Journal of Artificial Intelligence for Medical Sciences
Vol 7 (1–4), 2026, pp. 01–12
Journal home: https://www.nexus-press.com/journal/JAIMS
DOI: https://doi.org/10.55578/joaims.260408.001; eISSN: 2666-1470

RESEARCH ARTICLE

Graph-Enhanced Medical Question-Answering System Integrating Knowledge Graphs and Large Language Models

Yiqin Huang[1,‡], Qing Liu[1,‡], Meiling Wang[2], Zhengru Xie[1], Baopeng Ding[3,4,†], Ziwei Wang[5,†], Fenglong Yang[2,6,7,*]

[1]*School of Medical Imaging, Fujian Medical University, Fuzhou, China*
[2]*Department of Bioinformatics, Fujian Key Laboratory of Medical Bioinformatics, School of Medical Technology and Engineering, Fujian Medical University, Fuzhou, China*
[3]*Shanxi Datong University, Engineering Research Center of Coal-Based Ecological Carbon Sequestration Technology of the Ministry of Education, Datong, China*
[4]*Boskind Research (Shanxi) Technology Co., Ltd., and Shanxi Yanyun Herbal Medicine Technology Co., Ltd.*
[5]*Biomedical Sciences College & Shandong Medicinal Biotechnology Centre, Shandong First Medical University & Shandong Academy of Medical Sciences, Ji'nan, Shandong, China*
[6]*The Fujian Key Laboratory of Medical Bioinformatics, Institute of Precision Medicine, Fujian Medical University, Fuzhou, China*
[7]*Key Laboratory of Gastrointestinal Cancer (Fujian Medical University), Ministry of Education, Fuzhou, China*

ARTICLE DATA

Article History

Received 25 October 2025
Revised 12 February 2026
Accepted 02 April 2026

Keywords
Knowledge graph (KG)
Graph algorithms
Large language model (LLM)
Intelligent question answering

ABSTRACT

Objective: To address the challenges posed by the rapid growth of medical data and the fragmentation of knowledge, this study aims to construct a medical knowledge graph (KG) and provide efficient knowledge services for clinical practice.
Methods: A total of 44,157 entities and 291,170 relationships from an open-source database were integrated to build a local medical KG based on Neo4j. Graph algorithms including degree centrality, Louvain community detection, K-nearest neighbor, and Dijkstra's algorithm were applied to analyze the data. The retrieval results from the KG were combined with the Spark Lite model from IFlytek to develop a dual-channel question-and-answer system.
Results: Highly related entities, such as acute urethritis and blood routine tests, were successfully identified. The analysis yielded 35 disease communities and 17 department communities. Highly similar disease pairs, such as "lung abscess" and "pulmonary bullae," were discovered. Potential therapeutic pathways, such as "Erythromycin Ethylsuccinate Granules - Erythrasma," were uncovered, revealing clinical associations among various entities. The system is accessible at http://yangbiolab.cn:8054/.
Conclusion: Graph algorithms effectively mine key patterns and potential associations within medical knowledge, with several findings aligning closely with clinical practice. The integrated system offers an intuitive platform for exploring medical knowledge.

1. INTRODUCTION

The rapid advancement of medical informatization has led to an era of data abundance while simultaneously exacerbating the challenges of knowledge fragmentation in healthcare. As electronic health records, biomedical literature, and multi-omics data continue to proliferate [1], traditional information processing paradigms struggle to effectively organize and utilize this vast wealth of clinical knowledge. In this context, knowledge graphs (KGs) have emerged as a powerful solution, offering structured representations of complex medical relationships—such as disease-symptom-treatment networks [2]—that significantly enhance both the interpretability and computability of healthcare information [3].

In recent years, graph-based techniques have advanced significantly, with innovations such as Graph Neural Networks (GNNs) [4], contrastive learning [5], and multimodal pre-training [6] pushing the boundaries of medical knowledge extraction and analysis. These developments have enabled the creation

---

[*] *Corresponding author. Email: yangfenglong@fjmu.edu.cn*
[†] *Co-corresponding authors. Email: dingbp@sxdtdx.edu.cn; amber4@yeah.net*
[‡] These authors contributed equally to this work.

of clinically complex applications, including intelligent medical question-answering (QA) systems. Currently, hybrid KG embedding systems [7], GNN-enhanced frameworks [4], and models integrated with pre-trained language models [8] are commonly employed to design QA systems. While these approaches have demonstrated promising task-specific accuracy [9] and have substantially improved the performance of large language models (LLMs) in QA evaluations [10], they still face challenges such as managing semantic inconsistencies across heterogeneous data sources.

The construction of medical KGs has advanced along two primary trajectories: general-purpose KGs and domain-specific vertical KGs. Since Google formally introduced the concept in 2012 [11], KG technology has been widely adopted in both academia and industry. These structured representations, composed of nodes (representing real-world entities) and edges (denoting semantic relationships), essentially form subject-predicate-object knowledge systems [12]. In the medical domain, specialized KGs such as Unified Medical Language System (UMLS) [13], Drug Repurposing Knowledge Graph (DRKG) [14], and DrugBank [15] have become invaluable resources, with nodes representing medical concepts (e.g., diseases, symptoms, drugs) and edges encoding clinical relationships (e.g., "treats") [16,17].

In this study, we developed a graph-enhanced medical QA system that integrates a structured KG with a LLM to enable comprehensive retrieval and analysis of medical information. The system's technological innovations include: (1) accurate identification of entities such as symptoms, diseases, and drugs; (2) construction of Cypher query language interfaces to bridge natural language and graph database queries; (3) implementation of Neo4j graph mining techniques, including degree centrality analysis to identify key medical entities, Louvain clustering to reveal disease associations, the K-nearest neighbor algorithm for similarity inference, and Dijkstra's algorithm to map drug action pathways; and (4) establishment of a dual-channel QA response system that combines the generative capabilities of iFlytek Spark Lite with the structuring mechanisms of the KG.

2. RELATED WORK

With the rapid advancement of information technology, knowledge mapping technology has evolved from a conceptual idea to a specialized professional field. In 2012, Google integrated the KG into its search engine, marking a significant shift in information retrieval—from traditional "string matching" to intelligent retrieval based on "entity association" [11]. In the medical domain, the development of KGs exhibits a parallel trend from standardization to localization. Internationally, the UMLS has established a standardized biomedical concept association network by integrating multilingual medical terms worldwide [13]. The DRKG focuses on the complex relationships among drugs, diseases, and targets, supporting drug development and repurposing efforts [14]. DrugBank centers on detailed medicinal chemistry and pharmacology data, serving as a crucial bridge between bioinformatics and chemoinformatics [15]. In China, several Chinese medical KGs have been developed to accommodate the local medical system. For instance, the Chinese Unified Medical Language System (CUMLS) addresses the standardization of Chinese medical terminology [18]. The China Medical Knowledge Graph (CMeKG) encompasses multiple entity types, including diseases, symptoms, drugs, and examinations, along with their interrelationships, thereby supporting clinical decision-making [19]. Additionally, the Traditional Chinese Medicine Language System (TCMLS) provides a vital foundation for the standardized mapping of Traditional Chinese Medicine (TCM) terms to modern medical terminology [20].

In the process of KG construction, optimizing ontology matching and modeling graph structure relationships are crucial. The adaptive feature construction technique based on genetic programming [21] can enhance the similarity measurement in ontology matching within KGs. The successful application of graph convolutional networks in image restoration and other fields [22] demonstrates the significant potential of graph structure learning for relational data modeling. Additionally, incentive mechanism design in federated learning [23] and decision variable analysis methods for large-scale multi-objective optimization [24] provide important algorithmic and theoretical foundations for building collaborative, efficient, and scalable distributed medical intelligent systems.

