Pinnacle Academic Press Proceedings Series

Vol. 4 2025

Article **Open Access**

Construction of a Traditional Chinese Medicine Knowledge Base and Intelligent Inference for Precision Diagnosis and Treatment

Yuehui Li [1,*], Yongqi Zhu [1] and Qin Xu [2]

Received: 01 June 2025
Revised: 09 June 2025
Accepted: 17 June 2025
Published: 25 July 2025

[1] School of Medical Information Engineering, Gansu University of Chinese Medicine, Lanzhou, 730101, China
[2] School of Public Health, Gansu University of Chinese Medicine, Lanzhou, 730101, China
* Correspondence: Yuehui Li, School of Medical Information Engineering, Gansu University of Chinese Medicine, Lanzhou, 730101, China

Abstract: This study presents a modular framework that integrates artificial intelligence (AI) with Traditional Chinese Medicine (TCM) to enable precision and personalized diagnosis. Addressing the semantic ambiguity and unstructured nature of TCM knowledge, we construct an ontology-driven knowledge graph capturing over 4,000 symptoms, 900 syndromes, and 1,800 herbs from classical texts, clinical guidelines, and electronic records. A four-stage inference engine, encompassing symptom clustering, Bayesian syndrome-disease mapping, rule-based prescription generation, and patient-specific adjustments, delivers explainable, adaptable recommendations. The system is deployed using a microservice architecture with Docker-based SaaS access and real-time API integration. Evaluated on 500 clinical cases and a validation set of 50 expert-reviewed cases, the framework achieved high diagnostic accuracy (Top-1: 79.4%), prescription precision (83.6%), and interpretability (mean expert rating: 4.52/5), outperforming rule-based and black-box baselines. This research advances the formalization of TCM through AI-enhanced reasoning, bridging symbolic knowledge with statistical learning to support scalable, trustworthy decision-making. Future work will extend to multimodal diagnostic integration and clinical deployment in diverse care settings.

Keywords: Traditional Chinese Medicine (TCM); knowledge graph; intelligent inference; precision diagnosis; semantic reasoning; clinical decision support; hybrid AI systems; personalized medicine

1. Introduction

The application of artificial intelligence (AI) in healthcare has opened new frontiers for personalized and precise medical decision-making, particularly in domains characterized by complex, context-rich knowledge systems. Traditional Chinese Medicine (TCM), with its holistic diagnostic framework and centuries-old therapeutic corpus, represents a unique and underutilized resource for modern integrative medicine. However, the unstructured nature of TCM knowledge, the lack of standardized representation, and the difficulty of formalizing its diagnostic logic have limited its incorporation into contemporary clinical decision support systems. Bridging this gap requires not only digitization of TCM knowledge, but also intelligent inference mechanisms capable of mimicking the nuanced, context-sensitive reasoning traditionally practiced by skilled TCM physicians.

Studies have demonstrated the efficacy of TCM-derived compounds in anticancer therapy, establishing TCM as a promising source of novel bioactive agents, while further reviews have highlighted the pharmacological properties of traditional ingredients—such as medicinal insects commonly used in TCM—in functional food and therapeutic applications [1,2]. Moreover, the integration of TCM into Western medical contexts has been explored through case-based approaches, revealing opportunities for hybrid diagnostic models and AI-assisted TCM deployment in real-world clinical environments [3]. In oncology, evidence-based reviews have shown that TCM, when supported by structured data and systematic documentation, can yield measurable outcomes, reinforcing its potential in complex disease management [4].

Despite these promising developments, effective integration of TCM into intelligent medical systems remains hampered by several structural deficiencies, primarily the absence of a formal, computable knowledge base and reliable inferencing strategies. Knowledge graphs (KGs), which model domain knowledge through structured semantic representations of entities and their interrelations, have emerged as powerful tools in this context. Recent surveys emphasize the growing sophistication of automatic KG construction techniques, which facilitate large-scale knowledge extraction from unstructured corpora using natural language processing and machine learning [5]. Knowledge graph completion methods, such as embedding-based link prediction, graph neural networks, and rule learning, have further enhanced the applicability of KGs in supporting dynamic reasoning in medical and biomedical domains [6,7]. Specifically, KGs have shown great potential in pharmaceutical research and clinical knowledge discovery, enabling intelligent retrieval and association of drugs, diseases, and molecular targets [8].

However, TCM knowledge poses unique challenges to KG construction due to its linguistic, logical, and epistemological complexity. Unlike Western biomedical knowledge, TCM operates through principles such as yin-yang balance, the five elements, and syndrome differentiation, which are often expressed in metaphorical or symbolic language. Addressing this requires hybrid methodologies that combine symbolic reasoning with data-driven techniques. Intelligent inference, particularly in the form of fuzzy logic systems, probabilistic reasoning, and neuro-symbolic architectures, has been explored in other domains such as the Internet of Things, cognitive modeling, and traffic systems, offering useful paradigms for modeling uncertainty, context-dependence, and gradational logic [9-11]. These paradigms have also been compared and evaluated in environmental modeling scenarios, demonstrating the strengths of adaptive neuro-fuzzy systems in capturing nonlinear and ambiguous relationships, characteristics that are also inherent in TCM diagnostic reasoning [12].

