Building and Applying a Knowledge Graph for Interpretable Diagnosis and Treatment of TCM Arthralgia Syndrome

Kai Chang[1,2], Liangzhi Yuan[2], Jiali Hai[2], Wenping Deng[2†], Yong Xiao[2†]

1. Institute of Medical Intelligence, School of Computer Science and Technology, Beijing Jiaotong University, Beijing, China

2. College of Information Engineering, Hubei University of Chinese Medicine, Wuhan, Hubei, China

Abstract

The complexity, semantic ambiguity, and lack of standardized expression and interpretability in Traditional Chinese Medicine (TCM) syndrome differentiation and treatment severely limit its effectiveness in intelligent diagnostic support. To improve the accuracy and interpretability of intelligent assistance in the diagnosis and treatment of arthralgia syndrome, this study proposes a knowledge graph-integrated approach to syndrome differentiation and treatment. Firstly, based on real-world clinical records, we systematically extract and normalize knowledge related to arthralgia syndrome to construct the Traditional Chinese Medicine Arthralgia Syndrome Knowledge Graph (TCMASKG). This knowledge graph encompasses key elements such as etiology and pathogenesis, clinical symptoms, TCM syndrome, prescription-syndrome correspondence, and therapeutic principles, ultimately forming 4,515 entity nodes and 90,776 semantic relations, thereby structurally representing the interrelationships among knowledge in arthralgia diagnosis and treatment. Secondly, by integrating real diagnostic data with the knowledge graph, we design and train a graph-enhanced model for intelligent syndrome differentiation and treatment recommendation. This model not only significantly improves the accuracy of syndrome classification and herbal prescription recommendation but also provides diagnostic reasoning paths and evidence chains through knowledge graph visualization, thereby realizing transparency and interpretability in the diagnostic process. Experimental results demonstrate that the proposed method achieves superior performance in intelligent assistance for syndrome differentiation and treatment of arthralgia, highlighting the broad application prospects of knowledge-driven approaches in the modernization and intelligent development of TCM.

Keywords: Knowledge Graph Construction; TCM Arthralgia Syndrome; Interpretable Syndrome Differentiation; Machine Learning

1. Introduction

The diagnosis in traditional Chinese medicine (TCM) is guided by theories of Zang-Fu organs, meridians, Qi, blood, body fluids, etc., focusing on how to inspect and identify diseases

[†]Corresponding author: Wenping Deng (Email: chw_dwp@sina.com)
[†]Corresponding author: Yong Xiao (Email:xy_1015@126.com)

and syndromes. It serves as a bridge between the basic disciplines of TCM and various clinical specialties[1]. The essence of TCM theory and clinical practice lies in syndrome differentiation and treatment, which is also a primary method for diagnosis and therapy. Syndrome differentiation involves collecting patient data, symptoms, and signs through observation, listening, questioning, and pulse-taking, then analyzing and synthesizing to clarify the disease's cause, nature, location, and the relationship between pathogenic factors and body resistance to determine the disease. Treatment, based on the results of the syndrome differentiation, involves identifying the disease name according to its specific manifestations and prescribing medication according to the syndrome, adopting a method that treats both the root cause and symptoms[2]. The most crucial aspect of treating diseases is "to first examine its source, to first wait for the disease mechanism." The disease mechanism is the key core of disease occurrence and the fundamental principle of treatment. It reflects the patient's cause of disease, condition, nature, and even the trend of pathological changes, making "treatment according to the disease mechanism" a core aspect of the differentiation and treatment of the syndrome[3].

TCM arthralgia syndrome encompasses a range of musculoskeletal and joint disorders. These conditions are attributed to the disruption of Qi and blood flow due to external pathogenic factors, manifesting in various symptoms. Contemporary medical diagnoses such as rheumatoid arthritis, gout, shoulder periarthritis, tenosynovitis, fibromyalgia, and bone spurs fall within this category. The approach to treating these disorders with TCM has garnered growing interest for its research and clinical applications[4].

Since the introduction of the Knowledge Graph (KG) by Google in 2012, this technology has seen rapid development and widespread application across various domains, including healthcare [5]. Among the most well-known applications in medicine is the clinical decision support system. For instance, Robbins et al. have demonstrated the effectiveness of clinical decision systems based on Knowledge Graphs in HIV research[6]. Hasan et al. developed a prototype Knowledge Graph based on Louisiana's cancer registry dataset, enriching the data semantically and facilitating easy linkage with third-party data, helping to explain variations in cancer incidence, disparities and outcomes[7]. With the continuous deepening of information technology, more Knowledge Graphs support clinicians in developing personalized treatment plans, combining TCM's unique theories and personalized treatment system of syndrome differentiation and treatment, forming a unique prevention and treatment model for chronic and difficult diseases.

This study focuses on TCM arthralgia syndrome, collecting a large amount of data from famous TCM physicians' medical records and encyclopedic knowledge data on TCM arthralgia syndrome, cleaning, and structuring the data, and standardizing the content on diseases, symptoms, and syndrome analysis. Combining basic TCM theories, a semantic network between TCM arthralgia syndrome ontologies is constructed, integrating structured data with the ontology structure to form a Knowledge Graph. Based on the triple data of the Knowledge Graph, a quality assessment is conducted, assigning a reliability evaluation dimension to each triplet, thus forming a high-quality Knowledge Graph. This work establishes a generalizable approach for constructing interpretable syndrome differentiation and treatment models based on knowledge graphs, offering new paradigms for the digitization, explanation, and reuse of TCM knowledge, as well as references for studying other diseases in the TCM domain.

The main contributions of this study are threefold:

(1) Construction of the TCMASKG: Leveraging representative clinical records from renowned TCM practitioners, we constructed a highquality knowledge graph rooted in realworld data. A novel evaluation mechanism was introduced to assess the specificity and co-occurrence probabilities of relations, effectively minimizing potential inaccuracies and enhancing semantic con-

sistency.

(2) Development of an Intelligent Syndrome Differentiation and Treatment Model: Based on TCMASKG, we developed a graph-enhanced intelligent model that simulates the syndrome differentiation process of experienced TCM physicians. The model comprehensively analyzes multidimensional patient data to accurately identify subtypes of arthralgia syndrome and recommend appropriate treatment strategies.

(3) Enhancement of Interpretability in Clinical Reasoning: In contrast to traditional black-box models, our proposed framework explicitly links disease etiology and pathomechanisms to specific syndromes. This enables the model to infer underlying causes and internal pathological changes and generate individualized treatment strategies. As a result, the approach not only improves the accuracy and clinical utility of syndrome differentiation but also significantly enhances the transparency and interpretability of diagnostic and therapeutic decisions.

2. Related Work

The explosive growth of online content has led to increasing data diversity, heterogeneity, and unstructured formats, making it progressively more difficult to retrieve effective and relevant information. To address this challenge, the concept of the Knowledge Graph (KG) was proposed. Leveraging its strong capabilities in structured knowledge representation and semantic processing, the Knowledge Graph provides an intuitive and precise way to describe entities or concepts and their interrelations. Prominent large-scale general-purpose KGs such as Freebase [8], YAGO [9], and NELL [10] have demonstrated the effectiveness of this approach in various domains.

