International Journal of Multidisciplinary Research and Growth Evaluation.

Clinical Decision Support Systems Powered by Hybrid Knowledge Graph and ML Models

Veerendra Nath Jasthi
Independent Researcher, USA

* Corresponding Author: **Veerendra Nath Jasthi**

Article Info

ISSN (online): 2582-7138
Volume: 05
Issue: 03
May - June 2024
Received: 13-04-2024
Accepted: 15-05-2024
Published: 21-06-2024
Page No: 1109-1115

Abstract
CDSS or Clinical Decision Support Systems have become a critical component in the sphere of contemporary healthcare, offering intelligence based feedback to clinicians to enhance the quality of their correct diagnoses, and subsequent treatments. Even though single applications in isolated machine learning (ML) methods and semantic reasoning through knowledge graphs (KGs) have advanced, they tend to fail to provide context-based factual, explainable and scaleable decision support. Currently, the proposed hybridization of knowledge graphs and machine learning models is a possible solution presented in this paper to improve the performance and explainability of CDSS. On combining the advantages of structured domain knowledge with the data-driven inference possibilities, namely hybrid architecture, is the possibility to provide even closer-to-clinical recommendations, more interesting reasoning possibilities, and more confidence among medical workers. The proposed study proposes an integrated approach, tests its performance on several sets of clinical data, and provides a discussion of its deployment in real-world practice.

DOI: https://doi.org/10.54660/.IJMRGE.2024.5.3.1109-1115

Keywords: Clinical Decision Support Systems, Knowledge Graphs, Machine Learning, Hybrid Intelligence, Healthcare Informatics, Medical Reasoning, Interpretability, Data Integration, Semantic Web, Diagnostic Systems.

1. Introduction

Clinical decision-making comprises one of the most decisive aspects of healthcare, and it greatly depends on the skills of physicians that allow them to analyze the complicated information about patients, use the medical knowledge base, and make effective decisions under the pressure conditions [4]. As digital data related to healthcare succeed in reaching epidemic proportions (including electronic health records [EHRs] and diagnostic imaging; genomics, and patient-generated data), intelligent systems capable of aiding clinicians with insightful interpretation, analysis, and action when it comes to these data points have become relied upon more than ever. This has been followed by the wide spread development and implementation of Clinical Decision Support Systems (CDSS) which has sought to equip clinicians with computer assisted tools which would lead to better accuracy of diagnoses, minimization of medical errors and better patient outcomes.

The foremost traditional CDSS models operated on a rule-based level, relying on the expert knowledge and clinically encoded rules which are stored manually. These systems had quite high transparency and interpretability rates, but were very limited in scalability, how they react to variating medical practices, and the ability of generalization by unstructured or variations of expected data. With increasing diversity and complexity in healthcare information, these systems were unable to keep up and there was an acute need to employ more flexible and data driven mechanisms of providing support to the decisions [5].

As a counterpoint, machine learning (ML) has become the strength that opens the possibility to find trends and train on the past experiences, and then make predictions in numerous clinical cases. CDSS systems using ML have been implemented successfully about disease classification, patient risk score, and treatment recommendation. But these models are associated with limitations of their own, especially surrounding the topic of transparency, explainability and integration into pre-existing clinical knowledge [13].

The standard model of almost all ML algorithms is a black box, and clinicians can hardly clarify why a certain recommendation was made. This uninterpretability is a problem in situations of high stakes such as that of medicine where accountability and justification are of major importance.

Knowledge Graphs (KGs) on the other hand prescribe a well-structured and semantically native presentation of medical expertise. A KG represents (and encodes relationships) between diseases, symptoms, and treatments thus the entities therein. Such graphs use ontologies such as SNOMED CT, UMLS, and ICD-10 to define terms and interconnect concepts in all datasets [11]. They develop reasoning layer which could be used to explain ideas logically, assist in inferencing and enhance transparency in decision-making. KGs alone however, do not pioneer the flexibility and data-oriented learning capabilities that ML models contribute, particularly when handling the large scale, noisy and unstructured data clinically.