In parallel, intelligent QA systems have transitioned from rule-driven to data-driven approaches. Early systems, such as STUDENT [25] and MURAX [26], primarily relied on rule templates and were designed to solve query problems within specific domains. With breakthroughs in deep learning and pre-training techniques, QA systems gradually acquired the ability for deep semantic understanding. LLMs, exemplified by GPT and BERT, demonstrated strong language generation and contextual reasoning capabilities through massive text pre-training, driving the evolution of QA systems from an "information retrieval" to a "knowledge generation" paradigm. Considering the ambiguity and flexibility of the Chinese language, researchers have proposed various adaptive methods, such as encoder-decoder architectures based on attention mechanisms, which achieve precise mapping from natural language queries to machine-executable logic [27,28]. The KemQA framework effectively addresses the semantic parsing challenges posed by ambiguous expressions in Chinese through entity disambiguation and a structured query generation mechanism [29].

In recent years, the integration of KGs and LLMs has led to a new generation of medical QA systems. For example, Cao M.Y. et al. [30] constructed a KG for primary liver cancer and combined it with rule-based reasoning to enable automatic QA for specialized diseases. Chen Y.Y. [31] transformed natural language queries into graph database query languages to achieve precise retrieval from the KG for stroke. Additionally, some researchers have focused on Parkinson's disease management by using KGs to guide the consultation process and combining LLMs to provide personalized interactions, thereby enhancing the intelligence of chronic disease management [32]. These integrations not only leverage the advantages of KGs in structured knowledge representation but also utilize the capabilities of LLMs in natural language understanding and reasoning, driving the medical and health fields to evolve continuously toward deeper understanding, more accurate reasoning, and broader application scenarios.

3. MATERIALS AND METHODS

3.1. System Architecture

This intelligent QA system employs a three-layer fusion architecture (as illustrated in Figure 1) designed to facilitate deep integration between the medical KG and the LLM.

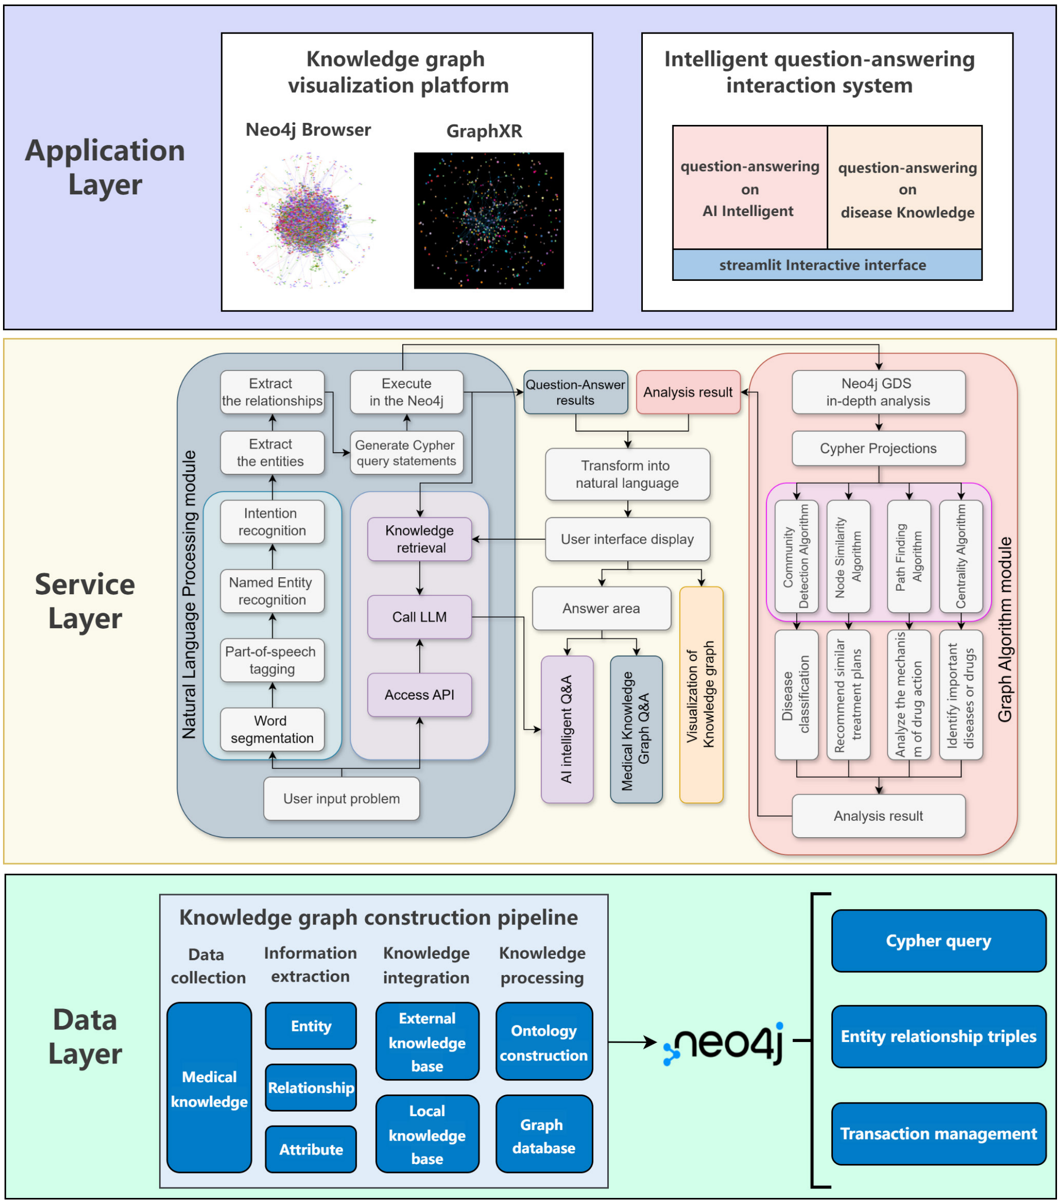

Figure 1 | System architecture diagram.

The data layer offers a structured knowledge foundation for the system and supports higher-level functions through a comprehensive KG construction process. This process involves key steps such as collecting multi-source medical data, extracting information, integrating knowledge, and modeling ontologies. Ultimately, the integrated local medical knowledge is persistently stored in the Neo4j graph database in a graph-structured format, ensuring both traceability and semantic consistency of the knowledge.

The service layer functions as the core computational engine of the system and comprises two main components: the Natural Language Processing (NLP) module and the Graph Algorithm module. The NLP module utilizes the Aho-Corasick algorithm (implemented via Pyahocorasick v2.1.0) to efficiently identify medical entities within user queries and classify user intentions. It converts these queries into structured Cypher queries to enable precise retrieval of relevant knowledge from Neo4j. Furthermore, this module integrates with the iFlytek Spark Lite API through the WebSocket protocol. Leveraging the structured evidence provided by the KG, it generates natural and fluent responses, effectively mitigating the inherent "hallucination" problem of the language model. The Graph Algorithm module supports a variety of efficient parallel graph analysis algorithms accessible via the Cypher interface, including degree centrality [33], Louvain [34], K-nearest neighbors [35], and Dijkstra's shortest path algorithm [36]. These algorithms provide computational support for clinical reasoning, relationship mining, and decision assistance.

The application layer features an interactive quality assurance interface developed using Streamlit (v1.43.2), which offers a dual-channel response display. This allows users to simultaneously view answers generated by the LLM alongside the original results retrieved directly from the KG, facilitating comparison, verification, and trustworthy assessment. Additionally, the system integrates the dynamic 2D/3D exploration environments of Pyvis (v0.3.2) and GraphXR, enabling users to intuitively navigate and understand complex medical relationships.

3.2. Data Layer

The data layer is constructed using an open-source dataset of the Chinese medical KG, which is publicly available. This dataset originates from the GitHub project: https://github.com/liuhuanyong/QASystemOnMedicalKG/. It contains a total of

44,157 entities spanning seven categories, including examinations, departments, diseases, drugs, foods, drug manufacturers, and symptoms (see Table S1 in Supplementary Information). Additionally, the dataset defines 291,170 relationships that capture various medical correlations such as causes, symptoms, treatments, drug interactions, dietary restrictions, recommended foods, department affiliations, and complications (see Table S2 in Supplementary Information). Among these, disease entities include nine attributes—such as name, cause, and cure probability—enabling rich semantic representation (see Table S3 in Supplementary Information).