In recent years, Knowledge Graphs have been widely applied in medical systems, enhancing clinical research and decision-making support for physicians, thereby continuously improving the quality of medical services [11, 12, 13]. Cheng et al. [5] provided a comprehensive review of recommendation systems based on Knowledge Graphs, systematically analyzing the foundational concepts and key reasoning mechanisms, including rule-based, embedding-based, and neural network-based methods. Jiang et al. [14] developed a medical recommendation system that integrates medical KGs with deep learning models, enabling personalized diagnosis and treatment suggestions for patients. However, most of the current medical KGs are constructed from standardized medical literature and ontologies, which may not fully reflect the variability, uncertainty, and complexity of real-world clinical scenarios [15]. This disconnect highlights the pressing need to incorporate large-scale real clinical data to improve the accuracy, adaptability, and practicality of medical Knowledge Graphs.

In the field of TCM, the inherently complex and context-dependent nature of knowledge presents further challenges for effective knowledge organization and application. To address this, researchers have proposed various methods for constructing and applying TCM-specific Knowledge Graphs. For example, Yang et al. [16]proposed an integrated Diagnostic Decision Support System (DSDS) that employs KG representation learning to recommend syndromes based on TCM principles. Similarly, Gao et al. [17] developed a TCM question-answering system using a Naive Bayes-based approach, extracting keywords from user queries and generating graph-based queries accordingly. These efforts have advanced the informatization of TCM and provided technical foundations for intelligent services.

Nevertheless, existing research often lacks interpretability from the perspective of syndrome differentiation and treatment, which is the central diagnostic and therapeutic paradigm in TCM [18, 19]. In clinical decision-making, interpretability is of paramount importance. Many current

intelligent systems, including prescription recommendation and syndrome classification models, function as black boxes, making it difficult for users—particularly physicians—to understand the rationale behind a given recommendation [20]. As a result, even if such systems demonstrate high accuracy, they may not gain the trust of clinicians due to their inability to provide transparent reasoning pathways [21, 22].

Although considerable progress has been made in constructing TCM Knowledge Graphs and promoting their application in digital health platforms, important challenges remain in the selection of foundational data, as well as in effectively addressing the interpretability and complexity of syndrome differentiation and treatment in TCM [23, 24]. The innovation of this study lies in the systematic use of real-world clinical case data from multiple renowned TCM physicians as the foundational resource for constructing the Knowledge Graph. This approach captures not only diverse theoretical perspectives but also reflects individual physicians' practical experiences in treating arthralgia syndrome. It ensures that the resulting graph embodies both theoretical richness and clinical practicality, making it more representative, credible, and generalizable.

By constructing a high-quality Knowledge Graph grounded in authentic medical cases and integrating it with advanced machine learning algorithms, this study aims to address the multifaceted challenges of individual patient variability and complex disease mechanisms in TCM syndrome differentiation and treatment, providing robust support for the intelligent, personalized, and interpretable implementation of TCM diagnosis and treatment.

3. Methods

3.1. Framework Overview

The construction of a knowledge graph involves extracting structured triplets from large volumes of structured and unstructured data, and storing them in a semantically meaningful manner. This process is organized into two layers: a logical layer, which defines the graph schema and structure, and a data layer, which stores the extracted knowledge [25]. Two common approaches exist for knowledge graph construction: top-down, where the schema is defined before data is populated, and bottom-up, where data accumulation drives schema development. This study adopts the top-down approach for constructing the knowledge graph [20].

We utilized high-quality standardized clinical case archives from senior TCM experts, combined with information extracted from authoritative online encyclopedias, and employed deep learning techniques to efficiently structure knowledge related to TCM arthralgia syndromes. The construction begins with the design of a standardized ontology that defines relevant concepts and their relationships. This step ensures terminological clarity and semantic coherence across the knowledge graph, preventing conceptual duplication and enhancing its usability in reasoning and practical applications.

In the data refinement phase, we adopt a hybrid method combining manual annotation and deep learning to ensure the integrity and completeness of the extracted knowledge. A relationship-based data evaluation mechanism is introduced to facilitate seamless integration of entities, attributes, and relationships. This not only improves the interpretability and scalability of the knowledge graph but also deepens the understanding and practical applicability of TCM knowledge, contributing to its digital transformation. The overall construction framework of TCMASKG is shown in Figure 1.

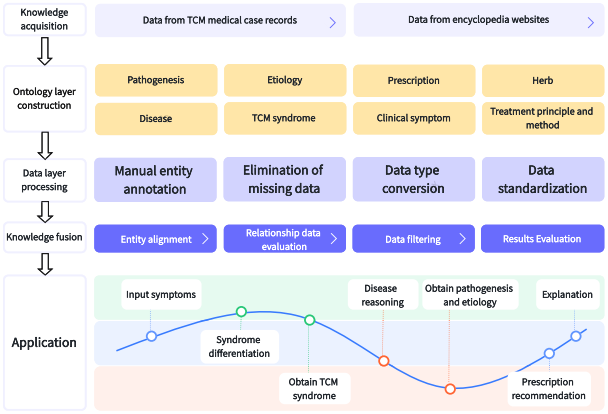


Figure 1: Construction framework of TCMASKG.

3.2. Ontology Layer Construction

The ontology layer, or schema, forms the backbone of the knowledge graph. It provides a standardized semantic framework that supports data integration, representation, and reasoning. To guide the schema design, we adopt the seven-step ontology development method proposed by Stanford University and align it with the international standard ISO 19465-2017, which classifies TCM clinical terminology. The resulting top-level ontology categorizes TCM knowledge into six major classes, each with well-defined attributes: disease, TCM syndrome, mechanism of disease, clinical symptom, treatment principle and method, and medication.

The ”Disease” class is subdivided into two categories: TCM disease and Western medicine disease, enabling integration and comparison between Eastern and Western medical knowledge systems. The mechanism of disease is modeled along two dimensions: etiology (causal factors) and pathogenesis (development and progression), capturing both the origin and dynamic evolution of disease. Regarding treatment, medications are categorized into prescription and individual herb, reflecting the diverse and distinctive therapeutic approaches of TCM. The schema design process is illustrated in Figure 2.

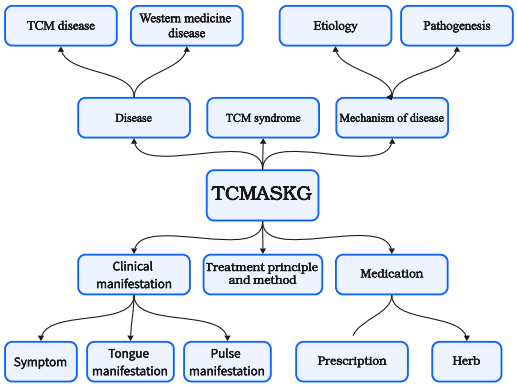


Figure 2: The overall schema design process of the TCMASKG.

To define semantic connections between entities, we established object property pairs along with their inverse relationships, as shown in Table 1. These relationships serve as the foundation for knowledge integration and logical reasoning processes.

Table 1: Semantic relationships of the TCMASKG.

