Research Article

A Machine Learning-Based Model for Predicting Therapeutic Responses to Integrated Traditional Chinese Medicine Interventions

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

Abstract: Integrated Traditional Chinese Medicine (TCM) interventions often involve multiple therapeutic modalities, individualized treatment adjustments, and heterogeneous patient characteristics, making treatment-response prediction difficult using conventional analytical approaches. This study proposes a machine learning-based framework for predicting therapeutic responses to integrated TCM interventions by incorporating patient characteristics, baseline clinical manifestations, treatment-related variables, and longitudinal response indicators. The framework considers the challenges of heterogeneous clinical data, incomplete observations, variable treatment combinations, and potential differences in response patterns among individuals. Several stages of data preprocessing, feature representation, model development, and performance evaluation are considered, with particular attention to the interpretability of predictive outputs and the clinical meaning of selected variables. Rather than assuming that a predictive model can fully capture the complexity of TCM intervention, the proposed approach is intended to provide probabilistic decision-support information that may assist in identifying potentially relevant response patterns and treatment-related factors. Further research is needed to evaluate model robustness, external validity, and clinical usefulness across different populations and intervention settings.

Keywords: *Traditional Chinese Medicine, Machine Learning, Integrated Intervention, Therapeutic Response Prediction, Clinical Decision Support, Personalized Medicine.*

1. Introduction

1.1 Background and Research Motivation

Integrated Traditional Chinese Medicine (TCM) interventions frequently involve the coordinated use of multiple therapeutic modalities, including acupuncture, Tuina, moxibustion, dietary regulation, lifestyle guidance, and other individualized approaches. Unlike interventions in which a single treatment component can be relatively clearly isolated, integrated TCM treatment often changes according to the patient's symptoms, physical condition, treatment response, and practitioner interpretation. This flexibility may be clinically meaningful, yet it also introduces considerable complexity when treatment outcomes are analyzed from a computational perspective. The same nominal intervention may not represent exactly the same therapeutic process across patients, while apparently similar patients may receive different combinations or sequences of interventions.

This issue becomes particularly evident when therapeutic response is considered as a heterogeneous rather than binary phenomenon. Improvement may involve changes in symptom severity, functional status, physiological indicators, quality of life, or other patient-reported outcomes, and these changes may occur at different rates. For example, research examining the effects of TCM non-pharmacological therapies on blood pressure in community patients with hypertension illustrates how clinical outcomes may need to be interpreted within a specific patient population and treatment context rather than being generalized independently of baseline characteristics [1]. Consequently, predicting therapeutic response requires consideration not only of treatment-related variables but also of individual characteristics and the interaction between patients and intervention patterns.

Machine learning provides a possible methodological approach to this problem because it can identify nonlinear associations and higher-order interactions within multidimensional clinical data without requiring every relationship to be specified in advance. However, the availability of machine learning algorithms does not itself resolve the underlying clinical and methodological difficulties. If the input variables are inconsistently recorded, if treatment categories are excessively simplified, or if the outcome definition fails to reflect clinically meaningful change, even a technically sophisticated model may produce predictions of limited practical value.

The motivation for the present study therefore arises from a dual consideration. On one side, there is a growing need to explore computational approaches capable of handling heterogeneous TCM intervention data; on the other, there is a need to avoid reducing individualized TCM practice to a set of mechanically optimized numerical variables. A useful prediction framework should preserve sufficient clinical meaning while remaining computationally tractable, and its predictive output should be interpreted as decision-support information rather than as a definitive statement about an individual patient's future response.

1.2 Research Problem and Objectives

The central research problem addressed in this study concerns how machine learning can be used to model therapeutic responses to integrated TCM interventions when patient characteristics, treatment strategies, and clinical outcomes are heterogeneous. Several related difficulties arise from this problem.

First, clinical datasets may contain variables originating from fundamentally different sources. Demographic characteristics, baseline symptoms, laboratory or physiological measurements, treatment modalities, treatment frequency, and patient-reported outcomes may all coexist within the same dataset, but they do not necessarily possess comparable clinical meanings. Second, treatment exposure is rarely completely uniform. Patients may receive different combinations of interventions or experience adjustments during treatment. Third, missing observations and irregular follow-up may introduce patterns that are related to clinical circumstances rather than occurring randomly. A patient who returns for additional evaluation, for example, may differ systematically from a patient who does not.