During the construction process, the original data was standardized, including the alignment of medical terms and entity disambiguation based on dictionary matching. Clinical standard terminology systems, such as ICD-10, were also referenced, effectively alleviating the problem of term ambiguity in Chinese medical texts and ensuring consistency between entities and actual clinical concepts. Subsequently, using the py2neo Python library (compatible with Neo4j community version v5.26.0), the standardized data was programmatically imported into the graph database. Systematic Cypher query commands were executed sequentially to complete node creation, relationship definition, and attribute assignment, ultimately constructing a clear metadata model (see Figure S1 in Supplementary Information). The complete medical KG supports interactive browsing and is accessible online at http://yangbiolab.cn:8054/.

3.3. Service Layer

3.3.1. Natural Language Processing

The NLP module serves as the core component of the service layer, responsible for interpreting and analyzing users' natural language queries. This system utilizes multiple layers of medical dictionaries—such as those for diseases, departments, drugs, drug manufacturers, and food—to identify medical entities within user queries. An Aho-Corasick tree (ACTree) was constructed using the Pyahocorasick library to enable efficient keyword matching, which is combined with entity disambiguation logic. The identified entities are then classified by type (e.g., disease, symptom, or department) for subsequent processing.

User intent, such as querying disease symptoms, causes, or preventive measures, is determined through recognized entities and predefined interrogative terms. The workflow includes the following steps:

1) Predefined medical questions: A set of medical question indicators has been established (see Table S4 in Supplementary Information).
2) Question Matching: Users input terms that correspond to predefined keywords to identify intent.
3) Cypher Query Generation: The intent classification results guide the targeted generation of Cypher queries for knowledge retrieval.

3.3.2. Graph Data Science Library

The Neo4j Graph Data Science (GDS) library provides an efficient, parallel graph algorithm framework accessible via a Cypher statement-based interface. It supports a wide range of graph computing tasks, including centrality analysis [33], community detection [37], node similarity calculation [38], path mining, and node embedding [39].

The algorithm formula used in this study is defined in Equations (1) to (4) (see Figure S2 for a legend explanation of the algorithm in Supplementary Information):

(1) Degree centrality identifies the most connected nodes in the network:

$$C_d(i) = \deg(i) \tag{1}$$

(2) Louvain algorithm iteratively merges nodes and communities to maximize modularity through a greedy strategy to discover the natural community structure in the network [37]:

$$Q = \frac{1}{2m}\sum_{i,j}\left[A_{ij} - \frac{k_i k_j}{2m}\right]\delta(c_i, c_j) \tag{2}$$

(3) K-nearest neighbors quantifies the semantic similarity between entities based on feature proximity:

$$y = \arg\max_{c}\sum_{i=1}^{K} I(y_i = c) \tag{3}$$

(4) Dijkstra's shortest path algorithm finds the shortest association path between entities in a weighted graph:

$$\text{if } \mathrm{dist}[u] + w(u, v) < \mathrm{dist}[v], \text{ then } \mathrm{dist}[v] = \mathrm{dist}[u] + w(u, v) \tag{4}$$

GDS offers two primary types of projections: native projection and Cypher projection. The native projection prioritizes efficiency and performance optimization for large-scale graph data science, whereas the Cypher projection emphasizes flexibility and customization. In this study, the Cypher projection method was employed to project multiple nodes and relationships for subsequent graph algorithm analysis. Prior to conducting the graph algorithm analysis, the system generated graph projections using Cypher technology, enabling repeated references through graph algorithms without the need for re-projection, thereby enhancing computational efficiency. This study constructed dedicated projections for different analytical dimensions, such as the disease multi-connection graph, drug multi-connection graph, and department multi-connection graph. For example, the disease multi-connection graph projection contains 26,948 nodes and 273,994 relationships, encompassing connections between diseases and all related entities (symptoms, examinations, drugs, departments, foods, etc.), thereby laying the foundation for subsequent relationship mining centered on diseases.

3.4. KG–LLM Integration Mechanism

The integration of KG and LLM is accomplished through a dual-channel response mechanism. When a user submits a query, the system performs two operations simultaneously:

First, the system identifies medical entities using the NLP module and automatically constructs a Cypher query to retrieve structured medical fact data from the Neo4j KG, providing an accurate answer based on evidence-based medicine. Simultaneously, the original query is transmitted via the WebSocket protocol to the IFlytek Spark Lite model, which generates a natural language response that includes clinical context, mechanistic explanations, and precautions.

The system displays both types of responses side by side in the Streamlit interactive interface, creating a contrast between "structured knowledge" and "generative explanations". This design not only ensures the accuracy of medical information, by leveraging the KG to mitigate potential hallucination issues associated with LLMs, but also enhances the completeness and clarity of the explanations, enabling users to simultaneously access core facts and relevant background information.

This hybrid mechanism effectively combines the rigor of KGs with the semantic flexibility of LLMs. It enhances the reliability of answers while supporting in-depth interpretation of complex medical queries, providing a feasible and effective approach to developing a medical intelligent QA system that is both accurate and user-friendly.

4. GRAPH ALGORITHM ANALYSIS RESULTS

4.1. Degree Centrality Reveals Key Clinical Entities

The degree centrality analysis revealed that acute urethritis (score: 66) had the densest connections in the KG, consistent with the disease's characteristic presentation of multiple systemic symptoms and complications in urological clinical practice. Other highly correlated entities included blood routine examination (score: 2127) as the most frequently associated test, dermatology (score: 647) as the most relevant department, abdominal pain (score: 1001) as the most common symptom, and levofloxacin hydrochloride capsules (score: 708) as the most commonly used medication. Regarding lifestyle factors, eggs (score: 2135) and liquor (score: 2754) were identified as the most recommended food and the most prohibited substance, respectively, highlighting the clinical importance of dietary factors in disease management (see Tables S5 and S6 in Supplementary Information).

Further visual analysis (Figure 2) reveals a dense network of connections between acute urethritis and various antibiotics, symptomatic treatments, and dietary recommendations. This structural feature provides a graphical foundation for comprehensive treatment decisions for this disease. The identification of these highly correlated entities not only confirms the consistency between the KG and actual clinical diagnosis and treatment models but also highlights, from a data perspective, the diagnostic pathways, core medications, and health

Figure 2 | Connection diagram of acute urethritis.

management elements that should be prioritized in clinical practice. This, in turn, enhances the clinical credibility and interpretability of the analysis results.

4.2. Louvain Community Survey Reveals the Structure of Clinical Specialties

Community detection within the medical KG was performed using the Louvain algorithm. Disease entities were grouped into 35 communities, while department entities formed 17 communities. Additionally, multiple other entity communities were identified, including examinations (3,128), drugs (2,034), foods (4,308), and symptoms (5,378). The analysis revealed that the disease and department communities exhibited a tightly coupled modular structure (Figure 3), indicating a significant clinical correlation.

Specifically, the clustering of departments and communities demonstrates multi-level clinical significance. The internal medicine specialties (such as cardiology and respiratory medicine) form a closely knit group, reflecting a collaborative management model for chronic disease diagnosis and treatment. The relationships between specialties and subspecialties (such as pediatrics and pediatric internal medicine) exhibit a clear hierarchical structure, embodying the specialized division of labor characteristic of the modern medical system. Additionally, the spatial proximity of the psychological and internal medicine departments reveals an integrated approach to managing psychosomatic comorbidities. In parallel, the parallel distribution of the traditional Chinese medicine and Western medicine departments highlights the clinical practice trend of integrating traditional and Western medical approaches. Notably, the emergency, rehabilitation, and other hub departments are situated at the intersections of multiple communities, aligning closely with their roles in triage, referral, and cross-specialty coordination during diagnosis and treatment.