Object Property	**Inverse Property**	**Domain**	**Range**
reflects	reflected_in	Western medicine disease	TCM Syndrome
has_symptom	symptom_of	Western medicine disease	Symptom
has_pulse	pulse_of	Western medicine disease	Pulse manifestation
has_tongue	tongue_of	Western medicine disease	Tongue manifestation
caused_by	causes	Etiology	Western medicine disease
manifests_as	manifested_through	Western medicine disease	Pathogenesis
treated_by	treats	Western medicine disease	Treatment principle and method
associated_with_western	associated_with_TCM	TCM Disease	Western medicine disease
has_symptom	symptom_of	TCM Syndrome	Symptom
has_pulse	pulse_of	TCM Syndrome	Pulse manifestation
has_tongue	tongue_of	TCM Syndrome	Tongue manifestation
treated_with	prescription_for	TCM Syndrome	Prescription
composed_of	part_of	Prescription	Herb

3.3. Data Layer Processing

3.3.1. Data Acquisition

The dataset for the TCMASKG was primarily sourced from the medical records of renowned practitioners in the field of TCM and online encyclopedias, with a particular focus on TCM arthralgia syndrome. A comprehensive collection of 2,165 medical records was acquired from the National Service Platform for Renowned TCM Practitioners [26]. These records included 39 structured fields such as demographics, symptoms, pulse and tongue diagnoses, medical history, syndrome differentiation, TCM and Western diagnoses, treatment principles, and prescriptions.

Records with missing syndrome labels or significant gaps in content were excluded, resulting in a refined dataset of 1,843 valid entries. Syndrome labels were further processed using punctuation-based segmentation and regular expressions to eliminate noisy or redundant terms.

To supplement this clinical data, Python-based web scraping techniques were employed to collect additional content from high-quality online resources such as Xunjiwenyao (https://www.xywy.com), Baidu Encyclopedia (https://baike.baidu.com/), and 39 Medicine (https://www.39.net).

This extended dataset brought the total number of valid records to 2,144, providing a comprehensive knowledge base for TCM arthralgia syndrome. Since this study exclusively utilized anonymized, publicly available data without any personal identifiers, ethical approval was not required.

3.3.2. Knowledge extraction

The dataset incorporates both structured and unstructured information. While structured fields such as diagnosis and medication are relatively straightforward to process, the expression of syndrome differentiation in TCM is often intricate and nuanced. To address this, we adopted a hybrid approach that combines deep learning techniques with rule-based methods such as regular expressions, enabling more effective parsing of both structured and unstructured data elements.

This dual-method strategy leverages the semantic understanding and contextual learning capabilities of deep learning with the efficiency and precision of rule-based pattern matching. Deep learning, in particular, demonstrates significant advantages in processing unstructured narrative data commonly found in medical records.

Initial data annotation was conducted using the collaborative electronic medical record annotation platform (www.tcmai.org) [27]. Four graduate students majoring in TCM were involved in the first-round annotation, following an ontology-driven schema tailored for arthralgia syndrome. The annotated data were subsequently reviewed and refined by three experienced TCM clinicians, ensuring domain-specific consistency and accuracy.

We utilized the BERT_BiLSTM_CRF model for named entity recognition (NER) to identify relevant TCM entities in the dataset. The model architecture comprises three core components: an embedding layer, a BiLSTM layer, and a CRF layer. This architecture is depicted in Figure 3.

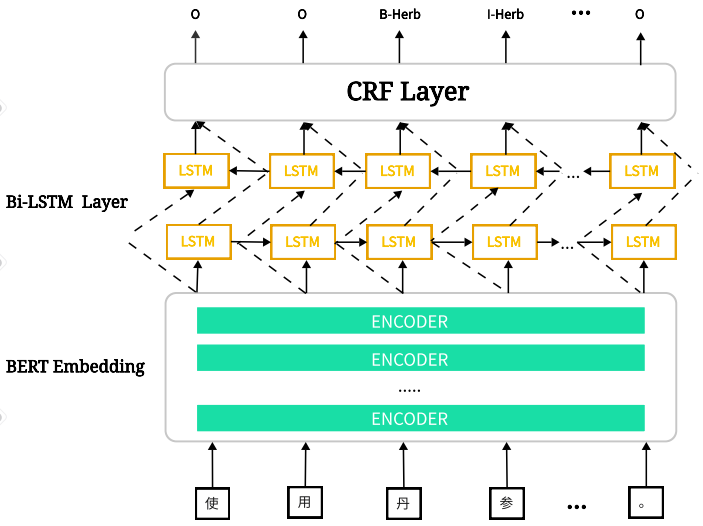


Figure 3: Architecture of the Bert_BiLSTM_CRF framework.

Embedding Layer: This layer maps discrete tokens (words or characters) into dense vector representations. In our study, we adopted the BERT pre-trained language model as the embedding layer, which leverages bidirectional attention to effectively encode context information.

This design is particularly suited to the TCM domain, where annotated NER datasets are scarce, and text semantics are highly context-dependent.

BiLSTM Layer: The Bidirectional Long Short-Term Memory network is employed to capture long-range dependencies in the sequence. Unlike traditional unidirectional LSTMs, BiLSTM processes input from both past and future contexts, enhancing the model's ability to understand word dependencies in complex medical narratives.

CRF Layer: The Conditional Random Field layer adjusts the output label sequence by maximizing the probability of the optimal global label sequence. This helps capture dependencies between adjacent labels and corrects inconsistencies in entity boundaries, further improving sequence labeling accuracy.

Model training was implemented using PyTorch 1.5.1 with acceleration provided by an NVIDIA A10 GPU. The corpus was divided into training and testing sets in a 7:3 ratio. Hyperparameters were selected based on preliminary experiments: a batch size of 32 was chosen to balance computational efficiency and model stability, while a learning rate of 1e-5 was adopted to facilitate gradual and stable convergence. To alleviate potential overfitting, the dropout rate was set to 0.5 during training, and early stopping was employed based on validation loss to prevent performance degradation.

Model performance was evaluated using precision (P), recall (R), and F1-score metrics. Precision represents the proportion of correctly predicted entities among all predicted entities, recall measures the proportion of true entities successfully identified, and F1-score balances these two metrics using the harmonic mean: F1 = 2PR/(P+R). Evaluation results for the three entity types are presented in Table 2, demonstrating satisfactory accuracy in recognizing TCM entities relevant to arthralgia syndrome. These parameter settings were informed by iterative tuning and empirical validation on the development set. In future work, we plan to incorporate more advanced methods to further evaluate model generalizability and robustness across different data partitions.

Table 2: Evaluation of entity extraction results.

Entity type	**Precision**	**Recall**	**F1-score**
Pathogenesis	0.87	0.83	0.86
Etiology	0.89	0.86	0.88
Symptom	0.92	0.88	0.90

3.3.3. Data Fusion

Due to the complexity and heterogeneity of knowledge extraction sources, inconsistencies in hierarchical structure, logical connections, and granularity of concepts often occur. Additionally, redundant knowledge and unbalanced expressions from different sources are common issues [28]. Variations in terminology across databases can result in semantic ambiguity—where identical named entities may refer to different concepts, or different named entities may refer to the same concept—thereby compromising the consistency and utility of the constructed knowledge graph.

To address these challenges, it is essential to perform entity alignment, which involves mapping source terms to standardized terms and constructing unified entities accordingly. In this study, we employed the third-party ontology system "BIOS" [29] for entity alignment. "BIOS"

is a comprehensive medical knowledge ontology constructed using cutting-edge technologies such as deep learning and text mining, covering over 54 million medical terms and 26 million medical concepts extracted from large-scale, authoritative medical texts.

We adopted a hybrid approach that combines string matching and semantic similarity calculation, incorporating both entity attributes and contextual language representations to achieve efficient and accurate alignment. For example, the symptom "vomiting pattern" may have synonymous expressions such as "omiting patterns" or "emesis pattern" Using entity alignment, these variations can be mapped to a unified standard term ID (e.g., CN10220518), thus achieving normalization across knowledge sources. This strategy enables automated and consistent integration of heterogeneous knowledge into the graph, significantly enhancing its accuracy and reliability. Experimental results show that the "BIOS" interface provides high precision in entity alignment and can effectively support the construction of a structured and standardized Traditional Chinese Medicine (TCM) knowledge graph.

