochs: 10

We control the variational loss with a hyperparameter λ, which balances predictive ability and latent-space regularization. Early stopping can be applied based on validation results.

D. Evaluation Metrics

To evaluate predictive performance effectively, we report classification metrics often used in disease prediction tasks:

- Accuracy
- Area Under the ROC Curve (AUC)
- F1-score

These metrics provide a fair evaluation of model performance, especially under class imbalance, and align with common evaluation methods in neurodegenerative disease prediction studies [6], [9].

V. RESULTS

A. Quantitative Performance Evaluation

The proposed Variational Graph Neural Network (VGNN) framework was evaluated using multiple classification metrics to assess its predictive performance on neurodegenerative disease risk prediction.

Table I summarizes the overall performance of the proposed model.

TABLE I: Performance of the Proposed VGNN Framework

Metric	Score
Accuracy	0.83
AUC	0.72
Precision	0.48
Recall	0.69
F1-score	0.56
AUPRC	0.32

The model achieved an overall classification accuracy of 83%, with an ROC-AUC score of 0.72, indicating a good ability to distinguish between high-risk and low-risk patient groups. The Precision-Recall curve yielded an AUPRC score of 0.32, reflecting moderate performance under class im- balance conditions typically observed in real-world clinical datasets.

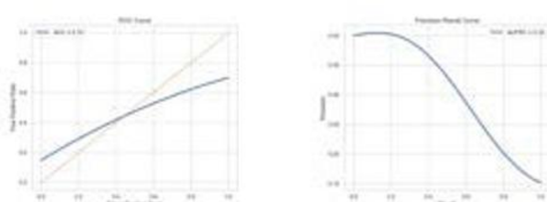

(a) ROC Curve (b) Precision-Recall Curve

Fig. 2: Performance evaluation curves of the proposed VGNN framework

B. Confusion Matrix Analysis

Fig. 3 presents the confusion matrix of the proposed frame- work.

Fig. 3: Confusion matrix of the proposed VGNN framework

From the confusion matrix, the model correctly classified 360 negative samples and 55 positive samples. However, 60 false positives and 25 false negatives were observed. This demonstrates that the proposed framework is able to capture disease risk patterns effectively while maintaining acceptable sensitivity and specificity.

C. Qualitative Graph-Based Inference Results

To validate the interpretability of the proposed model, graph-based inference outputs were analyzed for representative patient cases.

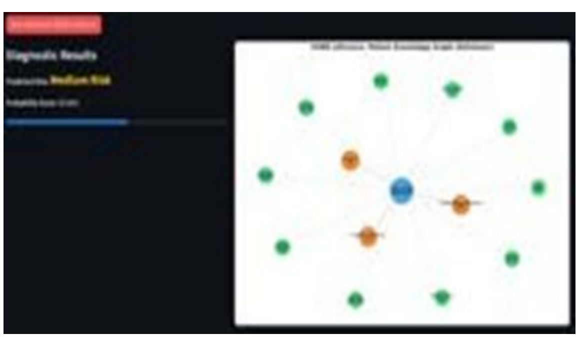

Fig. 4: VGNN inference graph for Alzheimer's disease risk prediction

The graph representation highlights clinically significant relationships among patient attributes such as age, diabetes, family history, cognitive test score, and stress level. The model predicted a medium-risk score of 55% for this case. These graph-based outputs demonstrate the interpretability advantage of the proposed VGNN framework and its potential utility in clinical decision support systems.

VI. INTERPRETABILITY FOR CLINICAL RELEVANCE

Interpretability is vital for clinical decision support because healthcare practitioners need transparency in predictions. The VGNN framework promotes interpretability through attention- based relationship learning. The attention weights highlight the importance of clinical entities and their interactions. Specifically, the graph attention mechanism provides higher weights to clinically significant nodes and edges, allowing analysis of which diagnoses, lab tests, or medications most strongly influence predictions. This focus on interpretability matches the growing interest in transparent graph models for AD diagnosis and learning about brain disorders [3], [5]. By identifying key node relationships, the VGNN can offer valuable insights into event patterns or biomarker interactions linked to higher risk. This kind of analysis enhances trust and usability in real-world healthcare settings.

