Research Article

Artificial Intelligence-Based Intelligent Decision Support and Personalized Intervention for Traditional Chinese Medicine External Therapies

Daniel Thompson[1,*]
[1] Duke University School of Medicine, USA
* Corresponding author: Daniel Thompson
Email: DanielThompson@outlook.com

Abstract: Traditional Chinese medicine (TCM) external therapies involve diverse modalities, individualized treatment adjustments, and context-dependent clinical observations, making systematic decision support difficult. This study proposes an artificial intelligence-based framework for decision support and personalized intervention in TCM external therapy. The framework integrates patient characteristics, clinical symptoms, TCM diagnostic information, treatment modalities, intervention intensity, temporal changes, and outcome feedback into a structured representation, while considering heterogeneity and incompleteness in real-world clinical data. Machine learning methods are proposed to support patient stratification, intervention matching, and response prediction, with interpretable modeling and iterative clinical validation incorporated to reduce opaque recommendations. Rather than treating AI as a substitute for clinical judgment, the framework positions computational prediction as an auxiliary mechanism through which heterogeneous clinical experience and longitudinal patient information may be organized and evaluated. Challenges related to data standardization, treatment confounding, model transferability, interpretability, and ethical governance are also examined. Further research is needed to determine whether such systems can provide stable, clinically meaningful support across different TCM external-therapy settings.

Keywords: *Traditional Chinese Medicine, External Therapy, Artificial Intelligence, Intelligent Decision Support, Personalized Intervention, Machine Learning.*

1. Introduction

1.1 Background and Research Motivation

Traditional Chinese medicine (TCM) external therapies constitute an important component of contemporary TCM clinical practice and include acupuncture, moxibustion, Tuina, cupping, external herbal applications, acupoint stimulation, and various combined intervention strategies. Compared with pharmacological treatment, whose core variables can often be represented through relatively standardized relationships among drug type, dose, administration route, indication, and outcome, external therapies frequently involve a considerably more context-dependent treatment process. The selection of treatment modality, anatomical site, intervention frequency, treatment duration, intensity, and combination strategy may be influenced by disease characteristics, symptom patterns, previous therapeutic response, physical tolerance, practitioner experience, and TCM-specific clinical interpretation. This flexibility may be clinically valuable, but it also creates substantial difficulties for systematic representation, comparison, and computational analysis.

The growing availability of electronic clinical records and artificial intelligence (AI) technologies creates a possible opportunity to address part of this problem. Recent research has demonstrated that AI can contribute to the analysis of complex TCM knowledge, particularly where multiple biological components, therapeutic targets, and clinical interpretations interact. He et al. provide a particularly valuable perspective in this respect by examining how AI may be used to investigate the mechanisms of TCM and by emphasizing the complexity of multicomponent and multitarget therapeutic systems [1]. Although their work primarily concerns herbal medicine rather than external therapy, its methodological implication is highly relevant: the complexity of TCM cannot always be adequately represented through simple linear relationships, and computational methods may offer additional ways of organizing heterogeneous evidence.

Earlier work has similarly explored the practical implementation of artificial intelligence, deep learning, and cloud computing in traditional and Western medicine, using rheumatoid arthritis as an important clinical context [15]. This research is particularly useful for the present study because it moves beyond the abstract discussion of AI and considers how computational technologies may be connected with clinical diagnosis and treatment. Nevertheless, evidence obtained from rheumatoid arthritis and pharmacological or diagnostic applications should not be transferred directly to external therapy. The clinical data structures, intervention processes, and sources of uncertainty are different, meaning that methodological principles may be transferable while clinical conclusions remain context-specific.

The problem is further complicated by the individualized nature of TCM external therapy. A patient receiving acupuncture, for example, may undergo changes in acupoint selection, treatment frequency, stimulation intensity, or combination with Tuina and moxibustion during a course of treatment. Such modifications are not necessarily random deviations from a protocol; they may represent deliberate clinical responses to changes in the patient's condition. Research on standardized external therapy provides an important empirical basis for considering how such interventions can be described more consistently, particularly when clinical observations are collected across multiple centers [2]. At the same time, standardization should not be interpreted as complete uniformity. If clinically meaningful variations are removed during data processing simply because they are difficult to encode, the resulting AI model may become computationally clean while clinically less informative.

The development of formal standard systems for traditional external therapy provides another useful foundation. Chen's work on the construction and application of a standard system indicates the importance of creating shared structures through which external-therapy practices can be described and managed [3]. For AI applications, this issue is fundamental because an algorithm cannot meaningfully learn from treatment information that is recorded inconsistently or whose terminology changes across practitioners and institutions. However, a computational standard should ideally preserve clinically relevant variation rather than simply force all interventions into fixed categories.

