Journal of Computational Medicine and Informatics
Vol. 1, No. 2, Apr - Jun 2025, pp. 36-44
ISSN: 3107-3735 (O), DOI: https://doi.org/10.17051/JCMI/01.02.05

Knowledge Graph-Driven Clinical Decision Intelligence System for Personalized Treatment Optimization

Ali Bostani

Associate Professor, College of Engineering and Applied Sciences, American University of Kuwait, Salmiya, Kuwait. Email: abostani@auk.edu.kw, 0000-0002-7922-9857

Article Info

Article history:

Received : 14.04.2025
Revised : 12.05.2025
Accepted : 10.06.2025

Keywords:

Knowledge Graph;
Clinical Decision Support System;
Personalized Treatment;
Graph Neural Network;
Electronic Health Records (EHR);
Healthcare Analytics;
Medical AI;
Treatment Recommendation;
Explainable AI;
Precision Medicine

ABSTRACT

The increased complexity and heterogeneity of healthcare data demand smart systems that have the ability to model relationships amongst clinical entities to facilitate accurate and personalized decisions. The conventional machine learning and deep learning methods do not adequately address the complex relationship among patients, diseases, symptoms, and treatment and their application in clinical settings is limited. The research will introduce a clinical decision intelligence system that optimizes personalized treatment using a knowledge graph to combine the structured graph representation with graph neural network (GNN)-based learning. The longitudinal electronic health record (EHR) data is used to build a comprehensive knowledge graph, with nodes representing clinical entities and the relationships defined by the edges representing semantic and correlation-based relationships. The suggested framework uses representation learning through graphs to understand contextual interactions, which allows better predictions of the most effective treatment strategies. The model is tested on the MIMIC-III dataset and compared to the baseline models, such as logistic regression, random forest, LSTM, and attention-based models. The experiments have shown that the new method has better performance (in accuracy, F1-score, and AUC) than other methods. Also, the graph-based model can be utilized to make interpretable decisions based on the relational reasoning scheme, which facilitates explainable clinical decisions. The intended system provides a scalable and efficient solution towards advanced clinical decision support in precision medicine and intelligent healthcare analytics.

1. INTRODUCTION

The intensive healthcare systems digitization has contributed to the extensive use of electronic health records (EHRs) that have resulted in the production of large, heterogeneous, and longitudinal clinical data. Patient demographics, diagnoses, laboratory results, medications, and clinical notes are some of the data sources that have great potential in data-driven clinical decision-making. Nonetheless, the effective use of such multifaceted data is one of the crucial issues because of the fragmented representations, high dimensionality, and lack of organized relational modeling among clinical entities [1], [2]. Traditional machine learning and deep learning solutions, such as logistic regression, random forests, and recurrent neural networks, are extensively used in prediction of diseases and recommendations of treatment options [3]–[5]. Although these models have good predictive performances, they generally handle clinical variables either in isolation or in a cascading fashion, not depicting the extensive interrelationships among patients, diseases, symptoms, and treatments. The weakness greatly limits their ability to aid individualized and contextual clinical decision-making [6], [7]. Current developments in both knowledge graphs (KGs) and graph neural networks (GNNs) offer a new direction to overcome these drawbacks by facilitating the representation and relational learning of complex healthcare data in a structured fashion [8]–[10]. Knowledge graphs represent clinical entities as a graph with nodes and interactions as edges, whereas GNNs can be used to perform representation learning with neighborhood aggregation and message passing. In spite of these developments, the current literature tends to either not even integrate longitudinal EHR data or does not utilize both semantic and statistical relationships and is also limited to offer

interpretability in clinical decision support systems [11], [12].

This paper aims to eliminate these constraints by suggesting a knowledge graph-based clinical decision intelligence system to maximize personalized treatment. The framework proposed develops a multifaceted graph of the knowledge based on longitudinal EHR data and learning with the use of graph neural networks to identify circumstantial associations between clinical entities. This integration facilitates a better level of prediction, and aids in explainable and personalized treatment recommendations.