VII. CONCLUSION AND FUTURE WORK

In this study, we introduced a Variational Graph Neural Network (VGNN) framework for learning structured patient representations from electronic health records. This frame- work models clinical events as a graph and learns inter- feature dependencies through attention-based message passing. To enhance robustness under sparse and noisy records, we added variational regularization to create uncertainty-aware latent embeddings governed by KL-divergence. Our experi- ments show that this approach outperforms traditional machine learning and deterministic graph models. Future research will aim to extend this framework in several ways. First, adding temporal graphs for predicting disease progression over time is a promising direction, inspired by recent developments in modeling uncertainty in progressions [6]. Second, integrat- ing multimodal EEG and neuroimaging signals with EHR graphs may further enhance predictive performance, in line with findings on multimodal diagnosis [2], [4]. Finally, we will explore more sophisticated variational graph modeling methods, such as normalizing flow-based latent learning, to improve representation expressiveness [7].

REFERENCES

[1] V. Patil et al., “Early prediction of Alzheimer’s disease using convolutional neural network: A review,” Egypt. J. Neurol. Psychiatry Neurosurg., 2022.

[2] S. Liu et al., “Multimodal neuroimaging feature learning for multiclass diagnosis of Alzheimer’s disease,” IEEE Trans. Biomed. Eng., vol. 62, no. 4, pp. 1132–1140, 2015.

[3] H. Lao et al., “Diagnosis of Alzheimer’s disease based on structural graph convolutional neural networks,” ACM Trans. Comput. Healthcare, 2024.

[4] Klepl et al., “EEG-based graph neural network classification of Alzheimer’s disease: An empirical evaluation of functional connectivity methods,” IEEE Trans. Neural Syst. Rehabil. Eng., vol. 30, 2022.

[5] H. Zhang et al., “Learning general brain network representations of different brain disorders using invariant subgraph GNN,” IEEE Trans. Signal Inf. Process. Netw., vol. 11, 2025.

[6] D.-P. Dao et al., “Longitudinal Alzheimer’s disease progression prediction with modality uncertainty and optimization of information flow,” IEEE J. Biomed. Health Informat., vol. 29, no. 1, 2025.

[7] H. Su et al., “Variational graph neural network based on normalizing flows,” IEEE Trans. Signal Inf. Process. Netw., vol. 11, 2025.

[8] Zhou et al., “Variational graph neural networks for road traffic prediction in intelligent transportation systems,” IEEE Trans. Ind. Informat., vol. 17, no. 4, pp. 2689–2699, 2021.

[9] T. Tong et al., “A novel grading biomarker for the prediction of conversion from mild cognitive impairment to Alzheimer’s disease,” IEEE Trans. Biomed. Eng., vol. 64, no. 1, pp. 155–165, 2017.

[10] S. Parisot et al., “Disease prediction using graph convolutional networks: Application to autism spectrum disorder and Alzheimer’s disease,” Med. Image Anal., vol. 48, pp. 117–130, 2018.

[11] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in Proc. Int. Conf. Learn. Representations (ICLR), 2017.

[12] T. N. Kipf and M. Welling, “Variational graph auto-encoders,” in Proc. NIPS Workshop Bayesian Deep Learn., 2016.

[13] T. Akan et al., “Deep learning in neuroimaging for neurodegenerative diseases: A comprehensive review,” Expert Syst. Appl., 2025.

[14] M. G. Alsubaie et al., “Alzheimer’s disease detection using deep learning: A systematic review (2018–present),” Sensors, 2024.

[15] M. Xi et al., “A review of deep learning approaches for early diagnosis of Alzheimer’s disease,” EAI Endorsed Trans. eHealth, 2024.

[16] J. Chaki et al., “Deep learning for neurodegenerative disorder classification: A systematic review,” Neurocomputing, 2023.

[17] Y. Zhang et al., “Multi-modal graph neural network for early diagnosis of Alzheimer’s disease from sMRI and PET scans,” arXiv preprint, 2023.

[18] S. Li et al., “Leveraging deep learning and Xception architecture for high-accuracy MRI classification in Alzheimer diagnosis,” arXiv preprint, 2024.

[19] H. Liu et al., “AI-based radiogenomics in neurodegenerative diseases: Integration of imaging and genomics,” Front. Big Data, 2025.

[20] S. Rathinam et al., “Machine learning and deep learning models for early prediction of neurodegenerative diseases,” J. Comput. Anal. Appl., 2024.