The study therefore aims to develop a conceptual machine learning framework capable of integrating heterogeneous patient and treatment information for therapeutic-response prediction. The objectives include identifying clinically meaningful predictor variables, constructing an appropriate representation of integrated TCM interventions, exploring suitable machine learning approaches, and establishing a validation strategy that considers both predictive performance and clinical interpretability.

Another objective is to examine the relationship between prediction and explanation. A model may produce relatively accurate predictions while providing limited insight into why a particular prediction was generated. For clinical decision support, however, a prediction without interpretable context may be difficult to incorporate into professional decision-making. The proposed framework consequently considers feature contribution, model behavior, and clinical plausibility as complementary dimensions of model assessment rather than treating predictive accuracy as the sole criterion.

1.3 Research Significance and Scope

The potential significance of this research lies in creating a methodological connection between individualized TCM intervention and data-driven clinical prediction. Rather than attempting to replace traditional therapeutic reasoning, machine learning may provide an additional analytical layer through which recurring patterns across patients can be explored. Such an approach could, to some extent, help identify patient characteristics associated with different response patterns and provide clinicians with additional information when considering treatment strategies.

The scope of the study is intentionally focused on therapeutic-response prediction rather than autonomous treatment recommendation. This distinction is important because predicting a possible response and deciding which intervention should be administered are not equivalent tasks. A predictive model may indicate that certain patient characteristics are associated with a particular response pattern under an intervention, but such an association does not necessarily establish causality.

The study also does not assume that all aspects of TCM can or should be transformed into conventional machine learning features. Some diagnostic and therapeutic concepts may be context-dependent and difficult to encode without losing clinically meaningful information. The framework therefore allows structured clinical variables to coexist with richer representations where appropriate. Considering these factors, the study is intended as a methodological foundation for further empirical investigation rather than a final claim concerning the universal predictive capability of machine learning in TCM.

2. Related Work and Research Background

2.1 Standardization and Representation of Integrated TCM Interventions

One of the fundamental prerequisites for predictive modeling is the availability of sufficiently consistent data representation. In TCM, however, standardization is complicated by the individualized nature of clinical practice. External therapies may involve different combinations of interventions, treatment frequencies, practitioner-specific adjustments, and patient responses. Research on standardized TCM external therapy for chronic soft tissue pain illustrates the effort required to establish more consistent intervention descriptions while retaining clinical applicability [2]. Such work provides a useful methodological basis for predictive modeling because a machine learning model requires treatment variables to be represented in a form that can be compared across cases.

The development of standardized time-based moxibustion approaches further illustrates that intervention representation may involve temporal and theoretical dimensions that cannot easily be reduced to a single treatment label. Work concerning the standardization of Linggui Bafa and time-based moxibustion demonstrates how treatment interpretation can involve temporal Qi transformation and meridian time sequence, suggesting that temporal context may sometimes be part of the intervention itself rather than merely an external variable [3].

The problem becomes more complex when integrated interventions are considered. Acupuncture, Tuina, and moxibustion may be used independently or in combination, and their therapeutic effects may not be separable in a straightforward manner. Research examining combined acupuncture, Tuina, and moxibustion for chronic multisite musculoskeletal pain provides an example of a clinical design in which multiple intervention components and outcome measures are considered together [7]. Such research can offer useful information for feature construction, although a clinical study designed primarily to assess treatment effects does not automatically provide an optimal structure for machine learning prediction.

Accordingly, the present framework treats intervention representation as a substantive methodological issue rather than a simple preprocessing step. Treatment modality, intensity, frequency, duration, sequence, and adjustment may need to be represented separately where the clinical data permit, while overly detailed encoding should be avoided when the available sample size cannot support reliable estimation.

2.2 Machine Learning and Predictive Modeling in TCM

The application of artificial intelligence to TCM has expanded from relatively structured diagnostic tasks toward more complex computational modeling, prediction, and personalized clinical applications. Deep learning and machine intelligence have been explored for discovering combination rules and pharmacodynamic characteristics in TCM, illustrating how computational methods may be used to examine complex relationships among multiple therapeutic components rather than focusing exclusively on isolated variables [4]. Although pharmacodynamic modeling and clinical therapeutic-response prediction represent different analytical problems, the methodological implication is relevant: complex TCM interventions may contain interaction structures that conventional single-variable analysis cannot adequately represent.