The key contributions of the work are as follows:

- A longitudinal structure of graph construction of knowledge which includes semantic and correlation-based relationships using longitudinal EHR data.
- A learning model of graph neural networks to improve clinical decision support.
- An individualized optimization of treatment based on relational reasoning.
- Experimental assessment on extensive clinical datasets in practice and comparing the results with the state-of-the-art baseline models.

2. RELATED WORK

2.1 Deep Learning for Clinical Prediction

The recent progress in the field of analytics in healthcare has widely applied deep learning methods to predict diseases and recommend treatments. Models based on recurrent neural networks (RNN) like Doctor AI have been suggested to capture temporal relationships of longitudinal EHR data [1]. In the same vein, RETAIN proposed an attention-based approach in enhancing interpretability in clinical predictions [2], and Deep Patient, an unsupervised representation learning based on patient embedding [3]. Although effective, the models are mostly based on sequential or tabular data representations and do not explicitly represent complex relationships with clinical entities (e.g., diseases, symptoms, medications). This weakness interferes with their capacity to make relational arguments to enable individualized clinical decisions [4], [5].

2.2 Knowledge Graphs in Healthcare

Knowledge graphs (KGs) have become of great interest in the representation of structured relationships among biomedical entities. They facilitate the combination of heterogeneous sources of data such as clinical records, biomedical ontologies and drug interaction databases [6], [7]. KGs are successfully used in activities like drug discovery, disease association analysis and clinical decision support. Nevertheless, most of the existing KG-based systems are based on fixed representations and rule-based inference which restricts the capacity to learn complex patterns using data. Moreover, they can usually use ineffective mechanisms to integrate both time and patient-specific data of longitudinal EHRs [8].

2.3 Graph Neural Networks for Clinical Applications

Graph neural networks (GNNs) are a generalization of deep learning to graph-structured input data, allowing learning a node representation by neuro-aggregating its neighbors. Most recent works have used GNNs in disease prediction, modeling patient similarity, and drug recommendation problems [9]–[11]. These methods show better performances by representing relational dependencies between clinical entities. However, there are a number of constraints of current GNN-based models in healthcare. The majority of methods build simplified graphs, and they do not consider both semantic and statistical relations. In addition, they fail to consider longitudinal dynamics of patient data and offer restricted interpretability toward clinical decision support [12], [13].

2.4 Research Gaps and Motivation

Based on the discussion above, the following key gaps are identified:

- L It does not have integrated frameworks that combine knowledge graphs and deep learning to support clinical decisions.
- Minimal use of longitudinal EHR data to graph-based models.
- Weak modeling of semantic and correlation based relationships.
- Unexplainability and lack of optimization of personalized treatment in current systems.

To overcome these difficulties, this paper presents a clinical decision intelligence system that uses a knowledge graph and combines graph representation with learning using graph neural networks [14], [15]. The suggested framework summarizes and encapsulates complicated clinical interactions, recommends individualized therapies, and enhances predictive accuracy and meaningfulness.

3. System Architecture

The suggested knowledge graph-based clinical decision intelligence system is a multi-step pipeline which converts raw clinical data into useful treatment recommendations as shown in Fig. 1. The architecture is made up of four modules that are interconnected: data preprocessing, knowledge graph construction, graph learning and decision intelligence. All modules are essential in facilitating correct and tailored clinical decision-making.

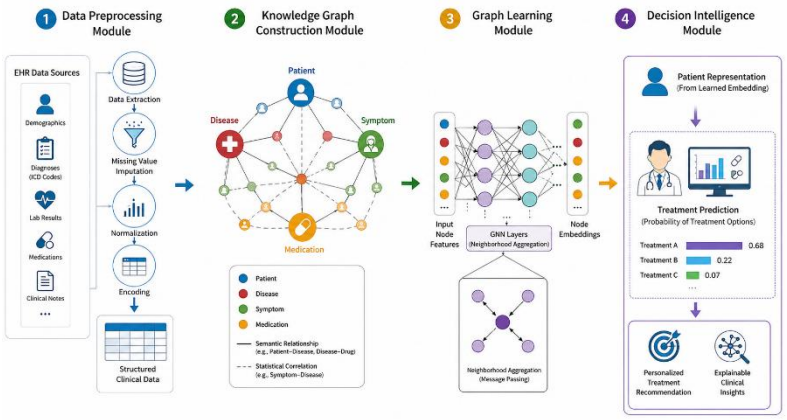


Fig. 1. System Architecture of the Proposed Knowledge Graph-Driven Clinical Decision Intelligence System

