Research Article

Knowledge Graph and Large Language Model-Based Clinical Decision Support for Traditional Chinese Medicine External Therapy

Ryan Cooper [1,*]
[1] Department of Computer Science, Autonomous University of Madrid, Spain
* Corresponding author: Ryan Cooper
Email: RyanCooper@outlook.com

***Abstract*:** Traditional Chinese medicine (TCM) external therapy involves heterogeneous clinical knowledge concerning symptoms, treatment indications, intervention methods, practitioner experience, and individualized treatment adjustments, which can be difficult to organize within conventional clinical information systems. This study proposes a knowledge graph and large language model (LLM)-based framework for clinical decision support in TCM external therapy. The framework integrates structured treatment knowledge, clinical observations, symptom–intervention relationships, and selected unstructured clinical narratives into a knowledge-oriented representation, while using an LLM to facilitate contextual retrieval, clinical reasoning, and natural-language interaction. Rather than assuming that LLM-generated recommendations are inherently reliable, the framework emphasizes knowledge grounding, traceable evidence, and clinician-centered interpretation. Particular attention is given to challenges involving semantic ambiguity in TCM terminology, incomplete clinical records, heterogeneous treatment descriptions, knowledge conflicts, and the potential generation of unsupported recommendations by language models. The proposed approach may improve the organization and accessibility of TCM external-therapy knowledge and provide more interpretable decision-support information for individualized intervention. Nevertheless, clinical reliability, generalizability, and safety require further evaluation through expert assessment and prospective clinical validation.

Keywords: *Clinical Decision Support, External Therapy, Traditional Chinese Medicine, Knowledge Graph, Large Language Model, Personalized Medicine.*

1. Introduction

1.1 Background and Research Motivation

Traditional Chinese Medicine (TCM) external therapy encompasses a broad range of non-pharmacological interventions, including acupuncture, moxibustion, Tuina, acupoint-based interventions, and other therapeutic approaches. In clinical practice, the selection and modification of these interventions often depend on symptom characteristics, treatment principles, practitioner experience, temporal considerations, and individualized responses. Such characteristics provide TCM external

therapy with considerable flexibility, but they also create challenges for systematic knowledge organization and computational representation. Research on time-based moxibustion, for example, demonstrates that treatment concepts may involve temporal patterns, meridian-related interpretations, and intervention rules that are difficult to represent through conventional clinical information structures [1].

Standardization has consequently become an important prerequisite for the digital transformation of TCM external therapy. Existing research has examined the development of standardized external-therapy approaches for chronic soft tissue pain and has demonstrated the practical importance of establishing relatively consistent treatment descriptions and intervention procedures [2]. Other work has further explored the construction and application of systematic standards for traditional external therapy within TCM health preservation, indicating that standardized terminology and intervention structures may facilitate the integration of traditional therapeutic knowledge with modern information systems [3]. These developments suggest that the digitalization of TCM external therapy should not simply focus on storing clinical records, but should also address how heterogeneous therapeutic knowledge can be organized into structured and computationally accessible representations.

The emergence of artificial intelligence (AI), particularly knowledge graphs and large language models (LLMs), provides a potential technical pathway for addressing this problem. A smart wellness model integrating TCM industry standards with digital systems illustrates how standardized traditional knowledge can be incorporated into broader digital health-management processes [4]. More specifically, recent research has explored the construction and application of TCM knowledge graphs based on large language models, providing a relevant methodological foundation for connecting structured knowledge representation with language-based information processing. This direction is particularly applicable to external therapy because clinical knowledge often involves relationships among symptoms, treatment methods, body regions, acupoints, therapeutic principles, and treatment responses.

1.2 Research Problem and Objectives

Despite these developments, several unresolved problems remain. First, TCM external-therapy knowledge is heterogeneous. The same intervention may be described using different terminology, while similar symptoms may lead to different treatment strategies depending on the patient's condition and the practitioner's interpretation. Second, important clinical information is frequently embedded in narrative records rather than standardized databases. Third, conventional keyword-based information retrieval is often insufficient for representing contextual relationships among symptoms, treatment indications, contraindications, and intervention adjustments.

LLMs provide strong natural-language processing capabilities and may facilitate the interpretation of unstructured clinical narratives. However, their ability to generate fluent responses does not necessarily imply that the generated information is clinically appropriate. A recent survey of LLM applications in TCM emphasizes both the emerging opportunities and the methodological challenges associated with applying language models to TCM knowledge and clinical tasks [9]. Accordingly, an appropriate clinical decision-support framework should not treat an LLM as an autonomous source of medical knowledge. Instead, the model should operate within a knowledge-grounded environment in which relevant clinical concepts and relationships can be retrieved, examined, and traced.