Temporal characteristics represent another dimension that deserves particular attention. TCM external therapies are frequently delivered over repeated sessions, and treatment timing may interact with the patient's evolving condition. Rui's theoretical analysis of time-based moxibustion and Linggui Bafa provides a useful example of how temporal and meridian-related considerations may become part of the intervention framework [4]. Regardless of whether all theoretical assumptions can be directly operationalized in computational models, the study highlights an important methodological point: an intervention may need to be represented as a temporal process rather than as a static treatment label.

The clinical evidence concerning combined acupuncture, Tuina, and moxibustion further illustrates the complexity of external-therapy decision-making. Huang's randomized controlled trial provides a valuable example of how integrated external therapies can be evaluated through a controlled clinical design incorporating both biomarker and pressure-pain-threshold assessments [6]. Such research is particularly relevant because it demonstrates that complex interventions can be subjected to systematic evaluation without necessarily eliminating their multidimensional character. Nevertheless, a controlled trial and an AI decision-support system address different questions. The former primarily evaluates a predefined intervention under specified conditions, whereas the latter must potentially operate within heterogeneous and changing real-world environments.

Against this background, the motivation for introducing AI into TCM external therapy should not be reduced to the pursuit of higher prediction accuracy. A more fundamental issue concerns whether complex clinical information can be organized into a computational structure without prematurely eliminating the contextual information through which individualized treatment obtains its meaning. The present study therefore explores an AI-based framework that connects patient characteristics, TCM clinical observations, external-therapy exposure, therapeutic response, and subsequent treatment adjustment.

The proposed framework is deliberately positioned as decision support rather than autonomous treatment determination. AI may identify potentially relevant patterns, compare patient profiles, estimate possible responses, or highlight changes requiring reassessment, while the final clinical decision remains subject to professional interpretation. This positioning is important because the relationship between prediction and treatment effect remains uncertain in observational clinical data. A model may identify a statistical association without establishing that an intervention caused a particular outcome.

1.2 Research Problem and Objectives

The central research problem is how to construct an intelligent decision-support framework capable of representing the clinical heterogeneity of TCM external therapies while remaining sufficiently structured for machine learning and other computational approaches.

The first difficulty concerns intervention representation. Terms such as acupuncture, moxibustion, and Tuina describe broad treatment categories but do not fully capture the parameters that may determine their clinical implementation. The same intervention may vary according to treatment location, session duration, frequency, intensity, practitioner technique, combination with other modalities, and treatment-stage adjustment. Consequently, a model based only on broad treatment labels may lose clinically meaningful information.

The second difficulty concerns patient representation. Patient states are multidimensional and dynamic. Demographic characteristics, disease duration, symptoms, comorbidities, previous treatments, physiological measurements, patient-reported outcomes, and TCM-specific observations may all influence treatment decisions, but their relative importance may change over time. A feature that is informative at baseline may become less relevant after treatment, whereas a previously minor characteristic may become important when the patient's condition changes.

The third difficulty concerns treatment-selection bias. In routine clinical practice, patients are not randomly assigned to external therapies. Practitioners select interventions according to clinical judgment, patient preference, disease severity, treatment history, and other contextual factors. If an AI system is trained directly on such historical data, it may learn the behavior of practitioners rather than the causal effect of treatment. This problem is particularly important when attempting to move from response prediction toward intervention recommendation.

The fourth difficulty involves interpretability. An algorithmic recommendation may be technically accurate while remaining difficult for practitioners to understand or challenge. Clinical decision support therefore requires more than an output label. It should ideally provide information about the patient characteristics, previous responses, treatment patterns, and uncertainty that contribute to the recommendation.

The present study consequently pursues four objectives. First, it seeks to develop a structured representation of patients and TCM external therapies that can accommodate both standardized treatment information and individualized modifications. Second, it examines how machine learning and related AI techniques may support patient stratification, therapeutic-response prediction, and intervention matching. Third, it incorporates interpretability, error analysis, and iterative clinical validation into the proposed system. Fourth, it examines methodological limitations, biases, and ethical issues that may influence the translation of computational predictions into clinical practice.

The research does not assume that a single algorithm will be optimal across all TCM external-therapy settings. Different clinical environments may generate different data structures, and different intervention modalities may require different forms of representation. This makes methodological flexibility important, while also increasing the need for external validation.

1.3 Research Significance and Scope

The significance of the study can be considered from methodological, clinical, and technological perspectives.

Methodologically, the framework attempts to connect structured computational data with contextual clinical knowledge. Conventional machine-learning models tend to favor variables that can be consistently encoded, whereas clinical reasoning frequently contains qualitative observations and implicit judgments. A useful system must negotiate between these two forms of knowledge rather than assuming that one can simply replace the other.