3.1 Data Preprocessing Module

The preprocessing module of data is charged with the task of converting raw electronic health record (EHR) data into a format that can be analyzed and categorized. Clinical datasets are usually heterogeneous, such as patient demographics, diagnosis, laboratory results, and medication records which is often incomplete and noisy. These issues are solved by imputating missing values with imputation methods, and normalizing numeric features to make variables consistently scaled. Categorical data like diagnosis and medication codes are coded into machine readable formats. The preprocessing stage is to make sure that the downstream modules are fed clean, standardized and high-quality data representations.

3.2 Knowledge Graph Construction Module

The module of knowledge graph construction takes the processed clinical information and transforms it to a structured graph format that represents clinical relationships between clinical entities. The node in this graph may be a patient, a disease, a symptom and a medication and the edges in the graph encode the interaction among the entities. The graph is built by combining semantic (e.g. patient-disease, disease-drug) and statistical correlations based on clinical data. This representation allows modeling the complex dependencies not reflected in the conventional tabular or sequential format of data. The system helps to make more informed decisions by structuring the clinical knowledge in the form of a graph and enables relational reasoning.

3.3 Graph Learning Module

The graph learning module utilizes the principles of Graph Neural Networks (GNNs) to learn high-quality representations of nodes between the knowledge graph generated. By aggregation in the neighborhood, the GNN is able to capture local and global structure information with the contextual information encoded by each node into the connected entities. This learning process allows the system to learn embeddings that capture complicated clinical associations, like interplay between diseases and treatments, or interrelationships between symptoms. Such learned representations are used as the basis of predictive modeling and decision-making activities.

3.4 Decision Intelligence Module

The learned graph embeddings are used to give customized recommendations on treatment in the decision intelligence module. The system forecasts optimal treatment plans, based on patient-specific representation, in the knowledge graph, predicting and refining the best-fit treatment plan by analyzing that specific representation in the knowledge graph. This module is probabilistic and generates outputs of treatment likelihood of the various treatment options and can be clinically interpreted through insights of relationship based on the graph architecture. Consequently, the system increases the accuracy of prediction in addition to the increased transparency of clinical decision-making processes.

4. Knowledge Graph Construction

Knowledge graph construction process is a crucial part of the proposed system, because it allows the complex relationships between clinical entities to be represented. In a difference to conventional tabular or sequential data structures, the knowledge graph models semantical and statistical dependencies, enabling a better relational reasoning in clinical decision-making, as shown in Fig. 2.

The clinical knowledge graph is defined to be:

$$G = (V, E) \quad (1)$$

where V denotes the set of nodes representing clinical entities, and E denotes the set of edges representing relationships among these entities.

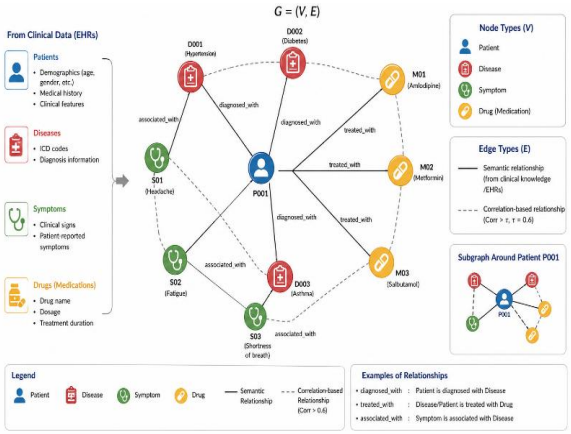


Fig. 2. Clinical Knowledge Graph Construction from Electronic Health Records

4.1 Node Definition

The nodes in the knowledge graph represent significant clinical objects that have been identified in electronic health records (EHRs). There are four major types of these entities:

- Patient Nodes: Denote distinct patients and consist of both demographic characteristics (e.g., age, gender) and integrated clinical characteristics based on the medical history.
- Disease Nodes: Are represented by diagnosed medical conditions, usually coded by standardized classification systems like ICD codes.
- Symptom Nodes: Nodes the clinical signs and symptoms that are observed and reported by the patient during her consultation or hospital visit.
- Drug Nodes: Depict drugs prescribed to the patients and where available they include dosage and treatment period.