The present study therefore proposes a knowledge graph and LLM-based framework for clinical decision support in TCM external therapy. The principal objective is to establish an architecture that connects structured TCM knowledge with natural-language clinical interaction. The framework is designed to

support knowledge retrieval, contextual interpretation, treatment-option organization, and clinician-centered decision support while retaining explicit attention to uncertainty, evidence traceability, and clinical responsibility.

1.3 Research Significance and Scope

The significance of the proposed framework lies primarily in its attempt to connect three previously separated components: standardized TCM external-therapy knowledge, computational knowledge representation, and natural-language clinical interaction. Rather than replacing practitioner judgment, the framework is intended to improve the accessibility and organization of relevant information.

The clinical relevance of such an approach can also be observed in research examining non-pharmacological TCM interventions in specific patient populations. For example, investigation of TCM non-pharmacological therapies for blood pressure control illustrates how external or non-drug interventions may involve multiple treatment characteristics and patient-specific considerations [6]. Similarly, research concerning the use of the TCM concept of "Harmonization of Body and Mind" for cancer-related anxiety demonstrates the importance of representing therapeutic concepts together with clinical symptoms and patient conditions [7]. These examples indicate that a computational decision-support system needs to accommodate both relatively structured clinical variables and conceptual knowledge that may originate from traditional medical theories.

The scope of this study is consequently limited to the design and analytical discussion of a knowledge-grounded clinical decision-support framework. It does not claim that an LLM can independently diagnose patients or determine treatment plans. Instead, the proposed architecture is intended to provide organized and interpretable information that can be reviewed by qualified clinicians. Clinical safety, effectiveness, generalizability, and prospective validation remain important areas for subsequent research.

2. Related Work and Research Background

2.1 Knowledge Representation and Standardization of TCM External Therapy

The digital representation of TCM external therapy requires a balance between standardization and preservation of clinical complexity. Excessive simplification may remove important contextual information, whereas insufficient standardization can make computational processing difficult. This tension is particularly apparent in external therapies involving acupuncture, Tuina, and moxibustion, where treatment selection may depend on symptom distribution, meridian examination, therapeutic principles, and patient characteristics.

Recent clinical research involving combined acupuncture, Tuina, and moxibustion for chronic multisite musculoskeletal pain provides an example of how multiple therapeutic modalities can be integrated within a structured clinical research setting [8]. Such work is relevant to knowledge representation because the therapeutic process contains relationships among patient symptoms, intervention modalities, treatment targets, and measurable outcomes. A knowledge graph can potentially represent these relationships as interconnected entities rather than treating each clinical description as an isolated text segment.

Other theoretical work on combined acupuncture, Tuina, and moxibustion for elderly chronic musculoskeletal pain further demonstrates the multidimensional nature of TCM external-therapy

reasoning, including the interpretation of meridian examination, tendon excess-deficiency, and qi-blood stasis [10]. From a computational perspective, these concepts should not necessarily be forced into simplistic binary categories. Instead, they may be represented as conceptual entities and relationships whose interpretation remains subject to clinical context.

A knowledge graph can consequently serve as an intermediate representation layer. It may connect symptoms, patient characteristics, treatment modalities, acupoints, therapeutic principles, temporal factors, contraindications, and reported responses. Such a structure can make relationships explicit while allowing the original clinical narratives to remain available for contextual interpretation.

2.2 Artificial Intelligence and Large Language Models in TCM

AI has increasingly been investigated as a means of integrating traditional medical knowledge with modern computational technologies. Current research has considered AI applications across areas including data mining, large language models, multimodal analysis, and clinical decision support. In particular, knowledge-graph construction based on large language models suggests that LLMs can contribute not only to natural-language generation but also to the extraction and organization of relationships from heterogeneous medical information.

Large language models are potentially valuable because clinical knowledge is not exclusively represented in structured databases. Medical records, treatment descriptions, practitioner notes, and literature frequently contain natural-language expressions that are difficult to convert into fixed database fields. An LLM may assist in identifying relevant concepts, normalizing terminology, retrieving contextual information, and generating explanations for clinicians.