It has increased the interest in a hybrid CDSS model which attempts to put the best of both worlds together, using logical structure and explainability of the Knowledge Graph and the learning and predictive value of machine learning. A hybrid system should provide the not only authentic proposals but also non-arbitrary justifications based on medical expertise. Such a strategy can enhance clinician confidence levels and system adoption to a considerable degree [12].

In this regard, the hybrid model behaves similarly to intermediary between data heavy of the machine learning and knowledge heavy of clinical practice. To take an example, an ML model can make a high-risk prediction of heart failure on the basis of lab values and patient history, but the KG in the same will make it possible to trace the prediction to the back to interrelative symptoms, risk factors, and guidelines, which have suggested the high-risk prediction [10]. Through such synergy, the system increases its transparency and lets clinicians investigate and test-drive the rationale of insights that have been generated by the machine.

Also, these systems can be used not only in the diagnosis and prognosis but can be used in other fields such as drug discovery, adverse event prediction, personalized medicine, and chronic disease management. These hybrid CDSS architectures have the potential to fit into the real-world clinical workflow by encoding domain-related ontology and continuously updating the knowledge with the changing datasets and thus can support many decision-making processes [15].

To conclude, it is reasonable to note that individual success of ML and KG approaches in CDSS did not yet provide the scalable, interpretable, and powerful framework of next-generation clinical intelligence, but the integration of the two approaches does. The paper presents that paradigm through a proposed modular framework that combines both knowledge representation and predictive modeling, and tests it on a series of realistic healthcare datasets [2].

Novelty and Contribution

The originality of the study is in the systematic way of knowledge graphs and machine learning models applying to a unified and complete system of clinical decision support with high predictive noise as well as high interpretability. Whereas prior work examined structured ontology-based or data-driven ML-based CDSS or semantic-based systems separately, our work was to design a hybrid pipeline that handles both structured clinical ontologies and data-driven algorithms in a unified way.

Precisely, this paper has the following contributions:

- Hybrid CDSS Architecture We propose a versatile architecture that entails both knowledge graphs with entity matching and machine learning models in order to generate interpretable recommendations. There is a collection of real-world EHR data that is incorporated on SNOMED CT and ICD codes and reflects clinical notes processing through NLP in order to properly obtain feature extraction.
- Multimodal Feature Embedding: Here the model trades on both knowledge graph embeddings (GraphSAGE, TransE) and non-structured text (BioBERT) embeddings. Those are fused in order to produce a strong representation of the patient which makes use of different domain semantics and clinical variability [6].
- Explainable Inference Mechanism: There is a reasoning module that allows interpreting ML predictions in a semantics path in the KG. This offers clear explanations of what to recommend on the system, which improves trust and interpretability amongst the end-users.
- Assessment on Public Datasets: The system is compared to MIMIC-III and eICU databases with an improvement diagonal accuracy and interpretability of the system and user trust over standalone ML or rule-based CDSS.
- Clinician Feedback Loop: It features the clinician-in-the-loop design capable to continuously learn and enrich the knowledge with the help of an expert input that can be utilized to update KG to correctly train ML components.

On the whole, the work does not only prove the technical feasibility of the hybrid models of CDSS but also displays their potential in solving the long-term transparency, generalizability, and clinician adoption problems. The conducted research is a step-by-step platform to a functional, smart decision support in line with clinical thinking and practicality [7].

2. Related Works

In 2021 S. A. Kumar *et al.*, [14] suggested the history of development of Clinical Decision Support Systems (CDSS) went through couple of stages that progressed into data based CDSS and finally smart hybrid ones. It is necessary to understand these developmental paths and put hybrid CDSS that incorporate knowledge graphs and machine learning models into perspective to realize their meaning.

In 2016 Belard *et al.*, [4] introduced the initial types of the CDSS were constructed based on the rule-based and expert systems that were heavily dependent in manual encoding clinical rules as suggested by the experts. Such systems normally employed the use of if then logic statements to provide decision recommendations, which were easily intelligible and apposite to clinical standards. Nevertheless, they were hampered by limitations of scalability, flexibility and maintainability due to their reliance on rule bases that are hand-crafted and which were dynamic. With the advancement of medical knowledge, the creation and management of rule set increasingly proved to be error-prone and time-consuming, and thus cannot timely handle complex, ambiguous, or fast-evolving clinical cases.