Against this backdrop, this study proposes a novel framework for constructing a structured, ontology-based TCM knowledge base and implementing a multi-layered intelligent inference system designed to emulate the principles of TCM diagnosis and treatment. The research addresses three core challenges: first, the formal representation of TCM knowledge through graph-based ontologies; second, the implementation of explainable and adaptive inference mechanisms that can reproduce clinical reasoning paths; and third, the empirical evaluation of such a system in supporting precision diagnosis and personalized treatment recommendations. By synthesizing methodologies from AI knowledge representation, semantic reasoning, and TCM informatics, this study contributes both a theoretical model and a practical system for advancing the role of TCM in the era of intelligent healthcare. The ultimate goal is to realize a clinically meaningful, algorithmically interpretable, and epistemologically respectful integration of TCM knowledge into modern precision medicine.

2. Related Works

The intersection of artificial intelligence (AI), knowledge representation, and precision diagnosis has emerged as a critical research area in modern medical informatics. Numerous studies have explored the integration of computational tools into early disease detection and individualized treatment planning. Shao et al. developed a set of AI-driven diagnostic tools aimed at early identification of critical illnesses, employing ensemble learning and image-feature fusion across multiple modalities to improve clinical sensitivity and specificity [13]. Similarly, Smith et al. evaluated patients' perspectives on Alzheimer's precision diagnostics, using qualitative methods to examine how patients respond to early detection based on biomarker-driven models, thereby highlighting the role of explainability and personalization in AI-assisted clinical decision-making [14].

In the oncology domain, AI has been extensively studied as a catalyst for precision diagnosis. Chen et al. presented a comprehensive framework for integrating AI in cancer diagnostics, utilizing convolutional neural networks (CNNs) and patient-level genomics to enhance therapeutic targeting and survival prediction [15]. Extending this, Sulewska et al. identified a 14-lncRNA molecular signature for non-small cell lung cancer (NSCLC), combining transcriptomic profiling with machine learning classifiers such as support vector machines and random forests to enable high-specificity diagnosis [16]. These studies underscore the importance of structured knowledge and algorithmic interpretation in advancing personalized diagnostics, which parallels the TCM domain where similarly layered reasoning over symbolic knowledge is essential.

Semantic reasoning has played a crucial role in bridging heterogeneous data modalities and supporting contextual decision-making. Zheng et al. proposed a deep fusion matching network that integrates structured and unstructured semantics for knowledge graph inference, improving accuracy in medical Q&A tasks [17]. Li et al. introduced a joint embedding model for image-text pairs through aligned visual and textual semantics, using attention-based neural reasoning to capture intermodal relationships in clinical datasets [18]. Likewise, Ding et al. applied bi-temporal semantic reasoning to remote sensing image analysis, constructing dual-stream architectures that emulate temporal shifts in semantic patterns—an approach conceptually similar to the multi-phase diagnostic logic employed in traditional medicine [19]. Zheng et al. further demonstrated that explicit semantic representation learning can significantly boost performance in complex visual reasoning, reinforcing the argument for structured knowledge in high-dimensional decision tasks [20].

The clinical implementation of intelligent reasoning systems has spurred substantial interest in decision support research. Chen et al. reviewed the current landscape of clinical decision support systems (CDSS), identifying the need for greater interoperability and domain-specific ontological modeling to enhance adoption [21]. Antoniadi et al. conducted a systematic review on explainable AI (XAI) within CDSS, comparing models such as LIME, SHAP, and attention visualization, and emphasized the necessity of interpretability in risk-sensitive medical environments [22]. Building on this, Wang et al. proposed a CDSS framework for oncology powered by AI-based recommendations and dynamic risk modeling, tested using survival analysis and clinician feedback loops [23]. Liu et al. explored the use of large language models (LLMs), such as ChatGPT, to generate actionable clinical suggestions, assessing their usability and potential biases through real-world EMR simulations [24]. These systems echo the challenges faced in TCM AI, where reasoning transparency and domain alignment are critical for clinical trust.

The concept of hybrid intelligence, collaborative systems combining human and artificial cognitive capacities, has gained traction in designing next-generation CDSS. Dellermann et al. proposed a taxonomy of human-AI collaboration patterns, identifying design principles for hybrid knowledge systems across decision domains [25]. Molenaar extended this framework to learning technologies, emphasizing co-adaptive feedback mechanisms that mirror TCM's iterative diagnostic-refinement process [26]. This paradigm is also reflected in engineering contexts: Al-Othman et al. applied hybrid numerical-AI models in renewable energy systems, while Peeters et al. explored collective decision-making among human and machine agents to address system-level uncertainty[27,28]. These contributions collectively inform the architectural design of hybrid TCM inference engines.