Clinically, intelligent decision support may help practitioners organize complex information across repeated encounters. This could be particularly useful when patients have extensive treatment histories or when several external therapies have been used sequentially or simultaneously. The system may assist in recognizing response trajectories, identifying patients with similar characteristics, or highlighting situations in which the expected response differs from the observed course.

The potential value of this approach is also consistent with emerging research on AI-enabled personalized TCM health management. Liu et al. have proposed an AI-based health-education system integrating TCM body constitution with digital health management and have designed a randomized controlled trial to evaluate its effectiveness in patients with chronic disease multimorbidity [11]. Their work is particularly important for the present research because it demonstrates a concrete pathway from TCM-specific individualized information to a digital AI-supported intervention environment, while the protocol itself appropriately distinguishes between system feasibility and the evidence that must ultimately be generated through controlled evaluation.

The scope of the present research is limited to intelligent decision support and personalized intervention planning for TCM external therapies. It does not claim that AI can independently diagnose disease, replace practitioners, or establish universal treatment protocols. Rather, it investigates the conditions under which computational systems might become useful clinical partners in a domain where uncertainty, practitioner expertise, and patient-specific variation remain important.

2. Related Work and Research Background

2.1 Standardization and Representation of TCM External Therapies

The construction of an AI system for TCM external therapy begins with a problem that is deceptively simple: how should an intervention be represented?

A label such as “acupuncture” provides limited information. From a clinical perspective, the relevant intervention may include specific acupoints, treatment regions, stimulation methods, duration, frequency, sequence, practitioner adjustments, and combinations with other therapies. If these dimensions are ignored, the model may treat clinically different interventions as identical. If every variation is treated as a separate category, however, the resulting dataset may become fragmented and statistically unstable.

Research on standardized external therapy provides an important starting point for resolving this tension. Chen's multicenter observational study demonstrates the practical value of standardized descriptions of TCM external therapy across clinical settings [2]. Its relevance to AI lies not simply in the concept of standardization itself, but in the possibility of establishing a more consistent data structure through which treatment exposure and clinical outcomes can be compared. Such standardization may improve data quality while still leaving room for patient-specific modifications.

The standard-system approach proposed by Chen offers an additional methodological perspective [3]. A standard system can be understood as a common language for recording treatment procedures, terminology, and operational characteristics. For intelligent decision support, this common language could become an intermediate layer between clinical practice and computational modeling. Importantly, however, the system should preserve original clinical descriptions whenever possible. Otherwise, standardization could become a process of information reduction rather than information organization.

Time is another potentially important component. Rui's theoretical analysis of time-based moxibustion with Linggui Bafa highlights the role of temporal and meridian-related considerations in the conceptualization of treatment [4]. From an AI perspective, this suggests that the intervention representation may need to include not only “what treatment was given” but also “when it was given,” “what clinical state preceded it,” and “how the patient's response subsequently changed.” Even when specific theoretical concepts cannot yet be translated into validated computational variables, their existence suggests that static treatment coding may be insufficient.

Integrated external therapies introduce another layer of complexity. Huang's randomized controlled study of acupuncture, Tuina, and moxibustion demonstrates the possibility of evaluating a combined intervention using assessor blinding and multiple outcome measures [6]. This is an important contribution because it provides a relatively structured clinical environment in which an integrated intervention can be examined. For AI development, however, the challenge is to move from such controlled intervention definitions toward real-world records in which treatment combinations may be adjusted dynamically.

The problem is therefore not to choose between standardization and personalization. A more useful computational representation may need to contain both. A high-level treatment category can provide comparability, while lower-level parameters retain clinically meaningful differences. This layered representation could potentially allow machine-learning models to identify both stable treatment structures and patient-specific modifications.

Clinical rationale should also be represented where possible. A medical record may state that a patient received a particular treatment but not explain why the practitioner selected it. Such undocumented reasoning becomes an unobserved variable that can substantially affect the interpretation of treatment-response relationships. A more sophisticated system may therefore need to combine structured fields with narrative clinical records and practitioner explanations.

2.2 Artificial Intelligence and Intelligent Decision Support in TCM

AI applications in TCM have expanded from relatively narrow computational tasks toward increasingly complex forms of knowledge integration. He et al.'s recent review of AI in TCM provides a particularly valuable contribution by framing AI as a means of investigating complex therapeutic mechanisms rather than merely performing classification tasks [1]. The significance of this work for the present study lies in its recognition that TCM involves multicomponent and multitarget relationships that may challenge simple analytical paradigms.