3.3.4. Relationship Data Evaluation

Traditional knowledge graphs typically represent relationships using triplets in the form of subject-predicate-object. While this structure captures basic relational semantics, it becomes inadequate for applications involving clinical decision support systems, which demand finer-grained, context-rich knowledge representations and computational reasoning. To address this limitation, we adopted the quadruple structure proposed in [30], which introduces a fourth dimension to represent attributes of each relational triplet. Specifically, we computed two attribute values: co-occurrence probability and specificity.

Co-occurrence probability reflects how frequently the object entity appears in association with a specific subject entity under a given relationship. It is defined as:

$$P(O_{ij} \mid S_i) = \frac{K_{co}(S_i, O_{ij})}{K(S_i)} \tag{1}$$

Where $K(S_i)$ is the total count of occurrences of subject entity S_i, and $K_{co}(S_i, O_{ij})$ is the count of co-occurrences of S_i and the object entity O_{ij} under relation P.

Specificity quantifies how indicative the object entity is for the given subject compared to all subjects:

$$specificity(S_i, O_{ij}) = \frac{P(O_{ij} \mid S_i)}{P(O_{ij} \mid S)}. \tag{2}$$

where

$$P(O_{ij} \mid S) = \frac{\sum_{l=1}^{|S|} K_{co}(S_i, O_{ij})}{\sum_{l=1}^{|S|} K(S_i)} \tag{3}$$

Where i= 1, 2, ..., S_i, j = 1, 2, ..., J_i, and J_i denotes the number of object entities associated with subject S_i under relationship P. Table 3 provides examples of symptoms associated with the syndrome Heat-Toxin Congestion, annotated with co-occurrence probability and specificity.

Compared with simple co-occurrence, specificity emphasizes the discriminative value of an object entity for a particular subject. In clinical TCM diagnosis, some symptoms, although generally infrequent, may have high indicative power for specific syndromes. Capturing this with specificity improves both the expressiveness and inference capabilities of the knowledge graph.

Table 3: Examples of symptoms associated with Heat-Toxin congestion.

TCM Syndrome	Symptom	Probability	Specificity
Heat-Toxin Congestion	Sore throat	0.035	10.81
	Diarrhea	0.035	36.51
	Stiffness and pain in the lower back	0.035	97.00
	Cough	0.14	1.52

3.3.5. Data Filtering

To ensure data quality, we implemented a filtering process using the calculated co-occurrence probability and specificity from the quadruples. The process includes:

Triplet-level filtering: If both the co-occurrence probability $P(O_{ij} \mid S_i)$ and specificity $specificity(S_i, O_{ij})$ fall below defined thresholds, the quadruple is discarded.

Entity-level filtering: Entities with frequencies below a predefined threshold within their category are removed.

Generally, the minimum threshold should be set to a value that ensures the quality of the data while not being too strict, leading to too many valid data being filtered out. Through multiple experiments and comparisons, we found that the appropriate minimum thresholds for co-occurrence probability and entity occurrence are 0.001 and 5, respectively, and the minimum threshold for specificity should be set to 1.2. This value also needs to be adjusted according to the characteristics of the specific data and the actual application scenarios.

Using these thresholds, over 20,000 invalid relational records were removed, while more than 40,000 high-quality quadruples were retained, significantly improving the overall integrity and utility of the graph. To validate filtering effectiveness, we invited TCM domain experts to compare node relationships before and after cleaning. Table 4 presents an example of symptom changes for the syndromes Wind-Cold-Dampness Obstruction and Cold Deposition.

Table 4: Data comparison before and after cleaning.

TCM Syndrome	Before Cleaning	After Cleaning
Wind-cold-dampness obstructing	Red tongue	Thin white coating on tongue
	Thin yellow coating on tongue	Redness and swelling
	Joint pain	Back pain
	Bitter taste in mouth	Thirsty
	Abdominal distension	Restless
	Spasms	Bitter taste in mouth
	Distending pain	Joint pain
Cold deposition	Cyanosis	Dark purplish lips and tongue
	White coating on tongue	White coating on tongue
	Deep and thin pulse	Deep and thin pulse
	Dark purplish lips and tongue	Cyanosis
	Pale complexion	Pale complexion
	Numbness	Dizziness and blurred vision

Specifically, the experts compared the basic statistical features of the data, and compared the consistency and accuracy of each node and its connected nodes. The results showed that after data cleaning, the quality of the data was significantly improved. The statistical features of the cleaned data were more reasonable, and the consistency and accuracy were also significantly improved. This confirms that quadruple-based processing, combined with co-occurrence probability thresholds, effectively refines the reliablity of TCMASKG. Moreover, a sensitivity analysis revealed that moderate variation in co-occurrence thresholds did not significantly impact the overall results, suggesting that the proposed filtering scheme is robust across practical scenarios.

3.4. Data Storage

The cleaned and structured data were stored using the Neo4j graph database. Storage mainly used the Python driver library in conjunction with Neo4j's Cypher language to carry out operations such as creation, querying, updating, and deletion of data. Nodes were used as entities, labels were used as the entity's belonging category, edges were used to represent relationships between entities, and the specific attributes of entities were stored as node properties. The relationships between entities and entities were constructed in the form of triples. This storage approach enables efficient graph traversal, flexible querying, and scalable management of TCM clinical knowledge, supporting subsequent tasks such as knowledge reasoning and intelligent clinical decision-making.

4. Results

4.1. Knowledge Graph Construction and Analysis

We constructed TCMASKG, a comprehensive knowledge graph that integrates diverse entities and their rich semantic relationships. The graph contains a total of 4,515 entity records, including entities such as diseases, symptoms, syndromes, and diagnoses. As shown in Figure 4, the graph includes 887 nodes representing Herbs, which are essential components in Traditional Chinese Medicine (TCM) treatment formulations. There are 484 nodes related to Treatment Principles and Methods, reflecting the diverse therapeutic strategies used in TCM diagnosis and therapy. The graph includes 482 nodes for Diseases, covering both TCM-specific disease names and their Western medicine counterparts. A total of 664 nodes are dedicated to Prescriptions, representing various herbal formulae tailored to specific clinical needs.

Through structural analysis of the graph, we identified 90,776 relationship entries, as illustrated in Figure 5. Among them, there are 14,006 associations linking Western medicine disease and Symptom, highlighting the symptomatology profile for each condition. A total of 3,552 connections exist between Treatment Principles and Methods and Western medicine disease nodes, offering insight into therapeutic approaches. There are 25,968 links connecting Herbs to Prescriptions. Furthermore, 2,018 links between Prescriptions and TCM Syndromes map prescriptions to the syndromic patterns they address.

Additionally, there are 3,570 direct relationships between Western medicine disease and TCM Syndrome nodes, embodying the fundamental TCM concept of treating diseases based on syndrome differentiation. Etiology nodes are linked to Western medicine disease nodes via 6,550 causal pathways, which thoroughly depict the internal and external pathogenesis mechanisms. Moreover, TCM Syndrome nodes connect to Symptom nodes through 20,510 relationships, forming a robust basis for pattern recognition and syndrome identification in practice.