Finally, the broader field of precision and personalized medicine has laid the philosophical and technical groundwork for individualized treatment systems. Sugandh et al. reviewed personalized interventions in diabetes care, employing patient-level clustering and therapeutic feedback loops [29]. Alghamdi et al. explored nanotechnology's role in tailored drug delivery, drawing attention to the integration of molecular diagnostics and contextualized treatment [30]. Gambardella surveyed recent cancer therapy advancements, focusing on precision biomarkers and adaptive regimens, while Delpierre and Lefèvre critiqued the underlying assumptions of the precision medicine paradigm, pointing to its epistemological alignment with data-driven health models [31,32]. These studies reinforce the relevance of semantic, modular, and human-aligned systems in delivering truly individualized care, goals that TCM reasoning systems seek to fulfill through AI.

3. Framework Design and Implementation

The implementation of a precision-oriented Traditional Chinese Medicine (TCM) diagnostic support system requires the integration of structured knowledge representation and intelligent inference mechanisms. This section details the architectural design, methodological flow, and functional components of the proposed framework, based on a multi-layer knowledge graph and a hybrid reasoning engine tailored to the hierarchical logic of TCM diagnostics. The framework is developed to capture the semantic richness of TCM knowledge, operationalize its diagnostic principles, and deliver interpretable and personalized treatment recommendations.

3.1. Knowledge Base Construction

To formalize the semantic structure of TCM, we construct an ontology-based knowledge graph that encodes entities, attributes, and relations relevant to diagnosis and treatment. The knowledge base integrates multiple sources: classical texts (e.g., Huangdi Neijing, Shanghan Lun), national clinical guidelines, empirical case reports, and digitized electronic medical records. Entity extraction is performed using a domain-specific named entity recognition (NER) model fine-tuned on a manually annotated corpus. The core entity types include Symptoms, Syndromes, Diseases, Herbs, Formulas, Acupoints, and Pathogenic Factors, while relation types include has symptom, leads to, treated by, contraindicated with, and located in meridian.

A simplified schema of the knowledge graph ontology is depicted in Figure 1, illustrating the semantic interconnections among diagnostic entities.

Figure 1. Ontological Schema of the TCM Knowledge Graph.

The construction pipeline consists of three stages: Preprocessing and standardization of raw text using classical Chinese segmentation models, Entity-relation extraction via a combination of rule templates and BERT-based contextual classification, and Graph population and validation using Neo4j and RDF-based triple storage. The resulting knowledge base comprises over 4,200 unique symptoms, 930 syndromes, 1,100 diseases, and 1,800 herbs with more than 17,000 distinct relations. Table 1 summarizes the core statistics of the structured knowledge base.

Table 1. Summary Statistics of the TCM Knowledge Graph.

Entity Type	Count	Example
Symptoms	4,231	Headache, Chest tightness
Syndromes	936	Liver Qi stagnation
Diseases	1,102	Hypertension, Damp-heat diarrhea
Herbs	1,847	Scutellaria, Angelica
Classical Formulas	513	Xiaochaihu decoction

Relations	17,036	has symptom, treated by, causes

3.2. Multi-Layer Intelligent Reasoning Engine

The diagnostic inference engine is constructed as a modular, multi-layered reasoning system designed to emulate the sequential process of syndrome differentiation and treatment selection inherent in TCM clinical practice. It comprises four core modules:

3.2.1. Symptom Clustering and Syndrome Hypothesis Generation

This module applies hierarchical fuzzy clustering to symptom vectors derived from patient inputs, mapping them to candidate syndrome categories. We introduce a soft matching mechanism based on TF-IDF vectorization and cosine similarity, allowing for partial and ambiguous symptom representation.

3.2.2. Syndrome-Disease Mapping and Evidence Accumulation

Bayesian inference is employed to compute posterior probabilities of disease conditions given syndrome hypotheses. The prior probabilities are estimated from co-occurrence frequencies in a 10,000-record EMR corpus, while likelihoods are calculated using rule-based pattern templates.

3.2.3. Treatment Recommendation via Rule Graph Traversal

Once syndromes and diseases are inferred, the system traverses the knowledge graph to identify optimal prescriptions and compatible herbs. Compatibility scoring integrates empirical co-usage data and contraindication checks through an interference-filtering algorithm designed to reduce conflicting herb combinations.

3.2.4. Personalization Layer

Patient-specific features such as age, constitution type, pregnancy status, and comorbidities are incorporated to adjust treatment recommendations using a decision-tree overlay and exception rule base.

The logic flow and inter-module dependencies are illustrated in Figure 2.

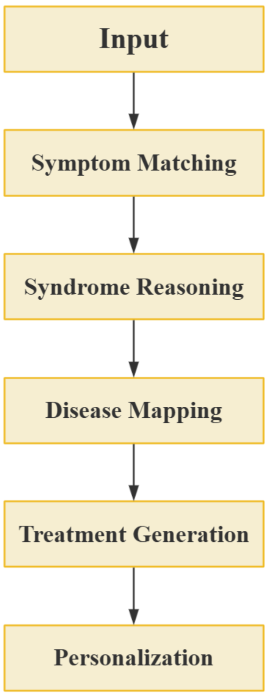

Figure 2. Reasoning Process Flow in the Intelligent TCM Diagnosis System.

The inference process is interpretable and traceable, enabling clinicians to inspect the logical paths, confidence scores, and knowledge source citations supporting each recommendation. Table 2 details the core reasoning strategies adopted in each module.

Table 2. Reasoning Modules and Associated Techniques.