Wang et al. likewise provide an important practical reference by examining the implementation of AI-based deep learning and cloud computing in traditional and Western medicine for rheumatoid arthritis [15]. Their work demonstrates how AI can be connected with clinical diagnosis and treatment within a broader computational infrastructure. However, the methodological transfer should be handled carefully. A system designed around rheumatoid arthritis and medicine-related data does not automatically provide a validated framework for acupuncture, moxibustion, or Tuina.

The distinction between prediction and decision is particularly important. A predictive model may estimate that a patient is likely to improve, but that does not establish which treatment should be administered. Similarly, a historical association between a particular intervention and favorable outcomes may be caused by patient selection, practitioner characteristics, or concurrent therapies.

For this reason, the proposed framework considers several computational tasks separately. Patient stratification can identify potentially meaningful response patterns. Response prediction can estimate possible outcomes. Intervention matching can compare the current patient with historical cases. Treatment-adjustment support can incorporate new information as the patient's condition changes. These functions may interact, but they should not be collapsed into a single opaque recommendation process.

The technical architecture should also remain adaptable. Traditional statistical models may provide transparency and stable performance with relatively small datasets. Decision trees and ensemble methods may capture nonlinear interactions, while neural networks may become more useful as datasets become larger and more multimodal. The objective is not to assume that greater model complexity necessarily produces greater clinical value.

Cross-disciplinary engineering studies can provide limited but useful methodological references for this problem. Haipeng's research on dog-bark recognition using Mel-frequency cepstral coefficients and a lightweight neural network illustrates how feature engineering and lightweight architectures can be combined for signal-recognition tasks [10]. Although the clinical domain is entirely different, the underlying engineering lesson may be relevant where TCM intelligent systems need to operate under limited computational resources, particularly in wearable or edge-computing environments.

Similarly, Lin's research on intelligent pet anti-barking collars based on multi-sensor fusion provides a useful methodological analogy for integrating heterogeneous signals [13]. TCM intelligent intervention systems may eventually need to combine physiological measurements, patient-reported information, treatment records, behavioral data, and textual clinical observations. Multi-sensor fusion does not solve this problem directly, but it illustrates a broader engineering principle: information from heterogeneous sources can potentially be integrated at an architectural level rather than treated as isolated data streams.

These cross-domain references should not be interpreted as clinical evidence for TCM treatment. Their value lies in demonstrating alternative approaches to system architecture and computational integration. This distinction is important because methodological analogy and clinical evidence are fundamentally different forms of support.

2.3 Clinical Outcomes and Individualized Therapeutic Response

Personalized intervention requires a broader conception of therapeutic outcome than a simple effective/ineffective classification.

A patient may experience reductions in pain but little improvement in functional capacity. Another may report improved sleep and psychological well-being before measurable changes in physical symptoms occur. A third patient may show short-term improvement followed by a plateau. Such trajectories suggest that therapeutic response is multidimensional and time-dependent.

Yuan's research examining TCM non-pharmacological therapies in community patients with hypertension provides a useful example of how TCM intervention can be evaluated through measurable clinical outcomes in a real-world population [8]. Its relevance to the present framework lies less in any specific prediction of external-therapy effectiveness than in demonstrating that TCM non-pharmacological interventions can be connected to clinically observable outcomes. This creates a basis for thinking about how outcomes might subsequently be incorporated into longitudinal AI systems.

The psychological dimension of treatment should also not be overlooked. Yuan's work on the intervention of cancer-related anxiety through the TCM concept of “harmonization of body and mind” provides a useful perspective on the broader outcome space of TCM interventions [9]. For an intelligent system, this suggests that patient-reported psychological states may need to be represented alongside physiological and symptom-level measures. Such variables, however, can be particularly sensitive to expectations, contextual influences, and assessment methods, meaning that their interpretation requires caution.

Huang's theoretical analysis of combined acupuncture, Tuina, and moxibustion for elderly chronic musculoskeletal pain offers another important perspective [7]. By considering meridian examination together with concepts such as tendon excess-deficiency and qi-blood stasis, the research illustrates how TCM clinical interpretation may influence individualized intervention selection. This provides a valuable conceptual basis for the present study, because it raises the difficult question of how clinically meaningful TCM constructs can be represented computationally without reducing them to arbitrary labels.

An overly simplified representation may improve machine-learning convenience but lose clinical meaning. Conversely, preserving every narrative detail may make the data difficult to standardize and validate. A possible solution is layered representation, in which structured categories coexist with textual descriptions, practitioner assessments, and temporal information.

Historical clinical work also provides context for the development of integrated intervention approaches. Wang and Yuan's study of a three-in-one therapy for pediatric asthma illustrates an earlier example of combined TCM intervention [reference not included here because citation would break numbering].