These findings confirm, at the structural level, the representational capacity of the KG for real-world clinical relationships. The interactive visualization, implemented using vis.js, enables dynamic exploration of the association network between diseases and departments. This provides a visual tool for quickly identifying differential diagnosis pathways and optimizing referral processes, thereby enhancing the practical value of the KG in clinical decision support.

4.3. K-Nearest Neighbors Reveal Disease Similarities

To quantify potential associations between disease entities, this study first generated low-dimensional vector representations based on disease attribute information using the FastRP node embedding algorithm. Similarity scores between disease pairs were then calculated using the K-nearest neighbor algorithm. The analysis results (see Table S7 in Supplementary Information) revealed that some disease pairs exhibited highly similar characteristics. For example, the similarity between "pulmonary abscess" and "pulmonary bullae" was 1.00, as was the similarity between "acute respiratory failure" and "acute respiratory failure in children", indicating that these diseases may share significant commonalities in their pathological mechanisms, imaging features, or clinical manifestations, providing a theoretical basis for developing integrated diagnostic and treatment strategies across different ages and disease types. Additionally, the complete similarity between "mercury poisoning" and "iron poisoning" offers insights for developing broad-spectrum antidotes or unified toxicological intervention plans.

Furthermore, the analysis identified pairs of diseases with extremely low similarity, such as “carcinoid” and “otogenic hydrocephalus”, with a similarity score of 0.00. These findings are valuable for distinguishing disease combinations with markedly different clinical manifestations and those that pose challenges in differential diagnosis. This, in turn, helps clinicians reduce the risk of misdiagnosis and improve diagnostic accuracy. The similarity analysis not only uncovers implicit association patterns between diseases at the data level but also provides structured data support for clinical applications such as early disease screening, precise differential diagnosis, and drug repositioning. It enhances the interpretability and practical utility of the KG in supporting clinical decision-making.

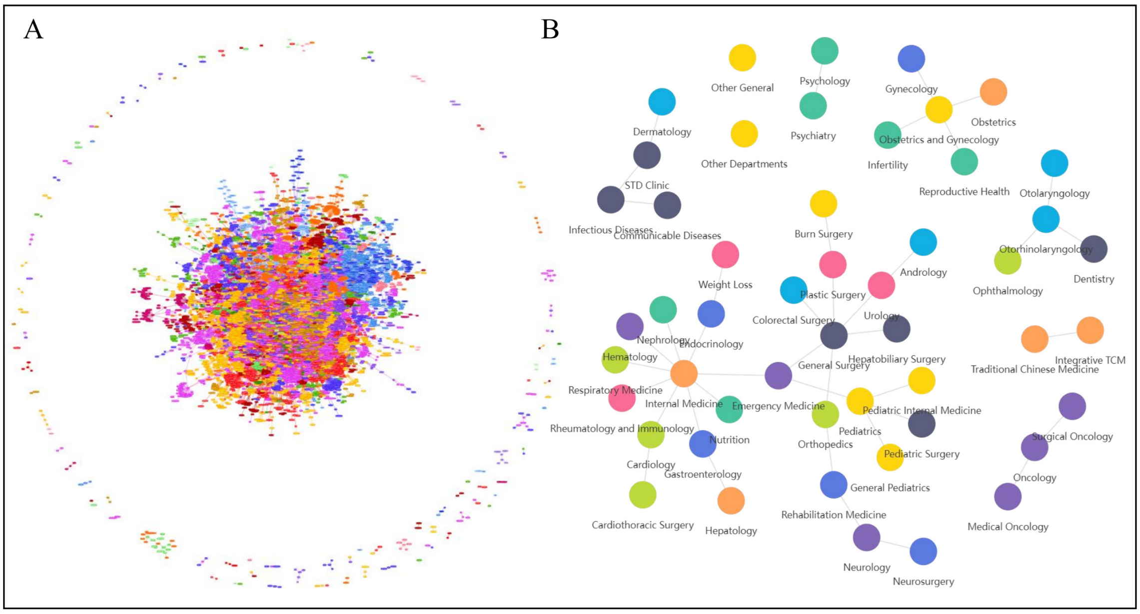

Figure 3 | Community detection visualization || A: Disease community visualization B: Department community visualization.

4.4. Dijkstra Explores Therapeutic and Nutritional Pathways

This study utilized Dijkstra's single-source shortest path algorithm to systematically analyze the association pathways among entities such as diseases, drugs, departments, and foods. In exploring the drug-disease pathways, analysis using "Erythromycin Ethylsuccinate Granules" as the source node revealed an association with "Erythrasma" via the path "Whooping cough→ Erythromycin Estolate Tablets" (Figure 4). This finding suggests that the drug may have potential for cross-indication treatment and provides evidence for the repurposing of antibacterial drugs.

In analyzing departmental relationships, the "Pediatric Internal Medicine" and "Pediatric Surgery" sections exhibit closely adjacent structural features within the path network, reflecting the natural clinical collaboration and referral connections between pediatric subspecialties. This finding provides a data-driven basis for optimizing the integration of pediatric medical resources and developing an efficient, collaborative pediatric diagnosis and treatment system.

In analyzing the nutritional support pathways for diseases, the shortest path analysis starting from "Eggs" reveals a direct connection to certain disease entities and further extends to specific dietary recommendations such as "Sea Cucumber and Duck Congee". This finding clearly demonstrates the structured guidance provided by the KG in disease nutrition management and lifestyle interventions.

The path analysis results above reveal potential semantic connections and clinical logic among medical entities from multiple dimensions. They not only offer new insights for expanding disease treatment strategies and optimizing departmental collaboration models but also establish a methodological foundation for constructing a knowledge-based health management support system.

5. INTELLIGENT QUESTION-ANSWERING SYSTEM

This study developed an intelligent QA system that integrates KGs and LLMs using the Streamlit framework. The system is accessible at http://yangbiolab.cn:8054/ and features a KG visualization interface available at http://yangbiolab.cn:7475/browser/. It consists of three modules: a system overview, a QA interface, and a KG exploration tool, arranged in a left-right layout with a sidebar and an interactive panel.

5.1. System Introduction

First, create an introductory interface for the intelligent QA system using the Streamlit framework, as shown in Figure S3 in the Supplementary information. Users can access both the Q&A interface and the KG through the sidebar tabs to interact with the system.

5.2. Question-Answering Interface

The intelligent QA interface enables users to input queries and provides real-time interaction, facilitating users' efficient acquisition of medical knowledge. The system employs two feedback mechanisms: one based on the artificial intelligence intelligent assistant (Spark Lite model), and the other based on the structured medical knowledge base. Currently, the system can respond to 18 types of medical questions, such as symptom inquiries, cause searches, and treatment cycle consultations (see Table S8 in Supplementary Information). According to the partial question-answer results displayed in Table 1, for the same question, such as "How can meningitis be diagnosed?", the AI intelligent assistant's answers cover a wide range and elaborately explain the definitions and procedures of various examination methods. However, there is a risk of factual hallucination, and