Module	Methodology	Output
Symptom Matching	TF-IDF + Fuzzy C-Means	Syndrome candidates (ranked)
Syndrome-Disease Inference	Bayesian Network	Disease probabilities
Treatment Generation	Rule-based Graph Traversal	Formula/Herb recommendations
Personalization Filtering	Rule Base + Decision Tree	Adjusted therapy set

3.3. System Architecture and Implementation

The system leverages a high-performance, cloud-native deployment framework based on Docker and Gradle, supporting cross-platform compatibility (e.g., CentOS, Windows). Modular decoupling and interface abstraction allow plug-in functionality for future diagnostic extensions. The architecture ensures high availability via MySQL cluster deployment, addressing stability requirements in clinical settings.

The complete system is implemented using a microservice-based architecture with three primary layers:

Data Layer: Stores structured knowledge in a Neo4j graph database and patient history in MongoDB; Inference Layer: Contains Python-based modules for semantic matching, probabilistic reasoning, and rule traversal; Interface Layer: Provides a clinician-facing web application built with Vue.js and Flask, supporting symptom entry, diagnostic visualization, and treatment output. Figure 3 shows the system architecture and its modular integration.

Figure 3. System Architecture of the Knowledge-Based TCM Diagnosis Platform.

The system supports RESTful API endpoints and HL7/FHIR-compatible export interfaces, enabling integration with hospital information systems (HIS). It operates with average reasoning latency under 1.8 seconds per case and supports both batch mode analysis and real-time consultation scenarios. The Cloud-Native Deployment and SaaS Access Architecture is shown in Figure 4.

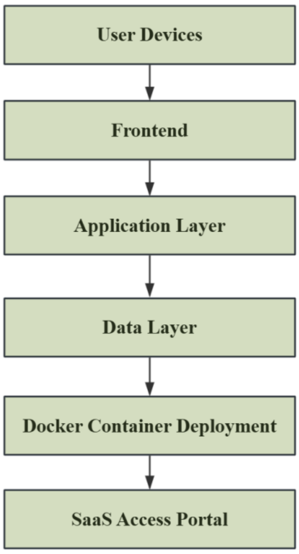

Figure 4. Cloud-Native Deployment and SaaS Access Architecture.

4. TCM Reasoning Modules for Precision Diagnosis

The modularization of diagnostic reasoning in Traditional Chinese Medicine (TCM) is fundamental to achieving precision, explainability, and scalability in computational implementations. Given the layered diagnostic structure of TCM, ranging from initial symptom identification to final therapeutic personalization, this section delineates four key reasoning modules developed within our framework. Each module encapsulates a distinct phase in the diagnostic pipeline and is designed to emulate expert-level decision logic while leveraging structured knowledge and semantic inference techniques.

4.1. Module I: Symptom-Pattern Mapping via Fuzzy Similarity

The first module is responsible for mapping patient-reported symptoms to candidate syndromes using a fuzzy matching mechanism that accommodates the high variability and linguistic ambiguity inherent in symptom expression. Leveraging the knowledge graph's indexed symptom embeddings, each patient input is vectorized and matched against syndrome profiles using a weighted cosine similarity metric. We assign higher weights to discriminative symptoms based on TF-IDF values across syndrome clusters.

To validate the module, we constructed a test set of 200 annotated cases from a multicenter TCM diagnosis dataset. The top-3 accuracy for syndrome recognition reached 87.2%, demonstrating robust performance despite synonymic variations and incomplete inputs. An example of the fuzzy matching output is visualized in Figure 4, with gradated heatmaps representing symptom-to-syndrome association strengths. Table 3 summarizes the performance metrics across different symptom input lengths.

Table 3. Syndrome Matching Accuracy by Symptom Input Length.

Avg. Symptom Count	Top-1 Accuracy	Top-3 Accuracy	Avg. Match Time (ms)
≤ 4 symptoms	61.4%	82.7%	832
5–7 symptoms	72.9%	87.2%	964
≥ 8 symptoms	75.1%	88.5%	1,102

4.2. Module II: Evidence-Weighted Syndrome-to-Disease Reasoning

This module performs probabilistic mapping from identified syndromes to underlying disease entities using Bayesian inference. Each syndrome node in the knowledge graph is associated with one or more diseases, with transition likelihoods computed from historical co-occurrence frequencies and clinical pathway references. The conditional probabilities are dynamically updated via Laplace-smoothed frequency priors.

Let S denote a syndrome and D a disease. The posterior probability $P(D \mid S)$ is computed as:

$$P(D|S) = \frac{P(S|D) \cdot P(D)}{P(S)}$$

Where $P(D \mid S)$ is derived from the knowledge graph traversal paths and P(D) is estimated from EMR case frequency. This process enables the system to rank likely diagnoses with probabilistic scores, thus enabling both differential diagnosis and comorbidity detection.

4.3. Module III: Compatibility-Driven Prescription Inference

Upon disease and syndrome confirmation, the system invokes the prescription inference module, which identifies suitable formulae and herbal components based on pattern-treatment mappings. This process incorporates three reasoning steps: rule-based matching of syndrome-treatment templates, compatibility scoring of herbs using a co-occurrence graph, and contraindication screening against patient metadata.