However, the application of LLMs to TCM presents distinctive challenges. TCM terminology may contain theoretical concepts that do not have direct one-to-one equivalents in contemporary biomedical terminology. In addition, the same term may acquire different meanings depending on the clinical context. Consequently, language-model processing should be supported by explicit domain knowledge rather than relying solely on statistical language patterns.

The integration of LLMs with knowledge graphs may provide a practical compromise. The knowledge graph supplies structured and traceable information, whereas the LLM provides flexible language understanding and interaction. This division of responsibilities may reduce the risk of allowing generated language to become detached from the underlying knowledge base.

2.3 Clinical Decision Support and TCM External Therapy

Clinical decision-support systems traditionally seek to provide clinicians with relevant information at appropriate points in the decision-making process. In the context of TCM external therapy, such systems may assist with treatment-option retrieval, patient-condition interpretation, intervention comparison, and identification of potentially relevant traditional therapeutic principles.

The potential value of this approach can be illustrated by research on AI-assisted TCM diagnostic and therapeutic models. Recent work on TCM prescription recommendation has explored the combination of large language models and graph neural networks, demonstrating a broader research direction in which graph-based relationships and language-model capabilities are jointly used to support TCM recommendation tasks [11]. Although prescription recommendation and external therapy are different

clinical domains, the methodological principle is relevant: structured relationships can provide an important knowledge foundation for AI-assisted reasoning.

For external therapy, the decision-support process could begin with patient information and clinical observations, followed by retrieval of relevant symptoms, therapeutic concepts, intervention methods, and supporting knowledge. The LLM could then organize the retrieved information into a natural-language explanation. Importantly, the system should distinguish between information directly supported by the knowledge base and interpretations generated by the language model.

Such a distinction is essential because clinical decision support should enhance rather than obscure professional judgment. The final treatment decision should remain under qualified clinical supervision, particularly where evidence is incomplete, terminology is ambiguous, or multiple treatment strategies are possible.

3. Methodology

3.1 Knowledge Graph Construction

The proposed framework begins with the construction of a domain-specific knowledge graph for TCM external therapy. The knowledge graph is designed to represent clinically relevant entities and the relationships among them. Potential entities include symptoms, diseases or health conditions, body regions, acupoints, intervention modalities, treatment principles, patient characteristics, contraindications, treatment responses, and evidence sources.

The construction process may begin with the collection of structured and unstructured sources, including clinical records, standardized treatment descriptions, academic publications, and expert-validated terminology. These sources can subsequently be processed to identify relevant entities and relationships. Terminological normalization is particularly important because different practitioners or publications may use different expressions for similar concepts.

The knowledge graph should also preserve provenance information. Each clinically meaningful relationship should ideally be associated with its source, evidence type, and relevant contextual limitations. This design allows the system to distinguish between established knowledge, expert-derived knowledge, theoretical interpretations, and information requiring further validation.

An important methodological principle is that the knowledge graph should remain extensible. TCM external therapy encompasses a wide range of interventions and theoretical frameworks, making it difficult to construct a complete ontology at the beginning of system development. An iterative approach may be more appropriate, allowing new concepts and relationships to be incorporated as clinical requirements and expert evaluations reveal gaps in the initial representation.

3.2 LLM-Based Knowledge Retrieval and Clinical Reasoning

The second component integrates the knowledge graph with an LLM. When a clinician enters a clinical query or patient description, the system first identifies relevant concepts within the input. These concepts are then used to retrieve related nodes and relationships from the knowledge graph.

The retrieved information provides a structured context for the LLM. Instead of asking the model to generate a treatment recommendation solely from its pretrained knowledge, the system supplies relevant

domain-specific information retrieved from the controlled knowledge environment. The LLM can then transform these structured results into a clinically understandable explanation.

This architecture can support several functions. First, it can summarize relevant patient characteristics and clinical observations. Second, it can identify potentially relevant external-therapy modalities. Third, it can explain the relationship between the patient's condition and the retrieved therapeutic concepts. Fourth, it can identify missing information that may need to be clarified before a clinical decision is made.

The reasoning process should remain transparent wherever possible. When a recommendation or treatment option is generated, the system should identify the underlying knowledge relationships and evidence sources. If the knowledge graph contains insufficient information, the system should communicate uncertainty rather than filling the information gap with unsupported statements.

This is particularly important in TCM because some concepts may be theoretical, context-dependent, or interpreted differently among practitioners. The LLM should therefore be used primarily as a language and reasoning interface over a structured knowledge environment rather than as an unrestricted autonomous clinical authority.