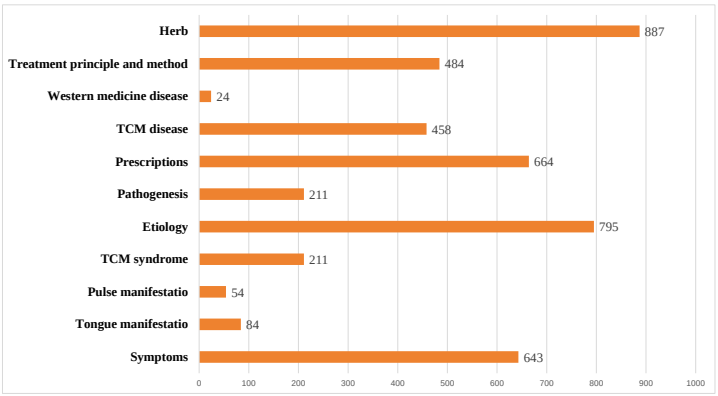


Figure 4: Statistics on the number of entities within the TCMASKG.

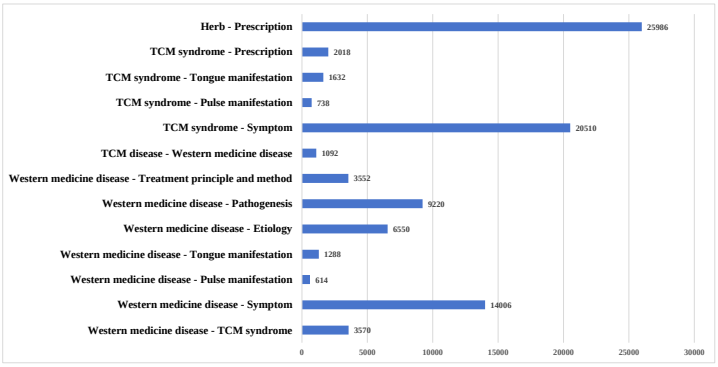


Figure 5: Statistics on the number of semantic relationships within the TCMASKG.

These interconnections form a large-scale and intricate knowledge network. In clinical decision support systems and medical research, such relational data can assist physicians in making diagnostic and therapeutic decisions, improving the quality and accuracy of patient care. A partial example of the graph stored in the Neo4j database is shown in Figure **??**. The figure presents the outcome of a Cypher Query Language (CQL) query related to the disease Nodular Erythema. As a disease node, Nodular Erythema is linked to a variety of associated nodes, such as those representing symptoms, treatments, etiological factors, and complications. This networked representation offers an intuitive visualization of the complex relationships surrounding this condition and facilitates further investigation. Leveraging the powerful querying capabilities of the Neo4j database, researchers and clinicians can quickly access relevant information and conduct deeper knowledge exploration and analytical tasks..

4.2. Syndrome Differentiation Model

In TCM practice, diagnosis and treatment are centered around the concept of syndrome differentiation. However, syndrome identification is highly subjective, relying on the physician's interpretation of patients' clinical manifestations. This leads to personalized and diverse syndrome distributions, which are closely related to the physician's knowledge base and clinical

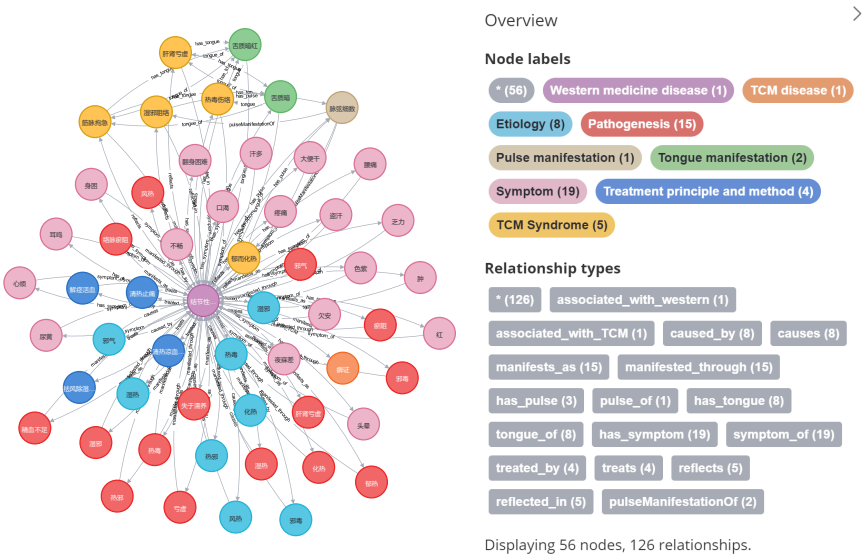


Figure 6: Partial visual presentation of the TCMASKG.

experience. Moreover, inter-individual variations among patients and differences in diagnostic styles among doctors present considerable challenges for accurate differential diagnosis, which in turn hampers the overall advancement of TCM syndrome-based diagnosis and treatment.

Within a TCM knowledge graph, a single syndrome can be associated with multiple symptoms and often corresponds to several prescriptions. Therefore, conducting accurate differential analysis based on this knowledge graph is a complex task. To address this, we utilized key clinical features—such as positive symptoms, tongue appearance, and pulse conditions—to identify and recommend all plausible syndrome patterns. Naive Bayes is commonly used in noise filtering tasks; here, we apply the same principle to filter and rank likely syndromes. The formula used to calculate the likelihood of a syndrome given a set of symptoms is:

$$P(D_i \mid S_1, S_2, ...S_m) = \frac{1}{J} \times P(D_i) \times \prod_{t=1}^{m} P(S_t \mid D_i). \tag{4}$$

where $P(D_i \mid S_1, S_2, ..., S_m)$ represents the posterior probability of syndrome D_i given a set of clinical manifestation. $P(D_i)$ denotes the prior probability of the syndrome D_i, obtained from the frequency of corresponding quadruples in the knowledge graph. If a particular quadruple does not exist, its probability is considered to be zero. The conditional probability $P(S_t \mid D_i)$ indicates the likelihood of observing symptom S_t under syndrome D_i. $J = P(D_i, S_1, ..., S_m)$ is the joint probability of syndrome D_i and all observed symptoms. However, since J is constant across all syndrome candidates for a given set of symptoms, it does not affect the relative ranking of probabilities and can be omitted during comparison.

Based on this formulation, for each test case, we first extract the relevant symptom features and compute the posterior probabilities of all possible syndromes. To avoid zero-probability issues when a symptom is not explicitly linked to a syndrome, we apply Laplace smoothing when calculating the conditional probabilities. The final predicted syndrome is the one with the

highest posterior probability. This probabilistic reasoning approach allows us to filter out unreasonable candidates and recommend the most plausible syndrome patterns, thereby improving the accuracy and interpretability of syndrome differentiation.

For instance, if a patient exhibits symptoms such as joint pain, muscle soreness, and numbness, these may correspond to multiple syndromes, such as Wind-Cold-Damp Arthralgia or Wind-Heat-Damp Arthralgia. Using our syndrome differentiation system, we can recommend potential syndromes and their associated pathogenesis based on the symptom profile. Physicians can then select the most appropriate treatment plan informed by the likelihood scores. This personalized approach significantly enhances the precision of syndrome differentiation and improves clinical decision-making. An example illustrating the operation of the syndrome differentiation model is shown in Figure 7.