The nodes are coupled with their feature vectors representing their attributes, and they can be used in graph learning downstream applications.

4.2 Edge Construction

Edges indicate how a relationship exists between the nodes and are very essential in the representation of interactions in the clinical field. In the proposed framework, the edges are formed based on semantic relationships as well as using statistical correlation.

Semantic Relationships

These are based on clinical knowledge and EHR records and they include:

- Patient-Disease: Refers to a patient that has been diagnosed with a particular disease.
- DiseaseDrug: Refers to treatment interrelationships between diseases and medications prescribed.
- SymptomDisease: Records the relationships between symptoms and underlying diseases.

Such relationships are based on the real-world clinical interactions and they are the backbone of the knowledge graph structure.

Statistical Correlation-Based Relationships

Besides the semantic links, there are also statistical dependencies among clinical variables that are introduced to the graph structure to enrich it. The definition of correlationbased edges is:

$$\mathrm{Corr}(x_i, x_j) > \tau \quad (2)$$

where $\mathrm{Corr}(x_i, x_j)$ represents the correlation coefficient between variables x_i and x_j, and τ is a predefined threshold controlling edge inclusion.

In this study, the threshold is set to $\tau = 0.6$, ensuring that only strong and clinically meaningful relationships are retained. It is done in this manner to enable the model to represent implicit connections, which may not be provided in clinical knowledge but can be found in data patterns.

4.3 Graph Characteristics

The graph obtained was:

- Heterogeneous, Multiple types of nodes and relationships.
- Graph, feature vectors, and attributed to each node.
- Thick but sparse, since edges are selected by a threshold.

This graphical representation allows efficient graph-based learning and facilitates extracting meaningful embeddings to downstream prediction tasks.

5. Proposed Methodology

This section shows the graph-based learning model that is proposed in clinical decision intelligence. The methodology uses the developed knowledge graph to acquire meaningful representations of clinical entities, and create individualized treatment suggestions. The general workflows include learning the graph representation, prediction of treatment and optimization of the model.

5.1 Graph Representation Learning

The suggested architecture uses a Graph Neural Network (GNN) to acquire low-dimensional vector representations (embeddings) of nodes in the clinical knowledge graph, as in Fig. 3. Given a graph $G = (V, E)$, each node $v \in V$ is associated with an initial feature vector derived from clinical attributes such as demographics, diagnosis codes, and medication history.

Let the node feature matrix be defined as:

$$X \in \mathbb{R}^{|V| \times d} \quad (3)$$

where $|V|$ denotes the number of nodes and d represents the feature dimension.

The GNN helps update the representations of nodes by iteratively aggregating the nodes in the neighborhood, enabling the node to understand the information held by its immediate neighbors. At layer k, the representation of node v is updated as:

$$h_v^{(k+1)} = \sigma\left(\sum_{u \in \mathcal{N}(v)} W^{(k)} h_u^{(k)} \right) \quad (4)$$

where $h_v^{(k)}$ denotes the embedding of node v at layer k, $\mathcal{N}(v)$ represents the set of neighboring nodes, $W^{(k)}$ is a trainable weight matrix, and σ is a nonlinear activation function such as ReLU.

The model combines local and global structural information by using many layers of and permits an encoding of contextual relationships in the clinical graph by each node. This is especially critical to formating interactions between diseases, symptoms and treatments, which are certainly relational in the nature.