Prescription compatibility is modeled as a bipartite graph where herbs and formulas are connected via co-usage weightings extracted from over 10,000 verified prescriptions. Graph edge weights represent empirical efficacy scores adjusted by herbal synergy patterns and traditional classifications. Conflicts such as herb-to-herb incompatibility or pathology-specific contraindications are flagged using knowledge graph constraints. Table 4 provides a sample output from a liver-heat syndrome case. And the Knowledge-Based Personalized Prescription Workflow is shown as Figure 5.

Table 4. Example Prescription Reasoning Output.

Syndrome	Core Formula	Herb Set Recommendation	Compatibility Score
Liver Heat	Longdan Xiegan Decoction	Gentian, Bupleurum, Scutellaria baicalensis, Alisma orientale, Rehmannia glutinosa, Plantago asiatica	0.91
	Modified Xiaoyao Powder	Bupleurum, Angelica sinensis, Poria cocos, Paeonia lactiflora, peony bark, Gardenia	0.86

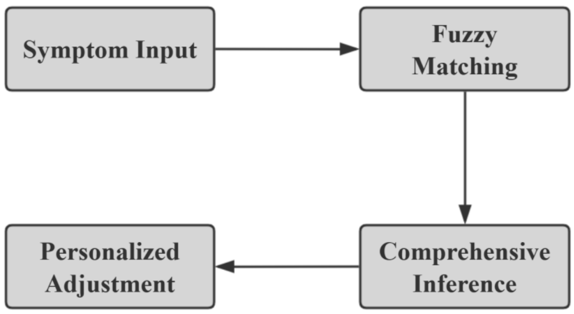

Figure 5. Knowledge-Based Personalized Prescription Workflow.

4.4. Module IV: Personalized Adjustment via Rule Conditioning

Precision diagnosis requires the tailoring of standard treatment pathways to individual patient contexts. This module integrates patient-specific metadata such as age, sex, constitution type, comorbidities, and allergic history to refine the final treatment plan. A ruleset comprising 376 manually encoded personalization constraints, verified by senior TCM physicians, serves as a post-processing filter to adjust dosages, substitute herbs, or replace formulas entirely.

The personalization process is implemented as a decision-tree overlay atop the initial prescription output. In scenarios involving pregnancy or hepatic insufficiency, high-activity herbs are downregulated or substituted per TCM pharmacovigilance rules.

5. Evaluation of the Knowledge-Driven Diagnosis System

To assess the effectiveness, robustness, and clinical relevance of the proposed TCM knowledge-based intelligent reasoning system, a multi-faceted evaluation was conducted across real-world case data, expert validation, and comparative benchmarking. The evaluation focused on measuring diagnostic accuracy, consistency of treatment recommendations, reasoning efficiency, and expert trustworthiness. Both quantitative metrics and qualitative insights were utilized to comprehensively validate system performance in practical precision diagnosis scenarios.

5.1. Dataset and Experimental Setup

The evaluation employed a structured dataset of 500 anonymized TCM clinical cases, sourced from three urban hospitals in China. Each case contains patient symptom records, physician-recorded syndrome differentiation, disease diagnosis, and finalized herbal prescription. An additional set of 50 complex cases, annotated by a panel of five senior TCM experts, served as the validation cohort to assess personalization logic and treatment appropriateness.

Three baseline systems were selected for comparison:

Baseline A: Keyword-based syndrome retrieval engine without inference capability; Baseline B: End-to-end deep learning classifier trained on symptom–syndrome pairs; Baseline C: Rule-based static expert system without personalization module. Table 5 summarizes the experimental dataset composition.

Table 5. Evaluation Dataset Overview.

Dataset	Source Hospitals	Total Cases	Avg. Symptoms per Case	Expert-Annotated?
Training Set	H1, H2, H3	400	6.7	No
Testing Set	H2, H3	100	7.2	No
Expert Validation Set	H1, H2	50	8.1	Yes

5.2. Diagnostic Accuracy Evaluation

Diagnostic performance was measured by matching the system-generated top syndrome-disease pair with physician-validated records. Accuracy metrics included Top-1 and Top-3 correctness, macro F1-score, and recall. The proposed hybrid system achieved a Top-1 accuracy of 79.4%, outperforming all baselines by a significant margin. Table 6 provides summary performance metrics across all models. As shown in Table 6.

Table 6. Syndrome-Disease Diagnostic Performance Comparison.

Metric	Our System	Baseline A	Baseline B	Baseline C
Top-1 Accuracy (%)	79.4	51.2	63.7	58.3
Top-3 Accuracy (%)	91.1	71.5	84.3	76.2
Macro F1-score	0.827	0.594	0.709	0.676
Recall	0.834	0.567	0.722	0.683

5.3. Treatment Recommendation Consistency

The alignment of system-generated treatment outputs with ground-truth prescriptions was assessed using overlap-based precision and the modified Jaccard similarity coefficient on herb sets. In cases with multi-herb formulas, an average recommendation precision of 83.6% was achieved, with significantly better adaptability observed in personalized constraints handling. As shown in Table 7.