3.3 Clinical Decision-Support Workflow and Validation

The complete workflow can be conceptualized as a sequence of patient-information input, terminology interpretation, knowledge retrieval, contextual synthesis, response generation, evidence presentation, and clinician review. Each stage provides an opportunity for validation and error detection.

During the input stage, relevant clinical information should be separated from ambiguous or incomplete descriptions. During knowledge retrieval, the system should prioritize relationships supported by reliable sources or expert validation. During response generation, the LLM should be constrained to the retrieved knowledge whenever possible. Finally, the clinician should be able to examine the supporting information before deciding whether a suggested intervention is appropriate.

Validation should occur at multiple levels. Technical validation may examine whether the system retrieves relevant knowledge and correctly identifies relationships. Semantic validation may evaluate whether TCM terminology is interpreted consistently. Clinical validation should involve qualified practitioners who assess whether the generated information is relevant, sufficiently supported, and clinically usable.

The evaluation process should also consider failure cases. These may include incorrect terminology mapping, incomplete retrieval, contradictory knowledge, inappropriate generalization from one patient population to another, and unsupported LLM-generated recommendations. Such cases are not merely technical errors; they provide information for improving the knowledge graph, retrieval strategy, and clinical interface through iterative development.

4. Analytical Discussion of the Proposed Framework

4.1 Potential Advantages and Practical Value

The proposed framework may provide several potential advantages for TCM external-therapy decision support. First, the knowledge graph can make complex relationships more explicit. Rather than storing symptoms, treatment methods, and clinical explanations as independent textual fragments, the system can connect them through structured relationships.

Second, the integration of an LLM may improve the accessibility of this structured knowledge. Clinicians do not necessarily need to formulate database-style queries. Natural-language interaction can allow users to describe clinical situations in a more familiar manner, while the underlying retrieval mechanism translates those descriptions into structured knowledge queries.

Third, the framework may facilitate individualized decision support. Standardization does not necessarily require identical treatment for all patients. Instead, standardized terminology and structured knowledge can provide a common representation through which patient-specific differences can be interpreted. This distinction is important for TCM external therapy, where individualized intervention is often an essential component of clinical practice.

The framework may also support broader digital health-management systems. Research on smart TCM has emphasized the integration of traditional medical standards with digital systems across different stages of health management [4]. A knowledge-grounded decision-support architecture could potentially serve as a specialized clinical layer within such broader systems, connecting treatment knowledge with longitudinal patient information and health-management processes.

4.2 Risks, Limitations, and Alternative Interpretations

Despite its potential, the framework has substantial limitations. The first concerns the quality of the knowledge graph itself. An LLM cannot compensate for incomplete, inconsistent, or poorly validated underlying knowledge. If the knowledge graph contains ambiguous relationships, the resulting system may simply transform those ambiguities into more fluent language.

The second limitation concerns the interpretation of TCM concepts. Some traditional concepts cannot be adequately represented through simple biomedical variables. Attempts to force such concepts into overly rigid computational categories could reduce rather than improve clinical fidelity. A layered representation that distinguishes empirical observations, clinical concepts, and theoretical interpretations may be more appropriate.

The third limitation concerns LLM hallucination and overconfidence. A response may appear clinically sophisticated while containing information that is not supported by the retrieved evidence. This creates a particular risk in medical applications because users may interpret fluent explanations as evidence of reliability. Consequently, evidence grounding, provenance tracking, uncertainty expression, and clinician verification should be treated as fundamental components rather than optional features.

There are also alternative interpretations of the proposed architecture. One possibility is that structured rule-based systems may remain preferable for certain narrowly defined clinical tasks because they can provide greater predictability. Another possibility is that LLM-based systems may become increasingly useful as clinical knowledge bases become larger and more standardized. These alternatives suggest that the most appropriate architecture may ultimately be task-dependent rather than universally determined by a single AI technology.

4.3 Future Development and Clinical Implications

Future development could proceed in several directions. First, the knowledge graph could be expanded through expert-reviewed clinical data and standardized terminology. Second, multilingual and cross-disciplinary terminology mapping could be developed to improve communication between TCM concepts and modern biomedical information systems. Third, multimodal information such as imaging,

physiological measurements, and longitudinal clinical observations could potentially be incorporated into the knowledge environment.

Clinical validation will remain particularly important. A technically functional system does not necessarily constitute an effective clinical decision-support tool. Future studies could evaluate whether clinicians can retrieve relevant information more efficiently, whether generated explanations are sufficiently interpretable, and whether the system reduces information omissions without encouraging inappropriate reliance on automated recommendations.