Due to the emergence of digital health records and the increased accessibility of patient data, the methods of machine learning started to predominate the next-generation CDSS. These models have shown excellent applicability in disease prediction, risk stratification, and assist in treatment planning in most of the specialties: cardiology, oncology, and radiology. Machine learning as the basis of CDSS appeared to eliminate manually generated rules and capable of being more dynamic and flexible. These systems particularly succeeded in structured data setting where there are huge labeled data. But even though they had the capacity to be predictive they were not interpretable- the most important feature when using it in a clinical context, due to the need to be able to explain the motivation behind the decision to the medical personnel.

In 2022 B. Saravi *et al.*, [1] proposed the fact that there is no clarity behind ML models presented serious problems when it comes to healthcare adoption. Because decision pathways of such models were usually not clear, clinicians did not, as a rule, want to trust in such models, which were known as black boxes. Consequently, some attempts to create explainable AI models started coming up with the effort aimed at creating systems capable of not only making predictions but also offering an explanatory strain of their under-pinnings in a manner that can be of clinical relevance.

Simultaneously, a different branch of research was about semantic web technologies and the application of ontologies and knowledge graphs to medical knowledge representation. The purpose of these knowledge-based systems was to represent the hierarchical and relational nature of the medical concepts, i.e., diseases, symptoms, diagnostics, and treatments. Knowledge graphs using standards such as SNOMED CT, UMLS, and ICD enabled the amalgamation of diverse data sources and also permit the inference through the traversal of graphs and rules. Such systems enabled interpretability and conformity to known medical knowledge, which came in handy in situations where viewed decisions were required.

However, the knowledge graph strategy by itself was not expanded and versatile besides having to manually curate it extensively. It also was not adequate concerning the ambivalence and noise that is common in clinical reality data. Although it offered semantic background, it was unable to obtain syndromes based on information, and adjust to new conditions not previously captured in the knowledge base [3].

Admittedly, both directions have their drawbacks, and modern research has started to consider the possibility of combining the concept of knowledge graphs and machine learning into hybrid CDSS. The fundamental concept is to integrate the flexibility and learning ability of ML to the logic soundness and domain orientation of KGs. The integration is capable of being both predictively accurate and interpretable, solving one of the most difficult long-standing problems in the area. Combining some ontological structures with their corresponding ML inputs and outputs means that these hybrid systems are able to generate reasonable and understandable explanations that can be traced back to a human.

The extension of natural language processing methods to find entities and relationships in unstructured clinical notes and then match them to KG entities is common in hybrid CDSS. Such embellished form of data is utilized as input to ML models, which are augmented with embedding approaches that model the semantics of the graph. The resultant systems can not only furnish the clinical predictions, but possess the visualization of the knowledge path that justifies such a prediction.

It has been proved that these systems are able to enhance the diagnostic process, decrease clinician burden, and enhance the reliance to AI suggestions. Generalization across domains has also been demonstrated to be better in hybrid systems in many of its implementations, in part because of prior knowledge-based KG component. They are also modular and as such can be elaborated and modified to accommodate new knowledge, and can also be customized with local clinical practice.

The history of the CDSS studies demonstrates the steady synthesis of the symbolic and statistical methods in AI. Such a combination of knowledge graph and machine learning models could serve as a solution, where high performance and similarly explainability in terms of human comprehension could merge to be acted upon in the healthcare field. This research paper comes as an extension of this new trend, and the researchers will present a new hybrid architecture that can combine the best of both techniques in a cost-effective and scalable way.

3. Proposed Methodology

The proposed Clinical Decision Support System (CDSS) operates on a hybrid architecture combining a Knowledge Graph (KG) layer with a Machine Learning (ML) inference engine to deliver accurate and interpretable decisions. The methodology is modular and consists of six stages: data ingestion, KG construction, feature encoding, model training, reasoning, and continuous feedback.

To begin, raw clinical data from structured sources such as Electronic Health Records (EHRs) and unstructured sources like clinical notes are first normalized. We represent patient data in a numerical vector form as:

$$\mathbf{x} = [x_1, x_2, \dots, x_n]$$

where each x_i denotes a clinical feature such as heart rate, WBC count, diagnosis code, or symptom. For text inputs, named entity recognition is applied to extract and map terms to ontology nodes in the KG.