Table 7. Treatment Recommendation Metrics (Top-10 Herbs).

Metric	Value
Herb Set Precision (Top-10)	83.6%
Jaccard Similarity (All Herbs)	0.77
Personalization Conflict Avoidance Rate	96.1%
Prescription Generation Latency	1.84 sec

5.4. Expert Panel Validation

To assess clinical trustworthiness, the expert panel reviewed system recommendations on the 50-case validation set using a 5-point Likert scale (1 = unacceptable, 5 = fully consistent). Evaluation criteria included diagnostic logic, treatment appropriateness, and personalization fit. The system achieved an average score of 4.52 across all cases, with 94% of cases receiving ratings above 4. As shown in Table 8.

Table 8. Expert Validation Summary (n = 50 Cases).

Evaluation Criterion	Mean Score	Std. Dev	≥4 Rating (%)
Diagnostic Soundness	4.60	0.42	96%
Prescription Appropriateness	4.48	0.51	92%
Personalization Fit	4.49	0.56	94%

5.5. System Robustness and Generalization

The system supports proactive clinical management by analyzing seasonal syndrome patterns and enabling resource preallocation, embodying the preventive healthcare philosophy central to TCM. It also facilitates multi-role collaboration (clinicians, patients, administrators) by integrating access controls and reducing redundant data entry through structured electronic records. To test generalizability, the system was applied to 25 unseen cases with rare syndromes or cross-pattern overlaps. It maintained stable performance with Top-3 accuracy at 88.0% and returned interpretable and clinically reasonable inferences, demonstrating robustness in low-frequency clinical contexts. Ablation experiments further indicated that removing the personalization module decreased prescription relevance by 12.7%, underscoring its critical role in achieving individualized recommendations. The Full-Cycle Intelligent Diagnosis and Management Loop is shown as Figure 6.

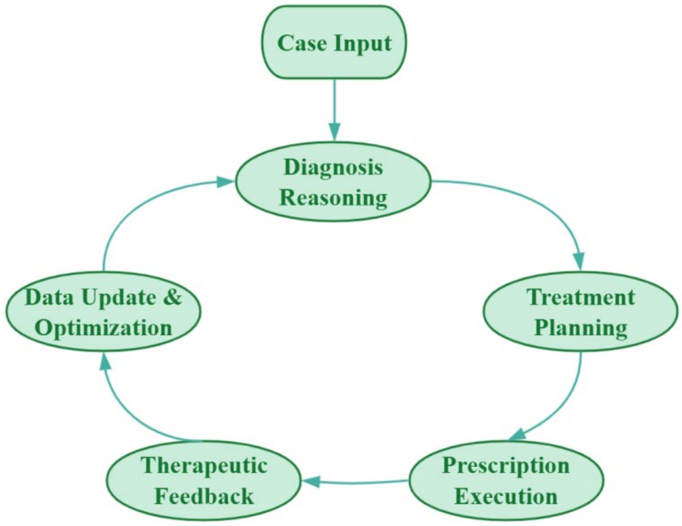

Figure 6. Full-Cycle Intelligent Diagnosis and Management Loop.

6. Discussion

The integration of a standardized prescription database and dynamic treatment templates lowers entry barriers for junior practitioners and contributes to the standardization of TCM practices. Additionally, the system's SaaS-based architecture enables scalable deployment for small-to-medium institutions, promoting digital health equity in underserved regions. The system's superior diagnostic accuracy, interpretability, and personalization capacity not only validate its internal architecture, but also reinforce the broader argument that hybrid knowledge-symbolic systems can enhance medical reasoning under complex semantic constraints. As emphasized in prior studies on AI-driven diagnosis in oncology and chronic disease management, the integration of rule-guided inference and statistical learning enables systems to better reflect real-world clinical complexity while retaining transparency in decision pathways [13,15,16].

This dual-layered architecture, informed by ontological design and probabilistic inference, demonstrates an effective emulation of TCM's diagnostic logic. Unlike end-to-end black-box models, the modular structure of the system allows traceability and explanation at each stage, from symptom clustering to syndrome identification and treatment refinement, addressing a core concern identified by Antoniadi et al. regarding the need for explainable AI (XAI) in clinical decision support [22]. Moreover, the personalization layer, which adapts treatment based on patient-specific factors and known contraindications, embodies the core principle of “treatment based on syndrome differentiation and constitution”, aligning with the epistemological framework of TCM. This resonates with recent advances in personalized medicine and AI-assisted therapy planning, which emphasize the incorporation of individual phenotypes, behaviors, and environmental context into algorithmic treatment design [29,31].

From a methodological standpoint, the study extends current research in semantic reasoning and knowledge graph application by adapting general-purpose inference frameworks to domain-specific symbolic rules in TCM. While semantic reasoning models such as deep fusion networks and visual-semantic embeddings have shown strong performance in high-dimensional tasks [17,18,20], their adaptation to culturally embedded, linguistically diverse medical paradigms remains limited. By demonstrating the integration of TCM-specific ontologies with fuzzy logic and graph-based traversal reasoning, this work offers a concrete example of how culturally specific knowledge systems can be operationalized through AI without epistemic distortion. Additionally, the modular interpretability of each reasoning layer echoes the principles of hybrid human-AI collaboration, where human domain knowledge is embedded as both constraint and context for machine reasoning [25,26,28].