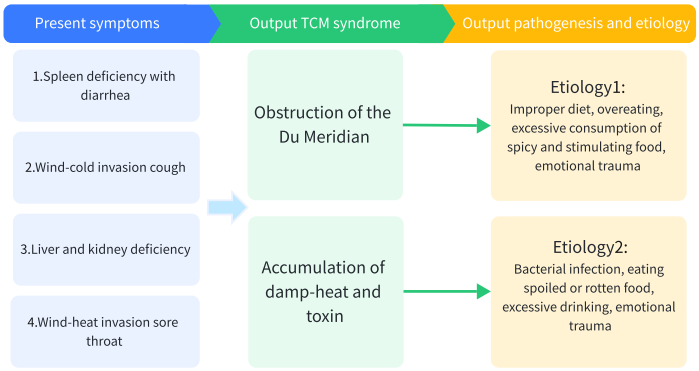


Figure 7: Examples of the process of syndrome differentiation.

To validate the effectiveness of the syndrome differentiation module, this study employed 958 medical case records that were not involved in the construction of the knowledge graph. The positive symptoms documented by physicians in these cases were used as input features for validation, while the corresponding TCM syndromes served as ground truth labels. Considering that the prediction model may yield multiple possible syndromes for a single case, we adopted the Top-K evaluation metric to comprehensively assess the performance of multi-label predictions.

In the evaluation formulas, i represents the i-th recommendation, $R(i)$ represents the set of correct corresponding syndromes in the i-th recommendation, $T(i)$ represents the set of true recommended syndromes in the i-th recommendation, N represents the total number of recommendations, and K indicates the number of top-ranked predictions considered. The evaluation metrics are defined as follows:

Precision@K:

$$Precision@K = \frac{\sum_{i=1}^{N} |R(i) \cap T(i)|}{\sum_{i=1}^{N} |R(i)|}. \tag{5}$$

Recall@K:

$$Recall@K = \frac{\sum_{i=1}^{N} |R(i) \cap T(i)|}{\sum_{i=1}^{N} |T(i)|}. \tag{6}$$

F1@K:

$$F1@K = \frac{2 \times Precision@k \times Recall@k}{Precision@K + Recall@K}. \tag{7}$$

These metrics comprehensively evaluate the model's performance in recommending relevant syndromes from multiple candidates. Experimental results show that when $K = 5$, the model achieves a Precision@5 of 81.5%, indicating a high level of accuracy. Moreover, as the value of K increases, the corresponding Precision@K, Recall@K, and F1@K also exhibit a continuous upward trend. The detailed performance metrics under varying values of K are illustrated in Figure 8.

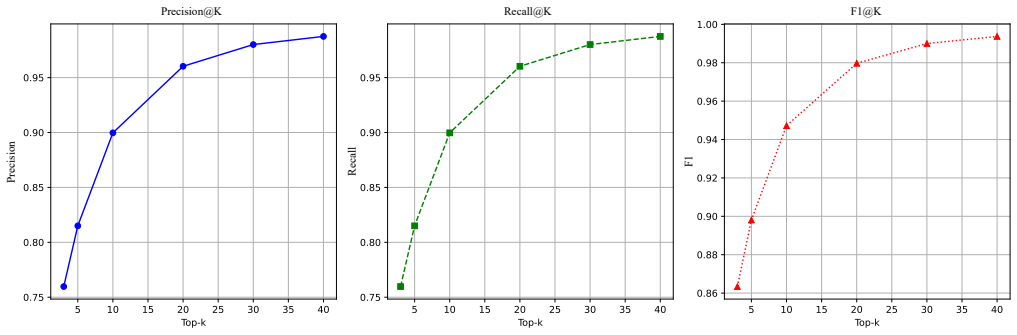


Figure 8: Performance of the syndrome differentiation model across different values of K.

The diagnostic system developed based on the TCM Arthralgia Syndrome Knowledge Graph (TCMASKG) demonstrates significant progress in the field of personalized Traditional Chinese Medicine (TCM) healthcare. It provides diversified syndrome differentiation schemes tailored to the individual clinical presentations of patients, effectively enhancing diagnostic and therapeutic accuracy by integrating physicians' domain knowledge with systematic symptom analysis. By leveraging the Naive Bayes algorithm for noise reduction, the system filters out erroneous predictions and improves the overall reliability of syndrome differentiation.

Additionally, the system offers interpretable explanations of disease pathogenesis, helping physicians gain a deeper understanding of patient conditions and facilitating the formulation of more precise and targeted treatment strategies. Compared to traditional diagnostic approaches, this system not only enhances diagnostic and therapeutic precision but also strengthens physicians' ability to apply syndrome differentiation theory in clinical practice, thereby representing a meaningful step toward intelligent and personalized healthcare in TCM.

4.3. Treatment and Prescription Recommendations

In TCM, the integration of disease identification and syndrome differentiation constitutes the core principle of diagnosis and treatment. Physicians are required not only to identify the disease but also to analyze its nature, etiology, and pathogenesis through the lens of syndrome differentiation. This methodology is essential for achieving a fundamental understanding and effective treatment, especially for complex and refractory cases. For example, in the treatment of ankylosing spondylitis, a symptom such as cold-aggravated lower back pain determines whether the condition should be categorized as "cold obstruction" or "damp-heat obstruction" thereby exemplifying the core concept of treatment based on syndrome differentiation.

This study adopts the same rationale and modeling approach for both syndrome differentiation and disease identification. By combining relational data assessment with the Naive Bayes

model, the correlation between symptoms and diseases can be quantified, enabling the identification of the most probable disease for a given symptom set. Furthermore, analysis of electronic medical records (EMRs) reveals a one-to-one correspondence between prescription names and disease diagnoses. While a single disease may be associated with multiple prescriptions in broader data, the one-to-one mapping extracted from EMRs serves as a useful reference for validating relationships between syndromes, diseases, and prescriptions, thereby improving the reliability of prescription recommendation.

The implementation process of the prescription recommendation algorithm is as follows:
Input: A set of patient symptoms, denoted as $SymptomList = \{S_1, S_2, S_3, \ldots, S_n\}$.
Output: A set of the most appropriate prescriptions for the patient.
Algorithm Steps:

1. Input *SymptomList* into the syndrome differentiation module to identify the most probable syndromes, and store the results in list Z.
2. Use list Z to query the knowledge graph for associated prescriptions D_1.
3. **If** D_1 is not empty:
   (a) Sort the retrieved prescriptions by relational relevance score.
   (b) Select the top 3 prescriptions and store them in *ZtoDlist*.
4. Use the Naive Bayes model to infer the most probable diseases from *SymptomList*, and store them in list B.
5. Use list B to query the knowledge graph for associated prescriptions D_2.
6. **If** D_2 is not empty:
   (a) Sort the retrieved prescriptions by relational relevance score.
   (b) Select the top 3 prescriptions and store them in *BtoDlist*.
7. Return *ZtoDlist* and *BtoDlist* as the final recommendation results.

When symptoms such as skin itching and swelling are input, the syndrome differentiation module may identify "Heat Toxin Accumulation" as the most likely syndrome, while the Naive Bayes model could generate the top three Western medical diagnoses: Systemic Lupus Erythematosus, Pyoderma, and Rheumatoid Arthritis. Based on the knowledge graph, both the syndrome and the diseases will have corresponding prescription sets. Physicians can compare the recommendations from both perspectives and select the most suitable prescription to recommend to the patient. The functional workflow of this treatment recommendation process is illustrated in Figure 9.