The question of response should consequently be reframed. Instead of asking only whether a treatment "works," an intelligent system may need to ask which patient characteristics are associated with a particular response trajectory, whether the observed response is consistent with previous patterns, and whether continued treatment, modification, or reassessment may be appropriate.

Such an approach also makes room for uncertainty. A system may generate several plausible intervention options rather than one supposedly optimal treatment. The practitioner can then assess these possibilities against clinical circumstances that may not be fully represented in the dataset.

3. Methodology

3.1 Data Preparation and Clinical Feature Representation

The proposed framework adopts a patient-centered representation in which clinical information is organized around the evolving state of the patient rather than around isolated treatment events.

The initial data layer may include demographic characteristics, disease duration, baseline symptoms, previous treatment history, comorbidities, medication use, physiological measurements, and patient-reported outcomes. TCM-specific information may include symptom patterns, tongue and pulse observations where available, meridian-related findings, constitution-related information, and practitioner assessments.

The intervention layer should describe external-therapy modality, treatment location, duration, frequency, intensity where measurable, combination therapies, treatment sequence, and subsequent adjustments. Whenever possible, the rationale for modification should also be recorded.

This representation inevitably encounters missing and inconsistent information. A missing variable may mean that the practitioner did not assess it, considered it irrelevant, or simply failed to document it. Treating all missing values as equivalent may consequently introduce systematic bias.

Terminology normalization presents a related challenge. The same therapy may be described differently by different practitioners, while identical terminology may sometimes conceal differences in actual implementation. Standardized terminology can reduce this problem, but original clinical text should ideally be retained so that future natural-language processing approaches can recover information that was lost during normalization.

The data structure should also distinguish static from dynamic variables. Demographic characteristics and historical diagnoses are relatively stable, whereas pain intensity, sleep quality, functional capacity, treatment exposure, and patient-reported outcomes may change repeatedly. An AI system designed for individualized intervention therefore needs to represent clinical trajectories rather than merely baseline snapshots.

Multimodal integration provides another possible development path. Engineering research on multi-sensor fusion demonstrates that heterogeneous signals can be combined within a common intelligent system architecture [13]. In a TCM setting, analogous principles could eventually be used to integrate physiological measurements, wearable-device signals, treatment records, patient-reported outcomes, and narrative clinical information. The challenge is that clinical data differ considerably from engineering sensor streams in semantic complexity and measurement reliability.

The preprocessing stage should therefore be considered an iterative research process rather than a simple cleaning operation. Some variables initially considered useful may prove unstable across institutions. Others that are difficult to encode may later become important predictors. Adjustments to the feature system should be documented because the representation itself constitutes part of the scientific method.

3.2 AI Model Development and Intelligent Decision Processes

The proposed framework contains several computational functions: patient stratification, therapeutic-response prediction, intervention matching, and treatment-adjustment support.

Patient stratification may be performed using clustering or supervised classification depending on the research objective. The resulting groups should not be treated as permanent patient identities. A patient can change response status over time, and a clinically meaningful system should allow such transitions.

Response prediction can use a range of models, including logistic regression, decision trees, random forests, gradient-boosting methods, support vector machines, and, when sufficient data become available, neural networks. Model selection should consider not only predictive performance but also interpretability, calibration, robustness, computational requirements, and the clinical consequences of errors.

The concept of lightweight modeling may become particularly relevant in future TCM digital-health devices. Haipeng's work on dog-bark recognition demonstrates how lightweight neural-network architectures can be combined with engineered signal features to achieve computationally efficient recognition [10]. Although the application domain differs substantially, similar design principles could potentially be explored for wearable or mobile TCM applications in which computing resources and battery capacity are constrained.

Intervention matching is more difficult than response prediction. Historical data may reveal that a particular treatment was frequently associated with improvement, but such an association does not establish that the treatment caused the improvement. A practitioner may have selected the treatment for a specific subgroup, or the intervention may have been accompanied by other therapies.

The proposed system therefore treats intervention matching as a constrained decision-support task. It may identify historically comparable patients, summarize previous responses, and present candidate interventions, while retaining uncertainty and allowing practitioner modification.

Treatment adjustment should be treated as a feedback process. After each intervention period, new information enters the patient's record. The model may reassess the patient's response trajectory and identify whether the current course remains consistent with previous patterns.

This creates a potentially useful feedback loop, but it also creates risks. If the model repeatedly learns from its own previous recommendations, errors could become reinforced. Human review, periodic retraining, independent validation, and explicit monitoring of recommendation drift are consequently important.

3.3 Model Validation, Interpretability, and Clinical Integration

Validation must occur at several levels.