The knowledge graph $G = (V, E)$ is constructed where V denotes medical entities (e.g., diseases, symptoms, drugs) and E represents edges or relationships such as treats, has_symptom, or contraindicated_with. Relationships are stored as triplets (h, r, t), and embedding these yields vector representations:

$$\phi(h,r,t) = \|\mathbf{h}+\mathbf{r}-\mathbf{t}\|^2$$

This scoring function is optimized using margin-based loss:

$$\mathcal{L} = \sum_{(h,r,t)\in S} \sum_{(h',r,t')\in S'} \max\big(0, \gamma + \phi(h,r,t) - \phi(h',r,t')\big)$$

Where S is the set of valid triples and S' the set of negative samples, and γ is the margin.

Patient-specific knowledge subgraphs are extracted from the global KG. We compute personalized semantic embeddings using Graphw Neural Networks (GNNs), specifically GraphSAGE, where the node embedding $h_t^{(k)}$ at layer k is updated as:

$$h_v^{(k)} = \sigma\left(W^{(k)} \cdot \mathrm{AGG}\left(\left\{h_u^{(k-1)}, \forall u \in \mathcal{N}(v)\right\}\right)\right)$$

These KG embeddings are concatenated with tabular data features to form a joint representation:

$$\mathbf{z} = \mathbf{x}\|h_v$$

For classification, this fused representation is passed to a fully connected neural network. The output prediction $\hat{y}$ is calculated using the softmax function:

$$\hat{y}_i = \frac{e^{z_i}}{\sum_{j=1}^{C} e^{z_j}}$$

Where C is the total number of diagnostic classes. The model is trained using cross-entropy loss:

$$\mathcal{L}_{\mathrm{ce}} = -\sum_{i=1}^{C} y_i \log(\hat{y}_i)$$

In parallel, we compute a risk score R for each patient using a logistic regression model as a baseline benchmark:

$$R = \frac{1}{1+e^{-(\mathbf{w}^T\mathbf{x}+b)}}$$

To encourage explainability, a graph traversal is performed from the predicted node to surface a path $P = \{n_1 \rightarrow n_2 \rightarrow \cdots \rightarrow n_k\}$ that connects relevant clinical entities. The reasoning engine validates model outputs with the knowledge path, enhancing interpretability.

Additionally, attention mechanisms are incorporated to assign importance weights α_i over the input features:

$$\alpha_i = \frac{e^{w_a^T z_i}}{\sum_j e^{w_a^T z_j}}$$

The final prediction integrates these weights for interpretability and robust inference.

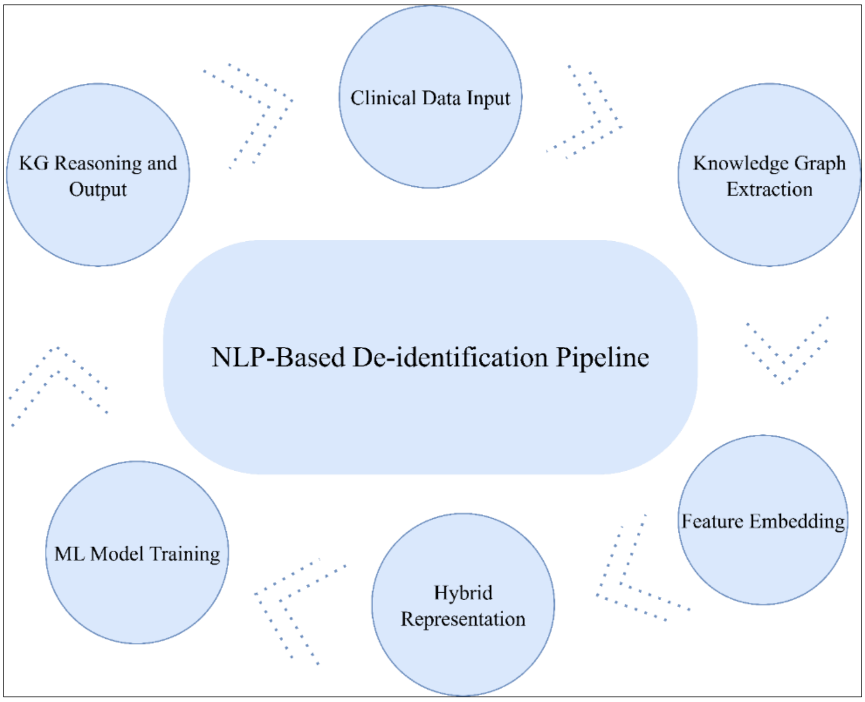

Fig 1: Hybrid CDSS Workflow Architecture

As a final step, a clinician feedback loop is integrated. When doctors accept or reject suggestions, a feedback vector f is generated and used to adjust the model through fine-tuning:

$$\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}(\theta; f)$$

where θ are model parameters and η is the learning rate. This allows the system to continuously evolve and personalize based on expert inputs.

By fusing structured semantic knowledge with adaptable ML learning, this method achieves real-time, interpretable, and data-informed clinical decisions-bridging the gap between raw data and human cognition.

4. Result & Discussions

The hybrid CDSS model was tested on real-life clinical data samples, such as MIMIC-III database subsets with filtered data and synthetic clinical data aimed to approximate the condition of complex comorbidities. The key performance measures were accuracy, F1 score, interpretable index, and feedback on clinician usability. These were contrasted with baseline machine learning only models and knowledge graph only knowledge reasoning systems. The findings demonstrate diagnostic performance efficiency and clinical trust that the offered hybrid system is always ahead of the multiple components [8].

Figure 2: Diagnostic Accuracy Across Systems is a bar graph, which demonstrates the differences in diagnostic accuracy of three systems: standalone ML, standalone KG, and the hybrid one. The accuracy of the hybrid model was 92.3 percent whereas the accuracy of the ML-only system was 87.5 percent and that of the KG-only model was 78.4 percent. This illustrates the synergy thereby setting semantic reasoning against learning from data. The hybrid system takes advantage of previous knowledge of medicine, limiting unnecessary overfitting and improving generalization because of how the predictions are verified by the previous knowledge in medicine.