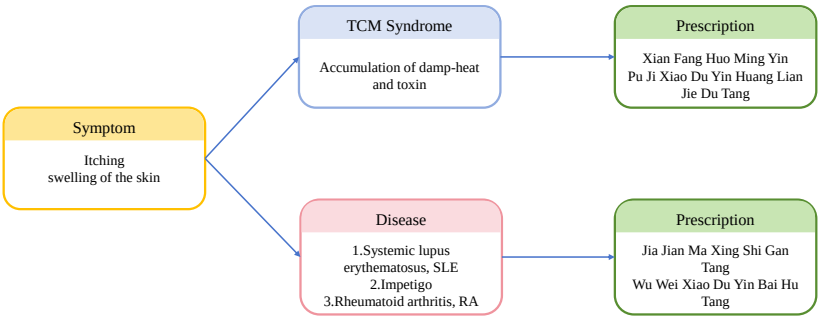


Figure 9: Example of the process of treatment and prescription recommendations.

In comparing tetrad modeling with traditional triad modeling, we found that the performance of the syndrome differentiation and treatment model using tetrad modeling is significantly better than the model using only triad modeling. This result is verified under different amounts of data. Specifically, when using a small amount of data, the syndrome differentiation and treatment model predicts more accurately than the triad model. As the amount of data increases, the performance advantage of the syndrome differentiation and treatment model becomes more prominent, with its predicted results much more accurate than the triad model. This indicates that adding the concept of tetrad during the modeling process can better explore the relationship between data and improve the performance of the model.

4.4. Expert Evaluation

To evaluate the reliability and practical applicability of our treatment system, we conducted a structured expert evaluation, a method widely adopted in various domains. We designed a 10-point Likert scale assessment covering ten quality dimensions across the data, schema, and application layers. Five experts in Traditional Chinese Medicine (TCM) were invited to participate as evaluators.

During the evaluation process, the system's construction methodology and core functionalities were presented to the experts, who were guided through the use of the knowledge graph for syndrome differentiation in TCM. Each expert anonymously scored the system on a scale of 1 to 10 for each dimension, where 1 indicated the lowest quality and 10 the highest. The results are summarized in Table 5.

Table 5: TCMASKG Quality Evaluation Scores.

Main Dimension	**Dimension**	**Metrics**
Data layer	Data completeness	8.21
	Data accuracy	8.16
	Data volume	6.15
Schema layer	Ontology Design	9.20
	Semantic Consistency	8.41
	Interoperability	7.41
Application layer	Consistency of results	8.14
	Ease of use	8.21
	Interpretability of results	8.72
	Knowledge Discovery	7.51

The expert panel awarded an average quality evaluation score of 8.01, signaling their positive appraisal of the treatment system's overall efficacy. In the domain of data integrity, the experts accorded a high score of 8.21, acknowledging the credibility of the data sources employed within the system. Contrastingly, the assessment of data volume yielded a lower score of 6.15, highlighting the impact of insufficient clinical electronic medical records on the system's usability and objectivity.

At the ontology layer, scores for ontology design and semantic consistency exceeded 9.20 and 8.41, respectively, markedly surpassing the median level. These findings suggest that the expert consensus views the system's ontology structure as excelling in rationality and expandability. In the application layer, the high rating of 8.14 for result consistency denotes superior performance

in data coherence maintenance. The scores for usability and result interpretability, exceeding 8.21 and 8.72, respectively, reflect the system's capacity to deliver understandable and highly accurate outcomes. Nonetheless, a knowledge discovery score below 3.56 points to the need for enhancements in the system's real-world clinical applications.

Although preliminary, these assessment outcomes offer a roadmap for refining our treatment system, underscoring the imperative of a broader evaluative framework that transcends technical considerations. In pursuit of augmenting the system's efficacy and utility in diagnosis and treatment, ongoing collaboration with experts, comprehensive research, and iterative refinements are envisaged.

5. Discussion

5.1. Construction of High-Quality Knowledge Graphs

Our research makes a meaningful contribution to the construction of high-quality knowledge graphs, surpassing the quality and consistency typically achieved through web scraping or heterogeneous data aggregation methods. By utilizing authoritative Traditional Chinese Medicine (TCM) case studies as the foundation, we meticulously categorized core entities—including diseases, symptoms, etiologies, and pathogenesis—and established their semantic relationships.

To enhance the clarity and structural coherence of the graph, we constructed a domain-specific ontology that formalizes the representation of TCM knowledge. In addition, we applied advanced data cleansing techniques such as quadruple modeling, co-occurrence probability, and specificity filtering to reduce noise and eliminate redundancy. These strategies significantly improved the internal consistency and overall precision of the knowledge graph.

Nevertheless, expert evaluations highlighted that the current volume of data remains limited. This reflects the need to scale up our dataset through further data acquisition and broader integration of clinical cases. As a next step, we plan to expand our collaboration with TCM hospitals and practitioners to collect larger, more diverse clinical datasets and support continuous updates to the knowledge base.

5.2. An Interpretable Model for Syndrome Differentiation and Treatment

In the area of syndrome differentiation, our study proposes an interpretable framework that explicitly incorporates etiology and pathogenesis as explanatory variables. This not only reveals the underlying mechanisms behind clinical manifestations but also enhances the specificity and reliability of treatment recommendations. The combination of syndrome differentiation and diagnostic reasoning enables the generation of logically coherent and personalized prescriptions.

By linking clinical symptoms to treatment strategies through well-structured causal and semantic relationships, our model offers improved transparency and decision support over traditional black-box approaches. This helps clinicians and patients alike to better understand and trust the system's recommendations [31].

However, the relatively modest expert score for knowledge discovery suggests that the system's current ability to uncover novel clinical insights is limited. To address this, we aim to incorporate data mining techniques such as association rule mining and causal inference in future work. These enhancements will enable more robust knowledge discovery.

6. Conclusions

This study innovatively integrates knowledge graphs with TCM to develop an interpretable and transparent treatment system for TCM arthralgia syndrome, capable of delivering personalized therapeutic recommendations. By leveraging knowledge graph technology and incorporating etiology and pathogenesis for clarity, the system enhances the reasoning from symptoms through TCM syndrome differentiation to precise prescription recommendations. Optimization techniques such as quadruples, co-occurrence probability, and specificity refine the algorithm, reducing noise and boosting accuracy. The constructed knowledge graph contains 4,515 entity nodes and 90,776 semantic relations, forming a robust knowledge base that supports accurate inference and decision-making. Experimental outcomes affirm the method's efficacy in offering highly accurate and practical personalized treatment plans. This research not only highlights the potential of knowledge graphs in medical treatment but also lays the foundation for future explorations aimed at enhancing the efficiency and effectiveness of medical therapies.

While the findings demonstrate significant potential, some challenges persist. These include the need for large-scale and high-quality data collection, further optimization and validation of the syndrome differentiation module, consideration of individual patient variability and treatment efficacy, and the necessity to verify the knowledge graph's applicability in broader clinical settings. Moving forward, robust collaborations with healthcare institutions and clinicians will be essential to gather diverse real-world clinical data for further empirical analysis and evaluation.

Author Contributions

Conceptualization: Kai Chang, Wenping Deng, and Yong Xiao; Methodology: Kai Chang; Data Acquisition and Collation: Liangzhi Yuan and Jiali Hai; Data Curation and Analysis: Liangzhi Yuan; Writing—Original Draft Preparation: Kai Chang and Liangzhi Yuan; Writing—Review and Editing: Kai Chang and Yong Xiao; Project Administration: Liangzhi Yuan and Jiali Hai; Funding Acquisition: Kai Chang and Yong Xiao. All authors have read and agreed to the published version of the manuscript.