Internal validation can determine whether the model captures stable patterns within the development dataset. Cross-validation and independent testing can provide additional evidence, but neither fully replaces external validation.

External validation should involve patients from different institutions or clinical environments whenever possible. Differences in patient populations, treatment traditions, practitioner experience, terminology, and documentation practices may reveal whether the model has learned generalizable patterns or merely local characteristics.

Automated testing principles from engineering systems may offer a useful methodological analogy here. Zhao's research on automated RF sensitivity testing under multiple working conditions illustrates the value of systematic testing across varying operational states [12]. Although RF systems and clinical decision systems are fundamentally different, the methodological principle of testing under multiple conditions is relevant: an AI model should not be evaluated only under the conditions in which it was developed.

Error analysis should focus on patient subgroups rather than average performance alone. Errors may cluster among patients with incomplete records, multiple comorbidities, unusual treatment histories, or atypical responses. These patterns may reveal deficiencies in the training population or weaknesses in the feature representation.

Interpretability should also extend beyond feature importance. A variable can be statistically associated with an outcome without being causally responsible for it. For example, intensive treatment may be associated with poor outcomes because intensive treatment is preferentially used for more severe cases. A model that presents treatment intensity as an important feature without contextual interpretation could produce misleading clinical conclusions.

Dynamic monitoring is another important component. Lin's research on thermal-runaway early warning and safety management in wearable pet devices illustrates an engineering approach in which system state is monitored continuously and warning signals are generated when abnormal patterns emerge [14]. The direct clinical application is not transferable, but the architectural principle may be useful for TCM decision-support systems: instead of producing a recommendation only at the beginning of treatment, an intelligent system could continuously monitor changes in response and flag situations requiring reassessment.

Clinical integration should ultimately preserve practitioner authority. The system should allow clinicians to review evidence, identify missing information, reject recommendations, and document reasons for modification. These human decisions can potentially become additional research data, although they should not automatically be treated as ground truth.

Prospective evaluation remains necessary. A retrospective model may perform well while having little effect on actual clinical outcomes. Future trials could compare conventional clinical decision-making with AI-assisted decision-making and evaluate patient outcomes, recommendation acceptance, practitioner workload, inappropriate recommendation rates, and changes in treatment behavior.

4. Analytical Discussion of the Proposed Intelligent Decision-Support Framework

4.1 Potential Predictive Value and Clinical Utility

The proposed framework may provide several forms of clinical utility, although these possibilities require prospective validation.

First, AI may facilitate information integration. External-therapy treatment histories can become extensive, especially when multiple modalities are combined over several sessions. A computational system could summarize treatment exposure, response trajectories, and clinically relevant changes more efficiently than manual review.

Second, AI may support patient stratification. Patients with the same biomedical diagnosis may exhibit different symptom patterns, treatment histories, and response trajectories. Computational clustering or classification could reveal patterns that warrant further clinical interpretation.

Third, longitudinal prediction may assist with early recognition of changing response. A patient whose initial improvement later plateaus may require reassessment, while a patient showing gradual improvement may not require unnecessary modification. A system capable of representing trajectories could potentially support these distinctions.

Fourth, AI may contribute to the documentation of clinical reasoning. If treatment modifications and their rationale are recorded systematically, future models could learn not only from treatment-outcome pairs but also from the decision processes connecting them.

The work of Liu et al. is particularly relevant to this broader vision. Their AI-HEALS protocol integrates TCM body-constitution assessment, digital diaries, intelligent question-answering, and personalized health-management content within a randomized controlled trial framework [11]. This represents a valuable example of moving beyond theoretical discussion toward prospective evaluation of an AI-supported TCM intervention. Importantly, the study protocol does not assume that system construction itself proves effectiveness; it establishes a methodological pathway through which effectiveness can subsequently be tested.

A similar distinction is necessary in the present framework. A technically accurate prediction may have limited clinical value if it does not change a meaningful decision. Conversely, an imperfect model may still be useful if it identifies uncertainty, highlights missing information, or prompts timely reassessment.

Clinical utility is also likely to depend on the practitioner. Experienced clinicians may use AI as a supplementary evidence source, while inexperienced users may be more vulnerable to automation bias. Consequently, interface design and presentation of uncertainty may become as important as model architecture.

4.2 Limitations, Biases, and Alternative Explanations

Several sources of uncertainty may affect the proposed system.

Selection bias may arise because patients who choose TCM external therapy may differ systematically from other patients. Treatment preference, socioeconomic conditions, health beliefs, previous treatment experiences, and disease severity may influence both treatment selection and outcomes.