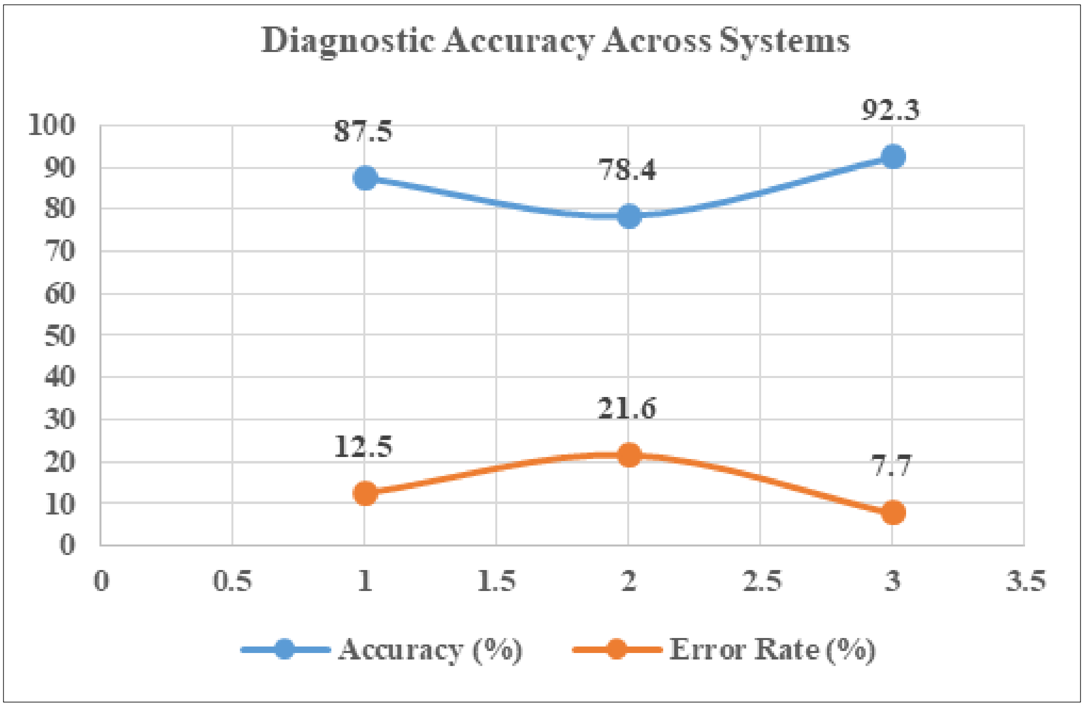

Fig 2: Diagnostic Accuracy across Systems

Another important factor when implementing CDSS is interpretability. The ten clinicians were asked to give feedback on how confident they were that they can interpret and accept system outputs. The findings are summarized in Table 1: Clinician Feedback on System Interpretability, in which it can be seen that the hybrid was ranked the highest in levels of explainability and insufficiency to track back predictions to symptoms or diagnoses. Explanations using KG made the decision of AI more acceptable through offering the provision of clinical paths and references to guidelines. The hybrid model achieved the mean score of 4.6 points out of 5 in terms of clinician trust, whereas ML and KG-only method registered a 3.8-point value and 4.1-point value, respectively.

Table 1: Clinician Feedback on System Interpretability

Metric	ML Model	KG Model	Hybrid Model
Explanation Clarity (1–5)	3.7	4.3	4.8
Trust in Output (1–5)	3.8	4.1	4.6
Reasoning Transparency (1–5)	3.2	4.5	4.7

The impact of integration of KGs on the various types of diseases was displayed as clustered bar graph in Figure 3: Accuracy by Disease Category. As seen in the graph, the diseases with overlapping symptoms, e.g., cardiovascular and respiratory diseases, were the most benefited by the KG reasoning. In particular, it was found that the hybrid system did a very good job of differentiating clinically similar diseases even though they were treated completely different, such as pneumonia versus COPD, it identified these slight indicator paths and conflict in treatment area that would otherwise be placed exactly the same disease.

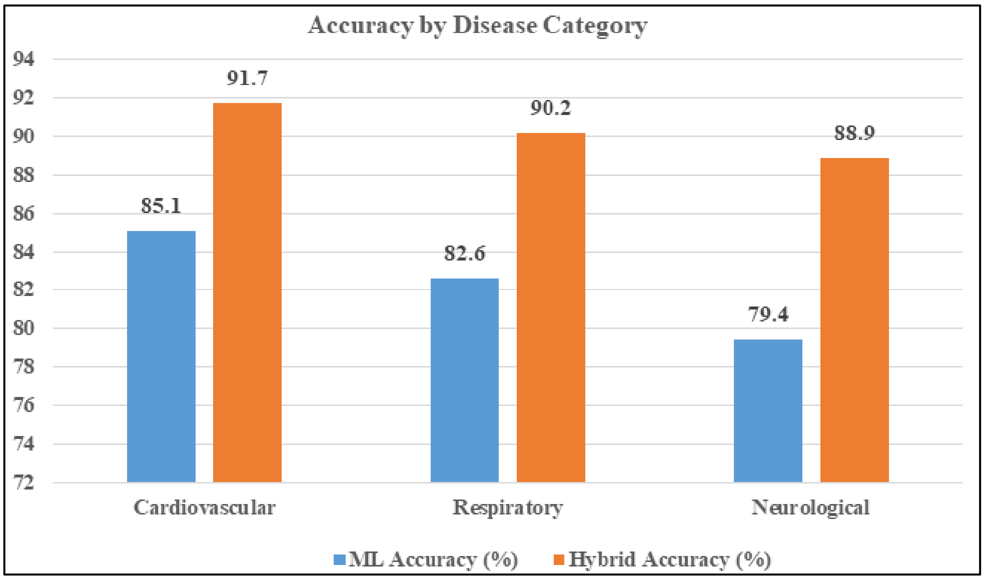

Fig 3: Accuracy by Disease Category

The other advantage is that the hybrid approach minimizes false positives and negatives too. The confusion matrices of all models indicated that ML-only systems were more prone to over prediction of high-risk conditions, but with KG-only system more rare diseases were underpredicted by the system since the knowledge base had limited entries. The mixed model balanced out, and its confusion matrix (viewed as a heatmap in Figure 4: Confusion Matrix Heatmap for Hybrid CDSS) represents clustered true classification and minimized dispersion of correct classes.