The framework may also benefit from continuous human-in-the-loop learning. Clinicians could identify incorrect relationships, missing concepts, and contextually inappropriate recommendations, allowing the knowledge graph and retrieval mechanisms to be progressively refined. Such a development model would treat AI not as a finished replacement for clinical expertise, but as a continuously evaluated computational infrastructure supporting knowledge organization and clinical communication.

Ultimately, the clinical value of the proposed approach will depend less on the apparent linguistic sophistication of the LLM than on the quality, traceability, and clinical validity of the knowledge infrastructure surrounding it.

5. Transitional Discussion

The proposed knowledge graph and LLM-based framework provides a conceptual bridge between the structured representation of TCM external-therapy knowledge and the flexible language capabilities required for clinical interaction. By combining standardized knowledge organization, contextual retrieval, evidence-aware language generation, and clinician-centered validation, the framework may offer a practical foundation for further development of intelligent TCM decision-support systems. At the same time, the unresolved issues of terminology ambiguity, knowledge incompleteness, evidence heterogeneity, and LLM-generated uncertainty indicate that future work should move beyond conceptual architecture toward systematic technical evaluation and clinical validation. Subsequent research can consequently investigate how the proposed framework performs under real-world clinical scenarios, how knowledge-grounding strategies influence the reliability of generated responses, and how continuous expert feedback can improve the safety and usefulness of AI-assisted decision support.

Data Availability Statement
Data will be made available on request.

Funding
This work was supported without any funding.

Conflicts of Interest
The author(s) declare no conflicts of interest.

Ethical Approval and Consent to Participate
Not applicable.

References

[1] *Rui, H. (2026). Theoretical Analysis on Standardization of Time-Based Moxibustion with Linggui Bafa for Intervention in Mild Cognitive Impairment: From the Perspectives of Time-Based Qi Transformation and Meridian Time Sequence. Insights in Social Science, 4(2), 102-113.*

[2] *Chen, Y. (2026). Standardized Traditional Chinese Medicine External Therapy for Chronic Soft Tissue Pain: A Multicenter Observational Study. Current Research in Medical Sciences, 5(3), 36-45.*

[3] *Chen, Y. (2026). Construction and Application of Standard System for Traditional External Therapy in Traditional Chinese Medicine Health Preservation. Journal of Innovations in Medical Research, 5(2), 23-31.*

[4] *Yao, C. (2026). Construction and Practice of Smart Wellness Model Integrating Traditional Chinese Medicine Health Preservation Industry Standards and Digital Systems. Insights in Social Science, 4(2), 76-86.*

[5] *Zhang, B., Li, R., Yin, K., Hua, S., Li, S., Jiang, M., ... & Li, P. (2026). Construction and application of traditional Chinese medicine knowledge graph based on large language model. Interdisciplinary Sciences: Computational Life Sciences, 18(2), 734-746.*

[6] *HuangYuan, D. (2025). The Impact of Traditional Chinese Non-Pharmacological Therapies on Blood Pressure Control in Community Patients with Hypertension. Current Research in Medical Sciences, 4(5), 8-15.*

[7] *Yuan, D. (2025). Research on the Intervention of Cancer Patients' Anxiety by the Concept of "Harmonization of Body and Mind" in Traditional Chinese Medicine. Journal of Innovations in Medical Research, 4(5), 19-25.*

[8] *Rui, H. (2026). Theoretical Analysis on Standardization of Time-Based Moxibustion with Linggui Bafa for Intervention in Mild Cognitive Impairment: From the Perspectives of Time-Based Qi Transformation and Meridian Time Sequence. Insights in Social Science, 4(2), 102-113.*

[9] *Zhang, R., Tian, D., & Wang, Y. (2025). Large language models in traditional chinese medicine: A short survey and outlook. AI Medicine, 2(1), 3.*

[10] *Huang, R. (2026). Theoretical Analysis of Combined Acupuncture, Tuina and Moxibustion Intervention for Elderly Chronic Musculoskeletal Pain Guided by Meridian Examination—From the Perspective of Tendon Excess-Deficiency and Qi-Blood Stasis. Journal of Research in Social Science and Humanities, 5(2), 32-41.*

[11] *Qi, J., Wang, X., & Yang, T. (2023, December). Traditional Chinese medicine prescription recommendation model based on large language models and graph neural networks. In 2023 IEEE International Conference on Bioinformatics and Biomedicine (BIBM) (pp. 4623-4627). IEEE.*