Recent research on AI-driven TCM has also considered a broader continuum extending from systems-biology-based mechanism discovery to real-world clinical evidence inference and personalized clinical decision support. This perspective is relevant to the present study because it places prediction within a wider clinical information ecosystem rather than treating machine learning as an isolated algorithmic exercise [11]. At the same time, such a broad AI framework also highlights a methodological difficulty: models developed for different purposes may operate at different levels of abstraction, and a mechanism-discovery model should not automatically be interpreted as a patient-level outcome predictor.

Another challenge is that TCM datasets are often smaller and less standardized than the large-scale datasets commonly associated with modern machine learning research. This creates a risk that a model may learn characteristics specific to a particular institution, practitioner group, or patient population rather than patterns that generalize to broader clinical settings. In addition, variables that appear statistically informative may partly reflect documentation habits or treatment-selection preferences.

The use of machine learning should consequently involve a distinction between predictive association and clinical explanation. A feature may contribute strongly to model prediction without representing a clinically causal factor. For example, treatment frequency may correlate with outcome not because frequency itself produces improvement, but because patients with certain clinical characteristics are more likely to receive frequent treatment. If such confounding relationships are not considered, a predictive model could inadvertently reproduce existing clinical selection patterns.

2.3 Clinical Outcomes and Individualized Response

A prediction model is only as meaningful as the outcome it attempts to predict. In TCM intervention research, outcomes may include symptom relief, functional recovery, physiological changes, quality of life, or composite measures. These outcomes are not necessarily interchangeable, and improvement in one dimension may occur without equivalent improvement in another.

For this reason, the present framework emphasizes clinically interpretable outcome definitions. A model predicting a broad category such as “effective” or “ineffective” may be easier to construct, but such a classification could obscure intermediate responses and temporal differences. In contrast, a more detailed response representation may provide greater clinical information but require larger and more consistently measured datasets.

The relevance of multidimensional outcomes can also be seen in research concerning the use of TCM concepts for psychological symptoms. Investigation of the “Harmonization of Body and Mind” concept in relation to cancer-patient anxiety illustrates how therapeutic outcomes may extend beyond a single physiological indicator and involve psychological and patient-centered dimensions [9]. Such findings suggest that predictive modeling should define the target outcome according to the clinical question rather than simply selecting whichever variable is most consistently available.

Individualized response further requires attention to baseline heterogeneity. Two patients may receive similar treatment but experience different outcomes because of differences in age, disease duration, symptom severity, comorbidities, lifestyle, previous treatment, or other characteristics. Machine learning may be useful in identifying such patterns, although further research is needed to distinguish genuine treatment-response heterogeneity from dataset-specific associations.

3. Methodology

3.1 Data Preparation and Clinical Feature Representation

The proposed methodology begins with the organization of patient-level clinical data into a structured analytical representation. The intention is not to transform every available clinical statement into a numerical feature, but to identify variables that have sufficient clinical relevance and recording consistency to support predictive modeling.

Potential input variables include demographic characteristics, baseline symptoms, disease duration, previous treatment history, physiological measurements, patient-reported outcomes, TCM-related clinical observations, treatment modalities, treatment frequency, treatment duration, and documented intervention adjustments. Where multiple treatment modalities are used simultaneously, the representation should preserve their combination rather than treating the entire intervention as an undifferentiated category.

Data preprocessing presents one of the more difficult stages of the proposed methodology. Missing values, inconsistent terminology, irregular observation times, and differences in documentation style may initially appear to be technical inconveniences, but they can also contain clinically meaningful information. A decision to remove incomplete records, for example, may unintentionally exclude patients with more complex clinical courses. Similarly, excessive imputation may create artificial regularity that was not present in the original clinical data.

The proposed framework therefore favors a clinically informed preprocessing strategy. Missingness should first be examined in relation to clinical circumstances, while categorical variables should be standardized according to their actual semantic meaning. Continuous measurements may require normalization where appropriate, but normalization should not eliminate clinically meaningful differences in scale. Feature selection should likewise combine statistical considerations with domain knowledge, particularly when variables are strongly correlated or represent overlapping clinical concepts.