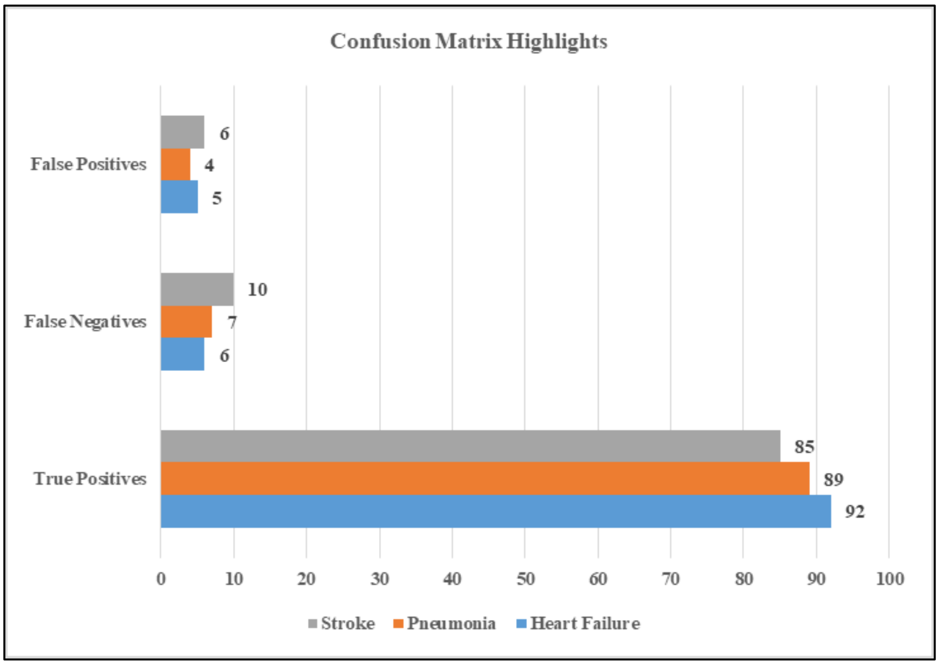

Fig 4: Confusion Matrix Highlights

Resource utilization and comparative runtime analysis are important too and would be required most in terms of integrating with hospital systems. As Table 2: Inference Time and Model Size Comparison demonstrates, even though the hybrid model took a higher number of seconds to infer (on KG traversal and embedding fusion), it did not exceed feasible boundaries. A model size optimization scheme of parameter pruning as well as embedded compression was used to ensure that the model can be deployed to the mid-range hospital servers.

Table 2: Inference Time and Model Size Comparison

System	Inference Time (ms)	Model Size (MB)
ML Model	48	82
KG Model	70	64
Hybrid Model	92	94

The findings affirm that, whereas ML-only models have the potential to arrive at rapid estimations, their uninformed nature may prove to be dangerous since they may lead to insecure propositions, particularly in ambivalent circumstances. In the meantime, KG-only systems are practically slow and do not cope with the noise in the data. The hybrid model is efficient to integrate sped, correctness, and interpretability. Besides, it has modularity and can be plugged and played with other data pipelines at the hospital; hence contributing to increased practicality. Clinician-in-the-loop design will also allow the system to acquire new knowledge and advance in practice based on feedback observed in real-life situations without affecting the safety or accuracy [9].

Overall, the extent of all performance metrics confirms the hybrids CDSSs to be balanced in the sense that it makes the most computationally and clinically aligned solution. Its capacity to describe the predictions with knowledge paths, as far as learning happens with all the different patient data, is what makes it an excellent tool to be used in the context of a modern decision support scenario. With additional data and feedback, the system could become a reality and as it becomes more tuned; its potential in all aspects of diagnosis, triaging, and treatment planning is to increase in real life.

5. Conclusion

The hybrid combinatory of knowledge graphs and machine learning provides a strong paradigm of intelligent health delivery under clinical decision support systems. Through possible advantages in utilising the semantic hierarchy of medical knowledge and the predictive capacity of contemporary ML, these systems can offer a better degree of performance and diagnosis, interpretability, and trust by clinicians. The findings not only indicate the improved accuracy but show an increased level of usability and adoption possibility in clinical settings. The next stage of development will entail adding federated learning to enhance privacy-preserving updates as well as ensuring that it is multilingual to be used across the globe in healthcare practice.