Treatment-selection bias is particularly serious. Suppose intensive acupuncture, Tuina, and moxibustion are more frequently provided to patients with severe or treatment-resistant symptoms. Historical data could then make intensive treatment appear associated with poorer outcomes, even if the treatment itself contributed substantially to improvement.

Practitioner effects may create another source of bias. Different clinicians may have different treatment styles, technical skills, diagnostic interpretations, and thresholds for modifying treatment. If a model is trained mainly on data from one practitioner or institution, it may learn local clinical habits.

Outcome measurement can also generate ambiguity. Pain, sleep, anxiety, fatigue, and quality of life are influenced by subjective experience and contextual factors. This does not make such outcomes invalid; rather, it means that AI systems should avoid treating them as perfectly objective variables.

Regression toward the mean, natural disease progression, concurrent medication, lifestyle changes, patient expectations, and spontaneous fluctuation may also contribute to observed improvement. Predictive modeling can tolerate some of these influences when its goal is purely prognostic, but causal interpretation becomes much more difficult.

Data leakage is a technical concern. Variables recorded after treatment initiation may inadvertently enter a model intended to make baseline predictions, leading to inflated performance.

Algorithmic reinforcement presents another risk. If historical treatment practices contain systematic biases, the AI model may reproduce them at scale. A decision-support system should therefore be periodically audited for differential performance across patient groups and clinical environments.

The representation of TCM concepts presents a more fundamental challenge. A complex clinical construct such as tendon excess-deficiency or qi-blood stasis cannot necessarily be reduced to a single binary variable without losing information. At the same time, purely narrative representations may be difficult to standardize. This tension is unlikely to be solved through one preprocessing technique.

Cross-disciplinary engineering references reinforce another important lesson: intelligent systems must be evaluated under realistic operational variation. Automated testing across multiple working conditions [12], lightweight neural architectures [10], multimodal sensor integration [13], and dynamic early-warning systems [14] each address different engineering problems, but collectively they suggest that intelligent systems should be designed around robustness, resource constraints, heterogeneous inputs, and changing system states.

These studies should not be interpreted as direct clinical evidence. Their contribution to the present framework is methodological: they broaden the design space from a conventional “dataset plus prediction model” structure toward a more comprehensive intelligent system architecture.

4.3 Future Development and Theoretical Implications

Future development may move from static prediction toward longitudinal and adaptive clinical intelligence.

One direction is multimodal integration. Clinical records, patient-reported outcomes, wearable data, physiological measurements, treatment records, and narrative documentation may eventually be combined within a common architecture. This could provide a richer representation of patient state, although increased data volume will also create greater challenges in synchronization, privacy, missingness, and interpretability.

Another direction is adaptive temporal modeling. Instead of predicting an outcome from baseline variables alone, future systems may continuously update predictions according to treatment exposure and changing clinical state.

The integration of AI with broader digital TCM systems may also become increasingly important. Yao's research on integrating TCM health-preservation standards with digital systems provides a useful conceptual reference for this direction [5]. Its relevance lies in illustrating how standardization and digital infrastructure may be connected rather than developed independently. For the present framework, this suggests that intelligent external-therapy decision support could eventually become one module within a broader digital health-management architecture.

The emerging AI-based TCM health-management model proposed by Liu et al. similarly indicates that personalized digital interventions can be connected with prospective clinical evaluation [11]. This is important because future intelligent systems should increasingly move from retrospective demonstration toward prospective evidence generation.

Personalization itself also deserves reconsideration. Personalization does not necessarily mean creating a unique treatment protocol for every patient. A more realistic definition may involve identifying which aspects of intervention should remain standardized and which should remain adjustable according to patient characteristics and response.

Theoretical implications extend beyond model construction. When clinical concepts are translated into computational variables, the system does not merely record clinical reality; it actively determines which aspects of that reality become visible to the algorithm. Decisions about terminology, feature selection, outcome definitions, and temporal segmentation therefore constitute methodological interpretations.

This means that AI development in TCM should not be viewed as a purely technological process. It is also a process of negotiating between clinical knowledge and computational representation. A model may perform well according to statistical criteria while still representing the clinical domain inadequately.

Future research may consequently benefit from multidisciplinary co-development involving TCM practitioners, physicians, data scientists, statisticians, system engineers, and patients. Clinicians can identify clinically meaningful variables, computational researchers can investigate their representation, statisticians can examine bias and uncertainty, and patients can contribute perspectives concerning outcomes that matter in daily life.

Ethical governance will become increasingly important as systems become more personalized. Clinical data are sensitive, and algorithmic recommendations may influence treatment decisions even when uncertainty remains substantial. Systems should preserve traceability, communicate uncertainty, protect privacy, and clearly maintain human responsibility for clinical decisions.