Acknowledgements

The research was supported in part by the Natural Science Foundation Innovation and Development Joint Fund of Hubei (No. 2025AFD574), the Specialized Information Standards of the State Administration of Traditional Chinese Medicine (No.GHC-2024-ZFGM-33), and National Natural Science Foundation of China (No. 82374302).

References

[1] CJ Hogeboom, KJ Sherman, and DC Cherkin. Variation in diagnosis and treatment of chronic low back pain by traditional chinese medicine acupuncturists. *Complementary Therapies in Medicine*, 9(3):154–166, 2001.

[2] Amy K Saenger and Robert H Christenson. Stroke biomarkers: progress and challenges for diagnosis, prognosis, differentiation, and treatment. *Clinical chemistry*, 56(1):21–33, 2010.

[3] RE Outerbridge. The etiology of chondromalacia patellae. *The Journal of bone and joint surgery. British volume*, 43(4):752–757, 1961.

[4] J Adam Rindfleisch and Daniel Muller. Diagnosis and management of rheumatoid arthritis. *American family physician*, 72(6):1037–1047, 2005.

[5] Xiaojun Chen, Shengbin Jia, and Yang Xiang. A review: Knowledge reasoning over knowledge graph. *Expert Systems with Applications*, 141:112948, 2020.

[6] Gregory K Robbins, William Lester, Kristin L Johnson, Yuchiao Chang, Gregory Estey, Dominic Surrao, Kimon Zachary, Sara M Lammert, Henry C Chueh, James B Meigs, et al. Efficacy of a clinical decision-support system in an hiv practice: a randomized trial. *Annals of internal medicine*, 157(11):757–766, 2012.

[7] SM Shamimul Hasan, Donna Rivera, Xiao-Cheng Wu, Eric B Durbin, J Blair Christian, and Georgia Tourassi. Knowledge graph-enabled cancer data analytics. *IEEE journal of biomedical and health informatics*, 24(7):1952–1967, 2020.

[8] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 1247–1250, 2008.

[9] M Fabian, Kasneci Gjergji, WEIKUM Gerhard, et al. Yago: A core of semantic knowledge unifying wordnet and wikipedia. In *16th International world wide web conference, WWW*, pages 697–706, 2007.

[10] Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam Hruschka, and Tom Mitchell. Toward an architecture for never-ending language learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 24, pages 1306–1313, 2010.

[11] Jing Xie, Xin Li, Ye Yuan, Yi Guan, Jingchi Jiang, Xitong Guo, and Xin Peng. Knowledge-based dynamic prompt learning for multi-label disease diagnosis. 286:111395.

[12] Xiaolong Li, Shuifa Sun, Tinglong Tang, Ji Lu, Lijuan Zhang, Jie Yin, Qian Geng, and Yirong Wu. Construction of a knowledge graph for breast cancer diagnosis based on Chinese electronic medical records: Development and usability study. 23(1):210.

[13] Maya Rotmensch, Yoni Halpern, Abdulhakim Tlimat, Steven Horng, and David Sontag. Learning a Health Knowledge Graph from Electronic Medical Records. 7(1):5994.

[14] Zhixue Jiang, Chengying Chi, and Yunyun Zhan. Research on medical question answering system based on knowledge graph. *IEEE Access*, 9:21094–21101, 2021.

[15] Lino Murali, G Gopakumar, Daleesha M Viswanathan, and Prema Nedungadi. Towards electronic health record-based medical knowledge graph construction, completion, and applications: A literature study. *Journal of biomedical informatics*, 143:104403, 2023.

[16] Rui Yang, Qing Ye, Chunlei Cheng, Suhua Zhang, Yong Lan, Jing Zou, et al. Decision-making system for the diagnosis of syndrome based on traditional chinese medicine knowledge graph. *Evidence-Based Complementary and Alternative Medicine*, 2022, 2022.

[17] Rongxuan Gao and Chen Li. Knowledge question-answering system based on knowledge graph of traditional chinese medicine. In *2020 IEEE 9th Joint International Information Technology and Artificial Intelligence Conference (ITAIC)*, volume 9, pages 27–31. IEEE, 2020.

[18] Yichong Zhang and Yongtao Hao. Traditional Chinese Medicine Knowledge Graph Construction Based on Large Language Models. 13(7):1395.

[19] Ziqiang Zheng, Yongguo Liu, Yun Zhang, and Chuanbiao Wen. TCMKG: A Deep Learning Based Traditional Chinese Medicine Knowledge Graph Platform. In *2020 IEEE International Conference on Knowledge Graph (ICKG)*, pages 560–564. IEEE.

[20] Tong Yu, Jinghua Li, Qi Yu, Ye Tian, Xiaofeng Shun, Lili Xu, Ling Zhu, and Hongjie Gao. Knowledge graph for tcm health preservation: Design, construction, and applications. *Artificial intelligence in medicine*, 77:48–52, 2017.

[21] Xiaoli Chu, Simin Wu, Bingzhen Sun, and Qingchun Huang. Data-driven quantification and intelligent decision-making in traditional chinese medicine: a review. *International Journal of Machine Learning and Cybernetics*, pages 1–16, 2024.

[22] Alex John London. Artificial intelligence and black-box medical decisions: accuracy versus explainability. *Hastings Center Report*, 49(1):15–21, 2019.

[23] Qiaoying Teng, Zhe Liu, Yuqing Song, Kai Han, and Yang Lu. A survey on the interpretability of deep learning in medical diagnosis. 28(6):2335–2355.

[24] Ziming Yin, Zhongling Kuang, Haopeng Zhang, Yu Guo, Ting Li, Zhengkun Wu, and Lihua Wang. Explainable AI Method for Tinnitus Diagnosis via Neighbor-Augmented Knowledge Graph and Traditional Chinese Medicine: Development and Validation Study. 12:e57678.

[25] Xiaolei Xiu, Qing Qian, and Sizhu Wu. Construction of a Digestive System Tumor Knowledge Graph Based on Chinese Electronic Medical Records: Development and Usability Study. 8(10):e18287.

[26] World Federation of Chinese Medicine Societies. National service platform for academic experience of famous tcm doctor. Available at: `http://123.57.245.75/`, 2012.

[27] Qunsheng Zou, Kuo Yang, Kai Chang, Xiaoping Zhang, Xiaodong Li, and Xuezhong Zhou. Phenonizer: A fine-grained phenotypic named entity recognizer for chinese clinical texts. In *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 3963–3970. IEEE, 2021.

[28] Jonathan Lawry, Jim W Hall, and R Bovey. Fusion of expert and learnt knowledge in a framework of fuzzy labels. *International Journal of Approximate Reasoning*, 36(2):151–198, 2004.

[29] International Digital Economy Academy. Biomedical informatics ontology system. Available at: `https://bios.idea.edu.cn`, 2023.

[30] Linfeng Li, Peng Wang, Jun Yan, Yao Wang, Simin Li, Jinpeng Jiang, Zhe Sun, Buzhou Tang, Tsung-Hui Chang, Shenghui Wang, et al. Real-world data medical knowledge graph: construction and applications. *Artificial intelligence in medicine*, 103:101817, 2020.

[31] Qing Ye, Rui Yang, Chun-lei Cheng, Lin Peng, and Yong Lan. Combining the external medical knowledge graph embedding to improve the performance of syndrome differentiation model. *Evidence-Based Complementary and Alternative Medicine*, 2023(1):2088698, 2023.