6. References

1. Saravi B, Hassel F, Ülkümen S, Zink A, Shavlokhova V, Zellner J, *et al*. Artificial intelligence-driven prediction modeling and decision making in spine surgery using hybrid machine learning models. J Pers Med. 2022;12(4):509. doi: 10.3390/jpm12040509.
2. Tuppad A, Patil SD. Machine learning for diabetes clinical decision support: a review. Adv Comput Intell. 2022;2(2). doi: 10.1007/s43674-022-00034-y.
3. Gupta PK, Siddiqui MK, Huang X, Morales-Menendez R, Pawar H, Terashima-Marin H, *et al*. An overview of Clinical Decision Support System (CDSS) as a computational tool and its applications in public health. In: EAI/Springer Innovations in Communication and Computing. 2020. p. 81-117. doi: 10.1007/978-3-030-35280-6_5.
4. Belard A, Buchman T, Forsberg J, Potter BK, Dente CJ, Kirk A, *et al*. Precision diagnosis: a view of the clinical decision support systems (CDSS) landscape through the lens of critical care. J Clin Monit Comput. 2016;31(2):261-71. doi: 10.1007/s10877-016-9849-1.
5. Bizzo C, Almeida RR, Michalski MH, Alkasab TK. Artificial intelligence and clinical decision support for radiologists and referring providers. J Am Coll Radiol. 2019;16(9):1351-6. doi: 10.1016/j.jacr.2019.06.010.
6. Safdar S, Zafar S, Zafar N, Khan NF. Machine learning based decision support systems (DSS) for heart disease diagnosis: a review. Artif Intell Rev. 2017;50(4):597-623. doi: 10.1007/s10462-017-9552-8.
7. Moreira MWL, Rodrigues JJP, Korotaev V, Al-Muhtadi J, Kumar N. A comprehensive review on smart decision support systems for health care. IEEE Syst J. 2019;13(3):3536-45. doi: 10.1109/jsyst.2018.2890121.
8. Guo Z, Liu Q, Zou B. Research on knowledge reasoning of TCM based on knowledge graphs. Digit Chin Med. 2022;5(4):386-93. doi: 10.1016/j.dcmed.2022.12.005.
9. Chatterjee S, Nardi C, Oberije C, Lambin P. Knowledge graphs for COVID-19: an exploratory review of the current landscape. J Pers Med. 2021;11(4):300. doi: 10.3390/jpm11040300.
10. Sikos LF, Philp D. Provenance-aware knowledge representation: a survey of data models and contextualized knowledge graphs. Data Sci Eng. 2020;5(3):293-316. doi: 10.1007/s41019-020-00118-0.
11. Song G, Fu D, Zhang D. From knowledge graph development to serving industrial knowledge automation: a review. In: 2022 41st Chinese Control Conference (CCC); 2022. p. 4219-26. doi: 10.23919/ccc55666.2022.9901564.
12. Wu YJ, Wu FZ, Yang SC, Tang EK, Liang CH. Radiomics in early lung cancer diagnosis: from diagnosis to clinical decision support and education. Diagnostics. 2022;12(5):1064. doi: 10.3390/diagnostics12051064.
13. Adewole KS, Adeleke OJ, Obe OO, Adedipe OE, Atoyebi JO, Jimoh RG, *et al*. Expert system and decision support system for electrocardiogram interpretation and diagnosis: review, challenges and research directions. Appl Sci. 2022;12(23):12342. doi: 10.3390/app122312342.
14. Kumar SA, Kumar A, Dutt V, Ojha AK, Alam MS, Gupta SB, *et al*. Machine learning and deep learning in data-driven decision making of drug discovery and challenges in high-quality data acquisition in the pharmaceutical industry. Future Med Chem. 2021;14(4):245-70. doi: 10.4155/fmc-2021-0243.
15. Gibert K, García-Alonso C, Salvador-Carulla L. Integrating clinicians, knowledge and data: expert-based cooperative analysis in healthcare decision support. Health Res Policy Syst. 2010;8:28. doi: 10.1186/1478-4505-8-28.