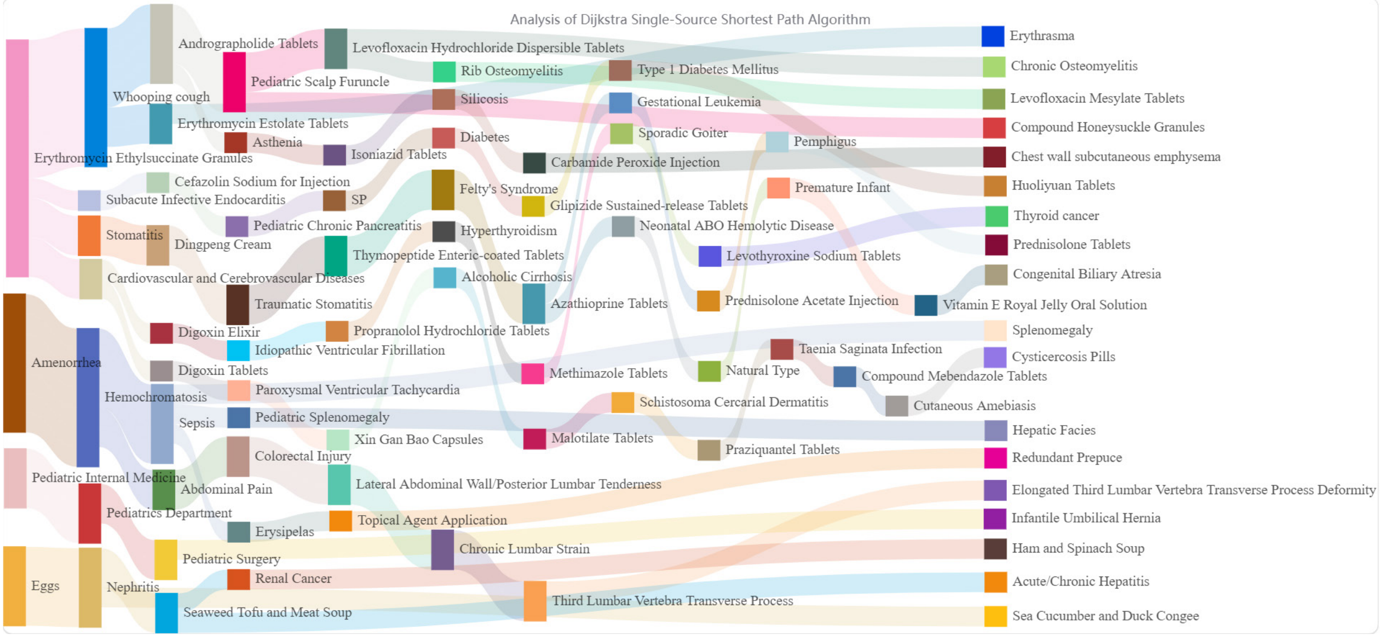

Figure 4 | Sankey diagram based on Dijkstra.

Table 1 | **Comparison of two-channel QA results**

Question	Medical knowledge base	AI intelligent assistant
How can meningitis be detected?	Meningitis can usually be detected by: Cerebrospinal fluid potassium; Urinary glutamine; urinary routine; cerebrospinal fluid bacterial culture; cerebrospinal fluid sodium; Fisher's finger test; cerebrospinal fluid aspartate aminotransferase; neck rigidity; nitrate reduction test; Examination of cerebrospinal fluid pathogens.	Meningitis is a serious disease that requires prompt diagnosis and treatment. Here are some commonly used meningitis examination methods: 1. Medical history inquiry: The doctor will ask about your medical history, symptoms, family history, etc., to know if you have risk factors for meningitis. 2. Physical examination: The doctor will perform a physical examination of the nervous system, including looking at your pupil response, muscle strength, sensation, reflexes, balance, etc. 3. Laboratory tests: Your doctor may ask you to do blood tests, urine tests, cerebrospinal fluid tests, etc. to determine if there is an infection or other abnormality. 4. Imaging tests: Your doctor may order a CT scan or MRI of your head to determine if there is damage to brain tissue caused by meningitis. 5. Cerebrospinal fluid test: Your doctor may draw cerebrospinal fluid for testing to determine if there is a bacterial infection or other abnormalities. In conclusion, if you suspect you have meningitis, seek medical attention promptly and receive diagnosis and treatment from a professional doctor.
Can leukemia be cured?	The probability of curing leukemia is (for reference only): 50–70%.	Leukemia is a malignant tumor that usually requires comprehensive treatment. Treatment methods include chemotherapy, radiotherapy, hematopoietic stem cell transplantation, etc. Cure rates are higher for patients diagnosed early, while lower rates for patients with advanced or relapsed stages. However, with the continuous advancement of medical technology, the treatment of leukemia has made great progress, and many patients have been able to survive for a long time and return to normal life. Therefore, if you are diagnosed with leukemia, it is important to actively receive treatment while maintaining a good mindset and lifestyle habits to help the body cope better with the disease.

it fails to establish an explicit association with the KG, making it difficult to verify the credibility of the information source. Responses based on the medical knowledge base directly rely on the structured data in the graph, clearly listing examination items such as "Cerebrospinal fluid potassium test" and "Urine glutamine determination". Their contents can be traced back to the entities and relationships in the KG, possessing good interpretability. However, such responses are relatively concise, lacking fluency in language expression and thoroughness in explanation.

Therefore, both methods have their own advantages and disadvantages, but they can complement each other to improve the effectiveness of the QA process. Moreover, the answers based on the knowledge base can all be traced back to the KG itself, reducing the occurrence of hallucinations and helping users acquire more accurate medical knowledge.

Figure 5 further illustrates the performance differences of the two mechanisms in real conversation scenarios. The user repeatedly raised the same question in two separate windows: "Which department should I go to for treatment of benzene poisoning?" As shown in Figure 5A, the response generated by the AI is more comprehensive, covering the background of the question, analysis, and preventive measures. In contrast, the answer in Figure 5B, based on the knowledge base, is more direct and highly aligns with the relevant KG.

5.3. KG Exploration

The KG interface comprises three modules: KG introduction, visual exploration, and Neo4j query. Users can switch between these modules by clicking the function buttons, facilitating efficient exploration of the KG.

Figure S4 in the Supplementary Information illustrates that the KG page consists of three modules: KG introduction, visual exploration, and Neo4j query. Users can switch between these modules by clicking the function buttons. This system supports the exploration of KGs using Neo4j and GraphXR tools, and uses graph algorithms for medical problem analysis. For example, Figure 6 is a local visualization of the association of hypertension diseases, showing that hypertension belongs to the cardiology department and is closely related to CT examination results and renal function tests.

In the intelligent QA system, users can conduct query exploration via Neo4j. In the case shown in Figure 7, the user entered the following query: "What medicine is used for diabetes?" After the "Generate and Execute Query" button is clicked at the lower right corner, the system successively outputs the query results based on Cypher as shown in Figure 7B and the query table based on local data and the KG as shown in Figure 7C, before finally presenting the visualization result, as shown in Figure 7D. The final result indicates that diabetes can be treated with drugs

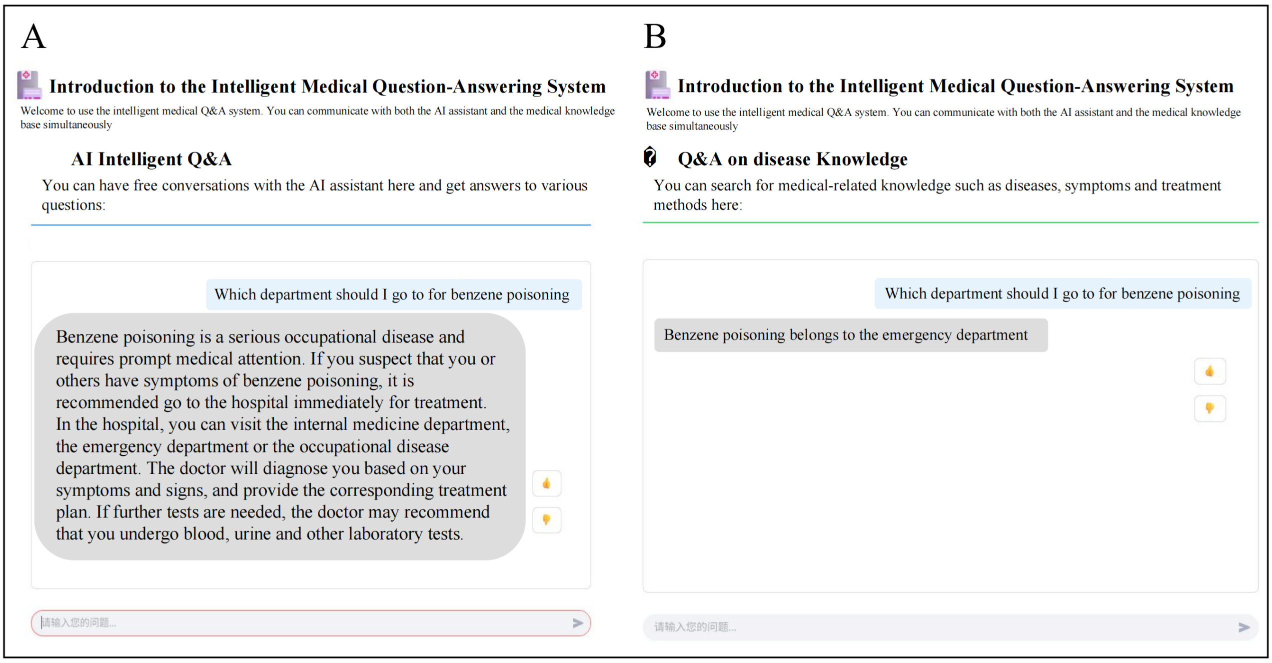

Introduction to the Intelligent Medical Question-Answering System

Welcome to use the intelligent medical Q&A system. You can communicate with both the AI assistant and the medical knowledge base simultaneously