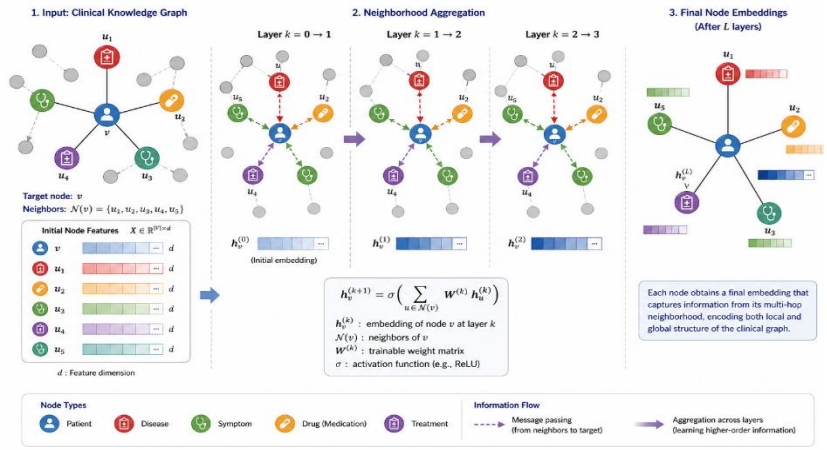


Fig. 3. Graph Neural Network-Based Node Representation Learning via Neighborhood Aggregation

5.2 Treatment Prediction Model

The trained node representations are used to make individualized treatment plans. The overall result of the GNN is a final embedding of a patient node, which combines both the specific clinical information of the patient and the network information based on the knowledge graph.

The prediction layer linearly predicts the learned embedding to a probability distribution of possible treatment choices with a softmax function:

$$\hat{y} = \text{softmax}(W h + b) \quad (5)$$

where h denotes the final embedding of the patient node, W is a learnable weight matrix, b is a bias vector, and $\hat{y} \in \mathbb{R}^C$ represents the predicted probability distribution over C treatment classes.

Such a formulation allows prioritizing treatment and choosing the most suitable intervention in the case of a particular patient. Through the use of relational information that is coded in the graph, the model makes more context-aware and personalized predictions compared to the traditional solutions.

5.3 Loss Function and Model Optimization

In order to train the proposed model, a cross-entropy loss is used to estimate the difference between the predictions of the treatment probability and the ground truth values. The loss function is given as:

$$L = -\sum_{i=1}^{N} y_i \log(\hat{y}_i) \quad (6)$$

where y_i represents the true label and $\hat{y}_i$ denotes the predicted probability for the i-th class.

This loss function is minimized through the training process, aiming at condensing the model parameters, the GNN weights as well as the prediction layer parameters. The parameters are updated iteratively using gradient-based optimization algorithms (stochastic gradient descent or Adam optimizer).

The model is trained to match patient embeddings to the correct treatment outcomes, thus enhancing prediction accuracy with time. Graph-based representation learning along with supervised optimization proves to be a successful method to incorporate the complex clinical relationships that can be converted into actionable clinically relevant treatment recommendations by the model.

6. Experimental Setup

This part explains the experimental design to determine the performance of the proposed knowledge graph-based clinical decision

intelligence system. It is performed at the cost of a real-world clinical dataset, where proper preprocessing methods and baseline models are used to make sure that the comparison can be fair and complete.

6.1 Dataset

The experiments are performed through the MIMIC-III (Medical Information Mart for Intensive Care) dataset that is a standard, publicly available benchmark dataset of healthcare studies. It includes health records of more than 40,000 patients who were admitted to intensive care units that were de-identified. The data set contains a variety of clinical data including patient demographics, diagnosis codes, test results, prescription of drugs and clinical notes. In these heterogeneous data sources, it is possible to have a valuable basis in building the clinical knowledge graph and testing the proposed model. The longitudinal data nature will enable the system to record trends over time and dynamic conditions of the patient, which are essential in optimizing treatment towards patients.

6.2 Data Preprocessing

The raw clinical data are then processed in a sequence of preprocessing operations before being trained into a model to guarantee quality and consistency of data. Any missing values that are prevalent in EHR data are managed with the use of suitable imputation methods to prevent bias and loss of information. Numerical attributes are standardized so that it exhibits the same scaling across variables, and thus, the model converges better at training. Categorical variables (such as diagnosis codes, medication identifiers, etc) are coded to numerical representations that can be consumed by machine learning models. The processed data is further converted to a graph form by mapping clinical entities to a graph node and providing relationship between two nodes. A graph of each patient is created, depending on their clinical history, and the model then provides an opportunity to acquire both relational and temporal dependencies.