Nonetheless, several limitations of the system warrant further exploration. First, while the knowledge base construction relied on a broad range of classical and contemporary TCM sources, the completeness and granularity of extracted relationships are inherently dependent on the quality of source standardization, a known challenge in medical ontology engineering [5,7]. Second, although the Bayesian reasoning model effectively captured syndrome-to-disease probabilities, it assumed conditional independence across features, which may oversimplify certain overlapping syndromes or comorbidities. Future iterations could incorporate graph neural networks to learn contextual interactions more robustly. Third, while expert validation demonstrated high trust levels, actual clinical deployment would require real-time interoperability with hospital information systems (HIS) and compliance with health informatics standards such as HL7 or FHIR, as discussed by Chen et al. and Wang et al. in their analyses of CDSS deployment frameworks [21,23].

Ultimately, this research affirms that precision diagnosis in TCM can be significantly enhanced by hybrid AI systems that align with its theoretical foundations and clinical practices. The proposed architecture not only performs competitively across conventional evaluation metrics, but also offers a transparent and extensible platform for domain-specific reasoning, adaptable to future advancements in multimodal data integration, patient engagement, and personalized health technologies.

7. Conclusion

This study presents a comprehensive, modular framework for constructing a structured Traditional Chinese Medicine (TCM) knowledge base and implementing a multilayer intelligent reasoning system for precision diagnosis and treatment recommendation. The framework integrates classical TCM theory with contemporary AI methodologies, combining ontological modeling, fuzzy semantic similarity, Bayesian inference, and personalization rules to emulate the logical sequence of syndrome differentiation and therapy formulation. Empirical results across a multi-center dataset and expert-reviewed validation demonstrate that the system outperforms traditional rule-based and black-box diagnostic models in accuracy, consistency, and clinical interpretability.

By formalizing TCM knowledge into a machine-readable graph structure and embedding domain-specific inference mechanisms, the system enables scalable, explainable, and patient-centered diagnostic support. Its modular design allows transparency at each reasoning stage and supports adaptation across a range of clinical conditions and institutional contexts. Moreover, the high concordance between system-generated outputs and expert assessments underscores the system's practical potential for integration into modern clinical workflows.

Looking ahead, future work will focus on expanding the knowledge base to include multimodal diagnostic cues (e.g., tongue images, pulse waveforms), enhancing the reasoning engine with graph-based deep learning models, and validating the system through longitudinal clinical trials. The convergence of semantic AI and traditional medicine, as demonstrated in this research, offers a promising path toward developing culturally competent, intelligent health systems capable of delivering precision care at scale.

References

1. C.-L. Yao, et al., "Traditional Chinese medicine (TCM) as a source of new anticancer drugs," *Nat. Prod. Rep.*, vol. 38, no. 9, pp. 1618–1633, 2021, doi: 10.1039/D0NP00057D.
2. E. Zhang, et al., "A minireview of the medicinal and edible insects from the traditional Chinese medicine (TCM)," *Front. Pharmacol.*, vol. 14, p. 1125600, 2023, doi: 10.3389/fphar.2023.1125600.
3. R. W. McGee, "Incorporating Artificial Intelligence and Traditional Chinese Medicine (TCM) into a Western Medical Practice: A Case Study," *Biomed. J. Sci. & Tech. Res.*, vol. 56, no. 3, pp. 48149–48154, 2024, doi: 10.26717/BJSTR.2024.56.008864.