AI Intelligent Q&A

You can have free conversations with the AI assistant here and get answers to various questions:

Which department should I go to for benzene poisoning

Benzene poisoning is a serious occupational disease and requires prompt medical attention. If you suspect that you or others have symptoms of benzene poisoning, it is recommended go to the hospital immediately for treatment. In the hospital, you can visit the internal medicine department, the emergency department or the occupational disease department. The doctor will diagnose you based on your symptoms and signs, and provide the corresponding treatment plan. If further tests are needed, the doctor may recommend that you undergo blood, urine and other laboratory tests.

Introduction to the Intelligent Medical Question-Answering System

Welcome to use the intelligent medical Q&A system. You can communicate with both the AI assistant and the medical knowledge base simultaneously

Q&A on disease Knowledge

You can search for medical-related knowledge such as diseases, symptoms and treatment methods here:

Which department should I go to for benzene poisoning

Benzene poisoning belongs to the emergency department

Figure 5 | An example of interactive question-answering interface.

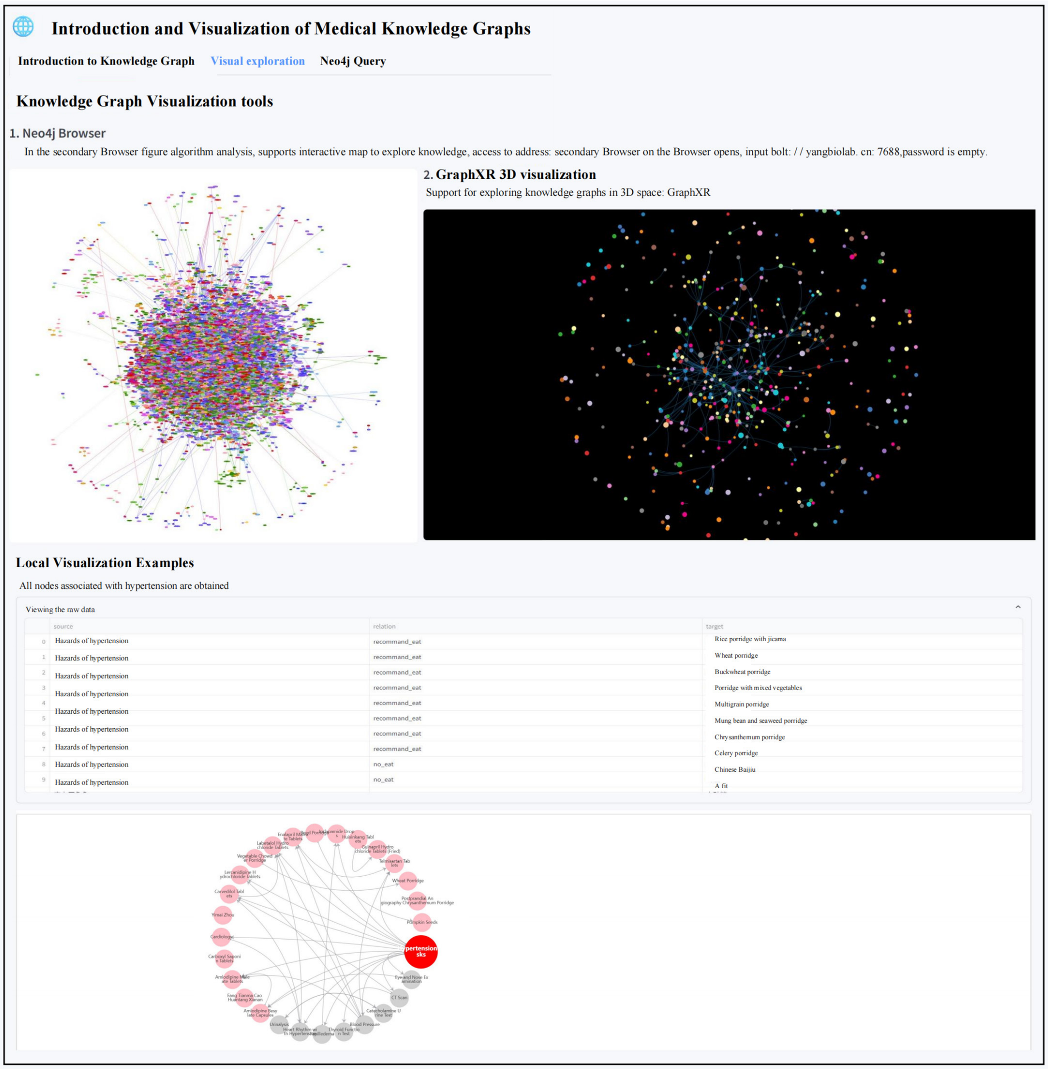

Introduction and Visualization of Medical Knowledge Graphs

Knowledge Graph Visualization tools

1. Neo4j Browser

In the secondary Browser figure algorithm analysis, supports interactive map to explore knowledge, access to address: secondary Browser on the Browser opens, input bolt: / / yangbiolab. cn: 7688,password is empty.

2. **GraphXR 3D visualization**

Support for exploring knowledge graphs in 3D space: GraphXR

Local Visualization Examples

All nodes associated with hypertension are obtained

Viewing the raw data

	source	relation	target
0	Hazards of hypertension	recommand_eat	Rice porridge with jicama
1	Hazards of hypertension	recommand_eat	Wheat porridge
2	Hazards of hypertension	recommand_eat	Buckwheat porridge
3	Hazards of hypertension	recommand_eat	Porridge with mixed vegetables
4	Hazards of hypertension	recommand_eat	Multigrain porridge
5	Hazards of hypertension	recommand_eat	Mung bean and seaweed porridge
6	Hazards of hypertension	recommand_eat	Chrysanthemum porridge
7	Hazards of hypertension	recommand_eat	Celery porridge
8	Hazards of hypertension	no_eat	Chinese Baijiu
9	Hazards of hypertension	no_eat	A fit

Figure 6 | Knowledge graph visualization platform.