6.3 Baseline Models

The results of the proposed approach are tested against various well-known baseline models that are widely applied in healthcare analytics to assess its performance. Among these are conventional machine learning algorithms like the Logistic Regression and the Random Forest that offer good baselines to classification problems. Further, deep learning models are also present to compare with, e.g., Long Short-Term Memory (LSTM) networks, which can be effectively used to model temporal sequences in case of clinical data, and RETAIN, an interpretable attention-based model aimed at healthcare prediction tasks. These standards reflect classical and state-of-the-art practices, and guarantee that the assessment will be most complete and significant.

6.4 Dataset Splitting Strategy

In order to achieve the stable assessment of the data, the dataset is separated into three parts: training, validation, and testing. In particular, 70 percent of the data is to be trained, 10 percent to be validated and 20 percent to be tested. Training set is applied to learn parameters of the model, whereas hyperparameter tuning and model selection are performed on the validation set. Testing set is invisible in the course of training, and is application only in the final examination of the model. This division will make the evaluation findings objective and applicable.

6.5 Hyperparameter Settings

The proposed model has a number of hyperparameters that determine its performance and are adjusted with the help of the validation set. The main hyperparameters of this study are:

- Learning rate: **0.001**
- Number of GNN layers: **2–3 layers**
- Hidden embedding dimension: **64**
- Batch size: **32**
- Number of training epochs: **50–100**
- Activation function: **ReLU**
- Optimizer: **Adam**

These values are chosen depending on empirical consideration to trade off between model performance and computational efficiency.

6.6 Evaluation Metrics

The proposed model is measured in terms of the classic classification metrics that are commonly used in healthcare analytics:

- Accuracy: The accuracy of predictions general across measures.
- Precision: Measures how often the positive cases are correctly predicted.
- Recall (Sensitivity): How much the true positive cases can be identified.
- F1-score: Gives an arithmetic average of the precision and recall.
- Area Under the ROC Curve (AUC): Tests the model in terms of the capacity of the model to differentiate among classes with varying thresholds.

These measures will present a broad analysis of the predictive performance of the model, especially when there is a lack of balance in the clinical data, where precision might not be adequate.

7. RESULTS AND DISCUSSION

Here, the performance analysis of the suggested clinical decision intelligence system based on a graph of knowledge and the analysis of the findings in comparison with basic models are detailed.

7.1 Performance Comparison

The performance of the proposed model is compared to that of various baseline models such as traditional machine learning and deep learning models, in terms of predictive performance. Table 1 summarizes the results.

Table 1. Performance Comparison of Different Models

Model	Accuracy	F1-score	AUC
Logistic Regression	0.78	0.75	0.80
Random Forest	0.82	0.80	0.84
LSTM	0.85	0.83	0.87
RETAIN	0.87	0.85	0.89
Proposed Model	0.91	0.89	0.93

The findings are clear evidence of the fact that the proposed model has the best performance of all the evaluation metrics. In particular, it has an accuracy higher than the RETAIN by about 4 percent and higher than the logistic regression by about 13 percent, which shows the effectiveness of integrating relational knowledge into the learning process.