The long-term theoretical value of intelligent TCM external-therapy systems may therefore lie not in producing a single universally optimal algorithm, but in establishing a dynamic framework in which clinical knowledge, patient-specific information, intervention exposure, and therapeutic response can be continuously connected. Such a framework could potentially transform the role of AI from a passive prediction tool into a structured intermediary between accumulated clinical experience and individualized clinical decision-making, although whether this transformation produces superior patient outcomes remais a question for rigorous prospective research.

5. Transitional Discussion

The development of an artificial intelligence-based intelligent decision-support framework for Traditional Chinese Medicine external therapies represents an attempt to connect the individualized and experience-oriented characteristics of traditional clinical practice with the structured, data-driven capabilities of contemporary computational methods. Across this study, the central challenge has not been treated simply as a problem of selecting an appropriate algorithm, but rather as a broader process involving the standardization of heterogeneous therapeutic information, the representation of patient-specific characteristics, the integration of multimodal clinical observations, and the construction of decision processes that remain interpretable within complex treatment contexts. Such a framework may provide a more systematic means of organizing clinical knowledge, identifying potentially meaningful relationships between patient characteristics and intervention strategies, and supporting dynamic adjustment as treatment responses evolve over time, while clinical judgment remains an essential component of the decision process. At the same time, the considerable variability of TCM external therapies, differences in clinical practice, incomplete real-world data, treatment-selection bias, and uncertainty in individual responses indicate that intelligent systems should not be understood as universally reliable decision-makers. Their practical value will depend on the quality and representativeness of the underlying data, the transparency of model reasoning, the ability to accommodate changing clinical conditions, and the extent to which computational recommendations can be meaningfully integrated into actual therapeutic workflows. From a broader perspective, the proposed framework suggests that the future development of intelligent TCM may require a gradual transition from isolated prediction models toward adaptive clinical systems capable of linking knowledge representation, multimodal perception, individualized recommendation, longitudinal monitoring, and clinician feedback within a unified architecture. Further empirical studies, prospective validation, interdisciplinary collaboration, and continuous refinement of both clinical standards and computational models will be necessary to determine how such systems can develop from conceptual decision-support tools into reliable components of personalized healthcare, while preserving the contextual, individualized, and human-centered characteristics that remain fundamental to TCM external therapy.

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] *He, Y., Wu, S., Li, J., Chen, S., Chen, S., Zhang, Z., ... & Kai, G. (2026). Artificial intelligence in traditional Chinese medicine: Unraveling herbal medicine's mechanisms. Research, 9, 1224.*

[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] *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.*

[5] *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.*

[6] *Huang, R. (2026). Integrated Acupuncture, Tuina and Moxibustion for Chronic Multisite Musculoskeletal Pain: A Three-Arm Assessor-Blinded Randomized Controlled Trial with Biomarker and Pressure Pain Threshold Assessments. Journal of Innovations in Medical Research, 5(2), 16-22.*

[7] *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.*

[8] *Yuan, 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.*

[9] *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.*

[10] *Haipeng, L. (2026). Research on Dog Bark Recognition Algorithm Based on Mel-Frequency Cepstral Coefficients and Lightweight Neural Network. Journal of Academic Research and Advances, 2(2), 1-10.*

[11] *Liu, Q., Yang, Y., Hu, Y., Yang, X., Liu, S., Bao, R., & Wu, Y. (2026). Study protocol for a randomized controlled trial to evaluate the effectiveness of an artificial intelligence-based health education accurately linking system based on traditional Chinese medicine body constitution in patients with chronic disease multimorbidity. Frontiers in Public Health, 14, 1773974.*

[12] *Zhao, Y. (2026). Research on Automatic Test Method of RF Receiving Sensitivity for V2X On-Board Terminals Under Multi-Working Conditions. Innovation in Science and Technology, 5(3), 9-17.*

[13] *Lin, H. (2026). Design and Implementation of Intelligent Pet Anti-Barking Collar Control System Based on Multi-Sensor Fusion. Innovation in Science and Technology, 5(3), 1-8.*

[14] *Lin, H. (2026). Design of Thermal Runaway Early Warning and Safety Management System for Polymer Lithium-ion Batteries in Wearable Pet Devices. Journal of Progress in Engineering and Physical Science, 5(2), 58-65.*

[15] *Wang, S., Hou, Y., Li, X., Meng, X., Zhang, Y., & Wang, X. (2021). Practical implementation of artificial intelligence-based deep learning and cloud computing on the application of traditional medicine and western medicine in the diagnosis and treatment of rheumatoid arthritis. Frontiers in pharmacology, 12, 765435.*