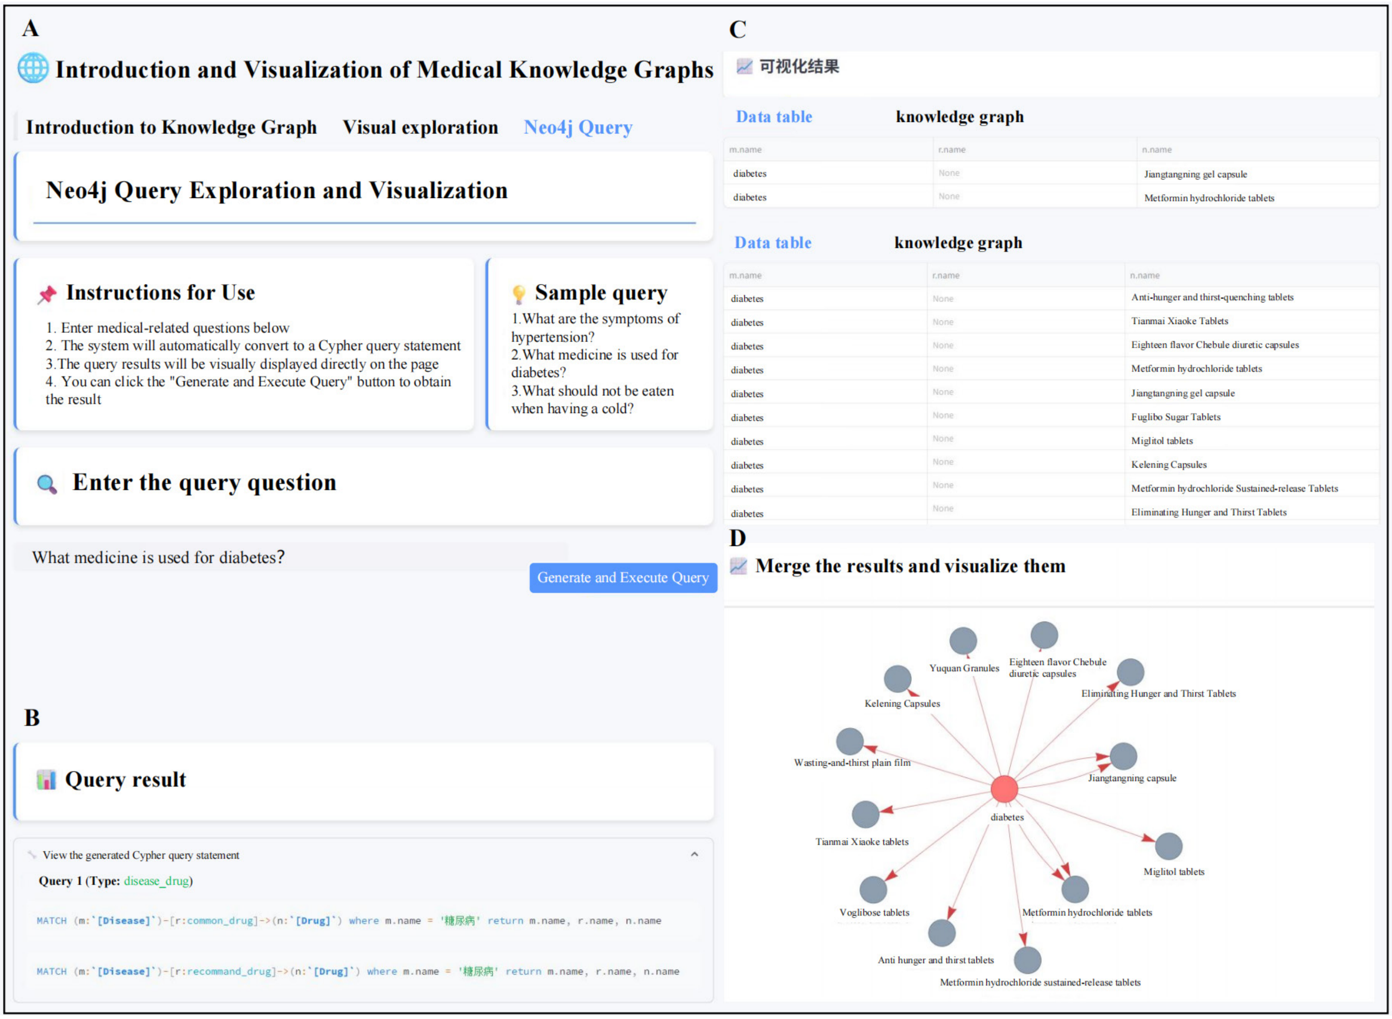

1. Enter medical-related questions below
2. The system will automatically convert to a Cypher query statement
3.The query results will be visually displayed directly on the page
4. You can click the "Generate and Execute Query" button to obtain the result

1.What are the symptoms of hypertension?
2.What medicine is used for diabetes?
3.What should not be eaten when having a cold?

What medicine is used for diabetes?

View the generated Cypher query statement

Query 1 (Type: disease_drug)

```
MATCH (m:`[Disease]`)-[r:common_drug]->(n:`[Drug]`) where m.name = '糖尿病' return m.name, r.name, n.name
```

```
MATCH (m:`[Disease]`)-[r:recommand_drug]->(n:`[Drug]`) where m.name = '糖尿病' return m.name, r.name, n.name
```

m.name	r.name	n.name
diabetes	None	Jiangtangning gel capsule
diabetes	None	Metformin hydrochloride tablets

m.name	r.name	n.name
diabetes	None	Anti-hunger and thirst-quenching tablets
diabetes	None	Tianmai Xiaoke Tablets
diabetes	None	Eighteen flavor Chebule diuretic capsules
diabetes	None	Metformin hydrochloride tablets
diabetes	None	Jiangtangning gel capsule
diabetes	None	Fuglibo Sugar Tablets
diabetes	None	Miglitol tablets
diabetes	None	Kelening Capsules
diabetes	None	Metformin hydrochloride Sustained-release Tablets
diabetes	None	Eliminating Hunger and Thirst Tablets

Figure 7 | Neo4j query visualization.

such as Yuquan granules and Xiaokeling tablets to help users understand medical knowledge more intuitively.

6. CONCLUSION

This study develops a graph-enhanced intelligent QA system that integrates a standardized medical KG and a LLM. Through multi-level graph algorithms (degree centrality, Louvain community detection, K-nearest neighbor similarity, and Dijkstra shortest path), a systematic analysis was conducted on 44,157 entities and 291,170 relationships, identifying highly correlated clinical entities such as acute urethritis, revealing the coupled community structure of diseases and departments, quantifying cross-disease similarity relationships, and mining potential paths for drug repurposing and nutritional intervention. The results show that the KG can effectively reflect the clinical diagnosis and treatment model, and verify the practical value of graph algorithms in medical knowledge mining.

The system innovatively integrates the precise retrieval capabilities of the KG with the semantic generation ability of the Spark Lite model, builds a dual-channel response mechanism, and realizes an interactive visual interface through Streamlit. It shows robust performance in entity retrieval, relationship reasoning, and knowledge interpretation.

The current system still has limitations such as limited knowledge coverage, insufficient processing of complex queries, and the absence of multimodal data. Future work will focus on the following directions: (1) integrating multiple knowledge bases such as UMLS and DrugBank to expand the knowledge coverage; (2) optimizing the algorithm's explainability through weighted relationships and clinical validation; (3) developing a hybrid architecture that combines graph neural networks and LLMs to enhance semantic reasoning capabilities. These improvements aim to enhance the clinical applicability and reliability of the system, providing a technical foundation for the transformation of intelligent medical decision support from research to application.

STATEMENTS AND DECLARATIONS

Ethical Approval

Not applicable.

Consent to Participate

Not applicable.

Consent for Publication

All authors have checked the manuscript and have agreed to the submission.

Conflict of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Authors' Contribution

YH, QL, and MW contributed to data collection. YH also led the study conception, performed data analysis, and drafted the original manuscript. QL assisted with figure preparation and integration and contributed to manuscript drafting and revision. MW and ZX contributed to reference supplementation and review. BD and ZW provided study guidance and suggestions. FY supervised the overall study and reviewed the final manuscript.

Funding

This work was supported in part by National Natural Science Foundation of China (Award No. 62102065), in part by Fujian Medical University Research Foundation of Talented Scholars (Award No. XRCZX2022003), in part by Taizhou Science and Technology Plan Project (Award No. 24ywa61), in part by Fujian Provincial Health Commission (Award No. 2022ZD01003).

Supplementary Materials

Supplementary data related to this article can be found at https://doi.org/10.55578/joaims.260408.001.

REFERENCES

[1] L. Wang, H. Hao, X. Yan, T.H. Zhou, K.H. Ryu, From biomedical knowledge graph construction to semantic querying: a comprehensive approach, Sci. Rep. 15 (2025), 8523.