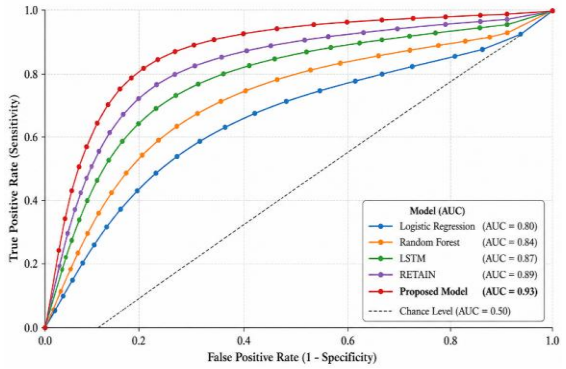


Fig.4. ROC Curve Comparison of Proposed and Baseline Models

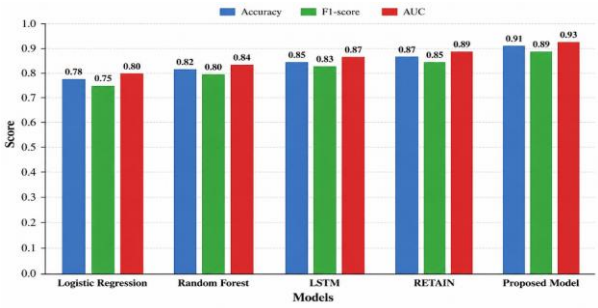


Fig.5. Performance Comparison Across Different Models

7.2 Result Analysis and Discussion

The high performance of the proposed model can be explained with several important factors. First, represented with the integration of a knowledge graph, relationships among clinical entities can be explicitly modeled, often being neglected in traditional machine learning and sequence deep learning methods. This relational representation enables the model to model dependencies among diseases, symptoms, and treatments in a better way. Second, Graph neural networks additionally improve the performance of learning context-aware embeddings, through aggregation of adjacent node information. In contrast to LSTM-based models which are mainly based on temporal relations, the model proposed captures both structural and contextual information which contributes towards a higher rate of predictions. Moreover, due to the presence of correlation-based edges within the knowledge graph, the model has the ability to discover latent associations not known in clinical records. This helps to enhance generalization and strength of the model. The proposed framework proves to be more effective when compared to the existing ones as the latter are mainly based on sequential or feature-based learning. The witnessed changes in AUC suggest increased capability to differentiate treatment results, especially significant in clinical decision support systems.

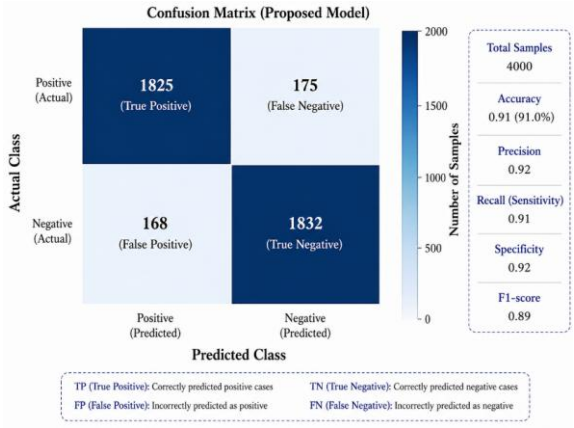


Fig.6. Confusion Matrix of the Proposed Model

7.3 Explainability and Clinical Interpretability

Besides predictive performance, the proposed system features increased interpretability which is a very important feature in a healthcare application. The structure of the knowledge graph can help to identify influential nodes and relationships that lead to model predictions. In particular, the model is capable of aiding explainability by allowing a process of examining node importance, during which it is possible to determine key clinical components (e.g., particular diseases or medicines) as key drivers of the

predicted treatment outcome. Also, the graph-based paths involved in reasoning will offer an insight on the relationship between various clinical factors and make the clinician know the reasoning behind the decisions made by the model. Furthermore, the GNN neighborhood aggregation mechanism implicitly places weights on the neighbors, which could be viewed as contextual relevance. This enables the system to indicate significant clinical relationships and hence greater transparency and confidence in the model.

7.4 Discussion Summary

On the whole, the experimental findings prove the hypothesis stating that the suggested knowledge graph-based method is much more effective than traditional and deep learning baselines. The structured graph representation and graph-based learning not only enhance predictive accuracy but also the interpretability of the system, which is appropriate in clinical decision support systems in the real world.