4. L. Peng, et al., "Real-World Evidence of Traditional Chinese Medicine (TCM) Treatment on Cancer: A Literature-Based Review," *Evid.-Based Complement. Altern. Med.*, vol. 2022, Art. no. 7770380, 2022, doi: 10.1155/2022/7770380.
5. L. Zhong, et al., "A comprehensive survey on automatic knowledge graph construction," *ACM Comput. Surv.*, vol. 56, no. 4, pp. 1–62, 2023, doi: 10.1145/3618295.
6. T. Shen, F. Zhang, and J. Cheng, "A comprehensive overview of knowledge graph completion," *Knowl.-Based Syst.*, vol. 255, p. 109597, 2022, doi: 10.1016/j.knosys.2022.109597.
7. X. Hao, et al., "Construction and application of a knowledge graph," *Remote Sens.*, vol. 13, no. 13, p. 2511, 2021, doi: 10.3390/rs13132511.
8. X. Zeng, et al., "Toward better drug discovery with knowledge graph," *Curr. Opin. Struct. Biol.*, vol. 72, pp. 114–126, 2022, doi: 10.1016/j.sbi.2021.09.003.
9. Q. Zhang, et al., "Intelligence inference on IoT devices," in *Learning Techniques for the Internet of Things*, 2023, pp. 171–195, doi: 10.1007/978-3-031-50514-0_9.
10. K. Friston, et al., "World model learning and inference," *Neural Netw.*, vol. 144, pp. 573–590, 2021, doi: 10.1016/j.neunet.2021.09.011.
11. M. Amini, et al., "An intelligent traffic congestion detection approach based on fuzzy inference system," in *Proc. 2021 IEEE 15th Int. Symp. Appl. Comput. Intell. Informatics (SACI)*, 2021, pp. 171–195, doi: 10.1109/SACI51354.2021.9465637.
12. R. Tabbussum and A. Q. Dar, "Performance evaluation of artificial intelligence paradigms—artificial neural networks, fuzzy logic, and adaptive neuro-fuzzy inference system for flood prediction," *Environ. Sci. Pollut. Res.*, vol. 28, no. 20, pp. 25265–25282, 2021, doi: 10.1007/s11356-021-12410-1.
13. J. Shao, et al., "Novel tools for early diagnosis and precision treatment based on artificial intelligence," *Chin. Med. J. Pulm. Crit. Care Med.*, vol. 1, no. 03, pp. 148–160, 2023, doi: 10.1016/j.pccm.2023.05.001.
14. H. S. Smith, et al., "Research participants' perspectives on precision diagnostics for Alzheimer's disease," *J. Alzheimers Dis.*, vol. 97, no. 3, pp. 1261–1274, 2024, doi: 10.3233/JAD-230609.
15. Z.-H. Chen, et al., "Artificial intelligence for assisting cancer diagnosis and treatment in the era of precision medicine," *Cancer Commun.*, vol. 41, no. 11, pp. 1100–1115, 2021, doi: 10.1002/cac2.12215.
16. A. Sulewska, et al., "A signature of 14 long non-coding RNAs (lncRNAs) as a step towards precision diagnosis for NSCLC," *Cancers*, vol. 14, no. 2, p. 439, 2022, doi: 10.3390/cancers14020439.
17. W. Zheng, et al., "A deep fusion matching network semantic reasoning model," *Appl. Sci.*, vol. 12, no. 7, p. 3416, 2022, doi: 10.3390/app12073416.
18. K. Li, et al., "Image-text embedding learning via visual and textual semantic reasoning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 1, pp. 641–656, 2022, doi: 10.1109/TPAMI.2022.3148470.
19. L. Ding, et al., "Bi-temporal semantic reasoning for the semantic change detection in HR remote sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, pp. 1–14, 2022, doi: 10.1109/TGRS.2022.3154390.
20. W. Zheng, et al., "Improving visual reasoning through semantic representation," *IEEE Access*, vol. 9, pp. 91476–91486, 2021, doi: 10.1109/ACCESS.2021.3074937.
21. Z. Chen, et al., "Harnessing the power of clinical decision support systems: challenges and opportunities," *Open Heart*, vol. 10, no. 2, p. e002432, 2023, doi: 10.1136/openhrt-2023-002432.
22. A. M. Antoniadi, et al., "Current challenges and future opportunities for XAI in machine learning-based clinical decision support systems: a systematic review," *Appl. Sci.*, vol. 11, no. 11, p. 5088, 2021, doi: 10.3390/app11115088.
23. L. Wang, et al., "Artificial intelligence in clinical decision support systems for oncology," *Int. J. Med. Sci.*, vol. 20, no. 1, pp. 79, 2023, doi: 10.7150/ijms.77205.
24. S. Liu, et al., "Using AI-generated suggestions from ChatGPT to optimize clinical decision support," *J. Am. Med. Inform. Assoc.*, vol. 30, no. 7, pp. 1237–1245, 2023, doi: 10.1093/jamia/ocad072.
25. D. Dellermann, et al., "The future of human-AI collaboration: a taxonomy of design knowledge for hybrid intelligence systems," *arXiv preprint arXiv:2105.03354*, 2021.
26. I. Molenaar, "Towards hybrid human-AI learning technologies," *Eur. J. Educ.*, vol. 57, no. 4, pp. 632–645, 2022, doi: 10.1111/ejed.12527.
27. A. Al-Othman, et al., "Artificial intelligence and numerical models in hybrid renewable energy systems with fuel cells: Advances and prospects," *Energy Convers. Manage.*, vol. 253, p. 115154, 2022, doi: 10.1016/j.enconman.2021.115154.
28. M. M. M. Peeters, et al., "Hybrid collective intelligence in a human–AI society," *AI Soc.*, vol. 36, pp. 217–238, 2021, doi: 10.1007/s00146-020-01005-y.
29. F. Sugandh, et al., "Advances in the management of diabetes mellitus: a focus on personalized medicine," *Cureus*, vol. 15, no. 8, 2023, doi: 10.7759/cureus.43697.
30. M. A. Alghamdi, et al., "The promise of nanotechnology in personalized medicine," *J. Pers. Med.*, vol. 12, no. 5, p. 673, 2022, doi: 10.3390/jpm12050673.
31. V. Gambardella, et al., "Personalized medicine: recent progress in cancer therapy," *Cancers*, vol. 12, no. 4, p. 1009, 2020, doi: 10.3390/cancers12041009.

32. C. Delpierre and T. Lefèvre, "Precision and personalized medicine: What their current definition says and silences about the model of health they promote. Implication for the development of personalized health," *Front. Sociol.*, vol. 8, p. 1112159, 2023, doi: 10.3389/fsoc.2023.1112159.