[2] G. Zhao, W. Gu, W. Cai, Z. Zhao, X. Zhang, J. Liu, MLEE: a method for extracting object-level medical knowledge graph entities from Chinese clinical records, Front. Genet. 13 (2022), 900242.

[3] H.-C. Hsueh, S.-C. Chien, C.-W. Huang, H.-C. Yang, U. Iqbal, L.-F. Lin, et al., A novel multi-level refined (MLR) knowledge graph design and chatbot system for healthcare applications, PLoS ONE. 19 (2024), e0302620.

[4] H. Zhu, A graph neural network-enhanced knowledge graph framework for intelligent analysis of policing cases, Math. Biosci. Eng. 20 (2023), 11585–11604.

[5] X. Mo, J. Pang, B. Wan, R. Tang, H. Liu, S. Jiang, Multi-relational graph contrastive learning with learnable graph augmentation, Neural Netw. 181 (2025), 106757.

[6] R. Johnson, M.M. Li, A. Noori, O. Queen, M. Zitnik, Graph artificial intelligence in medicine, Annu. Rev. Biomed. Data Sci. 7 (2024), 345–368.

[7] X. Zhou, S. Zhang, M. Agarwal, J. Akryod, S. Mosbach, M. Kraft, Marie and BERT—a knowledge graph embedding based question answering system for chemistry, ACS Omega. 8 (2023), 33039–33057.

[8] H. Park, J. Son, J. Min, J. Choi, Selective UMLS knowledge infusion for biomedical question answering, Sci. Rep. 13 (2023), 14214.

[9] C. Fu, X. Pan, J. Wu, J. Cai, Z. Huang, F.v. Harmelen, et al., KG4NH: a comprehensive knowledge graph for question answering in dietary nutrition and human health, IEEE J. Biomed. Health Inform. 29 (2025), 1793–1804.

[10] Y. Feng, L. Zhou, C. Ma, Y. Zheng, R. He, Y. Li, Knowledge graph-based thought: a knowledge graph-enhanced LLM framework for pan-cancer question answering, GigaScience. 14 (2025), giae082.

[11] A. Singhal, Introducing the knowledge graph: things, not strings, Official Google Blog, 2012.

[12] Q. Wang, Z. Mao, B. Wang, L. Guo, Knowledge graph embedding: a survey of approaches and applications, IEEE Trans. Knowl. Data Eng. 29 (2017), 2724–2743.

[13] M. Afshar, Y. Gao, D. Gupta, E. Croxford, D. Demner-Fushman, On the role of the UMLS in supporting diagnosis generation proposed by Large Language Models, J. Biomed. Inform. 157 (2024), 104707.

[14] K. Schatz, D. Korn, A. Tropsha, R. Chirkova, Workflow for domain- and task-sensitive curation of knowledge graphs, with use case of DRKG, 2022 IEEE International Conference on Big Data (Big Data), IEEE, Osaka, Japan, 2022, pp. 3692–3701.

[15] P. Chandak, K. Huang, M. Zitnik, Building a knowledge graph to enable precision medicine, Sci. Data. 10 (2023), 67.

[16] H. Wang, Q. Zu, M. Lu, R. Chen, Z. Yang, Y. Gao, et al., Application of medical knowledge graphs in cardiology and cardiovascular medicine: a brief literature review, Adv. Ther. 39 (2022), 4052–4060.

[17] Y. Lan, S. He, K. Liu, X. Zeng, S. Liu, J. Zhao, Path-based knowledge reasoning with textual semantic information for medical knowledge graph completion, BMC Med. Inform. Decis. Mak. 21 (2021), 335.

[18] D.Y. Li, T.J. Hu, J.L. Li, Q. Qian, W.Y. Zhu, Construction and application of the Chinese unified medical language system, J. Intell. 30 (2011), 147–151.

[19] H.Y. Zan, H.Y. Dou, Y.X. Jia, T.F. Guan, B. Odbayar, K.L. Zhang, et al., Construction of Chinese medical knowledge graph based on multi-source texts, J. Zhengzhou Univ. (Nat. Sci. Ed.), 52 (2020), 45–51.

[20] H. Long, Y. Zhu, L. Jia, B. Gao, J. Liu, L. Liu, et al., An ontological framework for the formalization, organization and usage of TCM-knowledge, BMC Med. Inform. Decis. Mak. 19 (2019), 53.

[21] X. Xue, Y. Mei, B. Zhao, M. Zhang, Adaptive similarity feature construction for ontology matching via multilayer hybrid genetic programming, IEEE Trans. Evol. Comput. 30 (2026), 519–533.

[22] T. Cheng, T. Bi, W. Ji, C. Tian, Graph convolutional network for image restoration: a survey, Mathematics. 12 (2024), 2020.

[23] R. Zeng, C. Zeng, X. Wang, B. Li, X. Chu, Incentive mechanisms in federated learning and a game-theoretical approach, IEEE Netw. 36 (2022), 229–235.

[24] L. Ma, M. Huang, S. Yang, R. Wang, X. Wang, An adaptive localized decision variable analysis approach to large-scale multiobjective and many-objective optimization, IEEE Trans. Cybern. 52 (2022), 6684–6696.

[25] T. Winograd, Five lectures on artificial intelligence, Stanford University, Stanford, CA, USA, 1974.

[26] J. Kupiec, MURAX: a robust linguistic approach for question answering using an on-line encyclopedia, Proceedings of the 16th Annual International ACM SIGIR Conference on

Research and Development in Information Retrieval, ACM, New York, NY, United States, 1993, pp. 181–190.

[27] L. Dong, M. Lapata, Language to logical form with neural attention, Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL, Berlin, Germany, 2016, pp. 33–43.

[28] R. Jia, P. Liang, Data recombination for neural semantic parsing, Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL, Berlin, Germany, 2016, pp. 12–22.

[29] R. Wang, M. Wang, J. Liu, W. Chen, M. Cochez, S. Decker, Leveraging knowledge graph embeddings for natural language question answering, In: G. Li, J. Yang, J. Gama, J. Natwichai, Y. Tong (Eds.), Database Systems for Advanced Applications, DASFAA 2019, Lecture Notes in Computer Science, Vol. 11446, Springer, Cham, 2019, pp. 659–675.

[30] M.Y. Cao, Q.Q. Li, Z.H. Yang, L. Wang, Y. Zhang, H.F. Lin, et al., A question answering system for primary liver cancer based on knowledge graph, J. Chin. Inf. Process. 33 (2019), 88–93.

[31] Y.Y. Chen, Construction and research of knowledge graph for stroke, syndrome, disease and medicine based on deep learning, Master's Thesis, Gansu University of Chinese Medicine, Lanzhou, China, 2023.

[32] Z.M. Yi, K. Li, Z.Y. Xu, Z.M. Xu, F. Gao, A construction method and device for personalized medication assistance question answering system for Parkinson's disease, Patent CN202410698116.0, Beijing, China, 2024.

[33] L.C. Freeman, Centrality in social networks conceptual clarification, Soc. Netw. 1 (1978–1979), 215–239.

[34] V. Blondel, J.-L. Guillaume, R. Lambiotte, Fast unfolding of communities in large networks: 15 years later, J. Stat. Mech. Theory Exp. 2024 (2024), 10R001.

[35] S. Lloyd, Least squares quantization in PCM, IEEE Trans. Inf. Theory. 28 (1982), 129–137.

[36] E.W. Dijkstra, A note on two problems in connexion with graphs, Numer. Math. 1 (1959), 269–271.

[37] H. Lu, M. Halappanavar, A. Kalyanaraman, Parallel heuristics for scalable community detection, Parallel Comput. 47 (2015), 19–37.

[38] D. Liben-Nowell, J. Kleinberg, The link-prediction problem for social networks, J. Am. Soc. Inf. Sci. Technol. 58 (2007), 1019–1031.

[39] B. Perozzi, R. Al-Rfou, S. Skiena, DeepWalk: online learning of social representations, Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, NY, United States, 2014, pp. 701–710.