CONCLUSION

This paper has introduced a clinical decision intelligence system, based on knowledge graphs, aimed at maximizing the optimization of personal treatment processes through a high-quality modeling of the multifaceted relationships between clinical entities. The proposed framework by combining the representation of knowledge graphs with the learning of graph neural networks had an important impact on overcoming the major shortcomings of standard machine learning and deep learning models, including the inability to learn the relational nature of heterogeneous healthcare data. The graph of clinical knowledge constructed allowed both patients, diseases, symptoms, and medications to be represented in a single grid, whereas the graph neural network made it possible to extract context-sensitive embeddings via neighborhood aggregation. As the experimental analysis of the MIMIC-III dataset proved, the proposed model outperformed the traditional baseline algorithms, such as logistic regression, random forest, LSTM and RETAIN, in various performance measures, such as accuracy, F1-score, and AUC. These findings substantiate the power of integrating relation modeling and graph learning to enhance predictive accuracy with regard to clinical decision support systems. Moreover, the framework improves the interpretability as it makes it possible to establish the critical clinical relationships and lines of reasoning and, thus, facilitate clear and explicable decision-making. Although such positive results have been recorded, there are a number of new opportunities that can be utilized. Dynamism and ongoing clinical decision support may be possible through the integration of the proposed system with real-time hospital information systems. Moreover, its predictive power could be greatly enhanced by extending the framework to include multi-modal data information, including medical imaging and genomic information. Lastly, more complex explainable artificial intelligence methods may be introduced to enhance the transparency and clinical reliability of the system and make it more appropriate to implement in the clinical setting.

REFERENCES

1. D. Chang, I. Balažević, C. Allen, D. Chawla, C. Brandt, and R. Taylor, “Benchmark and best practices for biomedical knowledge graph embeddings,” in *Proc. 19th SIGBioMed Workshop Biomed. Lang. Process.*, 2020, pp. 167–176.
2. P. Chandak, K. Huang, and M. Zitnik, “Building a knowledge graph to enable precision medicine,” *Sci. Data*, vol. 10, no. 1, p. 67, 2023.
3. D. Niraula *et al.*, “A clinical decision support system for AI-assisted decision-making in response-adaptive radiotherapy (ARCliDS),” *Sci. Rep.*, vol. 13, no. 1, p. 5279, 2023.
4. S. Köhler *et al.*, “The human phenotype ontology in 2021,” *Nucleic Acids Res.*, vol. 49, no. D1, pp. D1207–D1217, 2021.
5. A. Barredo Arrieta *et al.*, “Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI,” *Inf. Fusion*, vol. 58, pp. 82–115, 2020.
6. Z. Q. Li *et al.*, “Prediction of diabetic macular edema using knowledge graph,” *Diagnostics*, vol. 13, no. 11, p. 1858, 2023.
7. O. Loyola-González, “Black-box vs. white-box: Understanding their advantages and weaknesses from a practical point of view,” *IEEE Access*, vol. 7, pp. 154096–154113, 2019.
8. J. S. Amberger, C. A. Bocchini, A. F. Scott, and A. Hamosh, “OMIM.org: Leveraging knowledge across phenotype–gene relationships,” *Nucleic Acids Res.*, vol. 47, no. D1, pp. D1038–D1043, 2019.
9. Q. Zhu *et al.*, “An integrative knowledge graph for rare diseases, derived from the genetic and rare diseases information center (GARD),” *J. Biomed. Semantics*, vol. 11, no. 1, p. 13, 2020.
10. C. Rudin, “Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead,” *Nat. Mach. Intell.*, vol. 1, no. 5, pp. 206–215, 2019.
11. N. Scarpato *et al.*, “Evaluating explainable machine learning models for clinicians,” *Cogn. Comput.*, vol. 16, pp. 1436–1446, 2024.
12. B. Séroussi*et al.*, “Reconciliation of multiple guidelines for decision support: A case study

on the multidisciplinary management of breast cancer within the DESIREE project," *AMIA Annu. Symp. Proc.*, pp. 1527–1536, 2018.
13. M. A. Shapiro *et al.*, "AI-augmented clinical decision support in a patient-centric precision oncology registry," *AI Precis. Oncol.*, vol. 1, pp. 58–68, 2024.
14. C. Fang *et al.*, "Integrating knowledge graphs into machine learning models for survival prediction and biomarker discovery in patients with non-small-cell lung cancer," *J. Transl. Med.*, vol. 22, no. 1, p. 726, 2024.
15. M. P. Hendriks *et al.*, "Clinical decision support systems for multidisciplinary team decision-making in patients with solid cancer: Composition of an implementation model based on a scoping review," *Crit. Rev. Oncol. Hematol.*, vol. 195, p. 104267, 2024.