3.2 Machine Learning Model Development

Following data preparation, machine learning models can be developed to estimate potential therapeutic responses from baseline and treatment-related information. Candidate approaches may include logistic regression, decision-tree-based methods, random forests, gradient-boosting models, support vector machines, and other appropriate supervised learning algorithms. The purpose of comparing multiple model families is not simply to identify the model with the highest numerical performance, but to examine whether different modeling assumptions lead to consistent or substantially different interpretations.

This comparative process may reveal practical difficulties that are easily overlooked in an idealized modeling workflow. A more flexible model may capture nonlinear interactions but become difficult to interpret, whereas a simpler model may provide clearer relationships while failing to represent complex treatment-response patterns. Similarly, a model that performs well under internal validation may decline substantially when evaluated on data from another clinical setting.

The framework therefore places emphasis on model selection according to the intended clinical use. Where interpretability is particularly important, relatively transparent models may be preferable even if their predictive performance is modestly lower. Where complex interactions appear clinically plausible, more flexible models may be explored, but their outputs should be accompanied by appropriate interpretability methods.

Feature importance and model explanation should also be interpreted cautiously. An important feature is not necessarily a causal determinant of treatment response. Explanatory analyses may identify variables associated with predictions, but additional clinical research is required before such associations can be interpreted as mechanisms of therapeutic effect.

3.3 Model Validation and Clinical Interpretability

Validation should occur at several levels rather than being reduced to a single performance metric. Internal validation can examine the stability of the model within the development dataset, while cross-validation may provide a more reliable estimate of performance under repeated sampling. Where sufficient data are available, an independent test set should be retained to evaluate model behavior on previously unseen cases.

External validation is particularly important for TCM prediction models. Data obtained from a single institution may reflect local treatment preferences, practitioner expertise, patient demographics, and documentation conventions. A model that performs adequately within that environment may not necessarily generalize to another institution. Differences in treatment protocols may also alter the relationship between predictors and outcomes.

Clinical interpretability should be assessed alongside predictive performance. Clinicians need to understand not only what the model predicts but also which variables appear to contribute to that prediction and whether the reasoning is compatible with the clinical context. However, interpretability should not be mistaken for proof of causality. A visually intuitive explanation can still reflect biased or confounded relationships.

The proposed validation process should therefore include error analysis. Cases with unusually inaccurate predictions may reveal limitations in feature representation, outcome definition, data quality, or treatment heterogeneity. Rather than simply removing these cases as outliers, researchers should investigate whether they represent clinically meaningful subgroups. This iterative process may require repeated adjustments to the feature structure and model architecture, making model development less linear than a conventional “data-model-result” workflow might suggest.

4. Analytical Discussion of the Proposed Prediction Framework

4.1 Potential Predictive Value and Clinical Utility

A machine learning model for integrated TCM interventions may provide value by identifying response patterns that are difficult to detect through conventional clinical observation alone. In particular, the model may reveal interactions among baseline characteristics, treatment combinations, and clinical outcomes that are not easily represented through single-variable analyses.

Such predictions could potentially support patient stratification. Patients with broadly similar diagnoses may nevertheless exhibit different response patterns, and a predictive model could provide additional information concerning which characteristics are associated with those differences. This information may be useful when clinicians discuss treatment expectations, select monitoring strategies, or determine whether a treatment course should be reassessed.

The practical value of prediction, however, should not be measured only through statistical accuracy. A model with slightly lower predictive performance but greater transparency and easier integration into clinical workflows might be more useful than a highly complex model whose outputs are difficult to interpret. Conversely, excessive emphasis on interpretability could result in models that fail to capture clinically relevant nonlinear relationships.

Considering these factors, the most appropriate role of the proposed framework may be as an additional source of structured clinical information. Its predictions could complement practitioner judgment, particularly in cases where historical clinical data provide useful analogies, while leaving the final treatment decision within the clinical context.

4.2 Limitations, Biases, and Alternative Explanations

Several limitations may affect the reliability of machine learning-based therapeutic-response prediction. The first is selection bias. Patients included in a dataset may differ systematically from those who receive treatment elsewhere, and treatment decisions themselves may be influenced by practitioner experience or patient preference. A model trained on such data could consequently learn the characteristics of a particular clinical environment rather than general therapeutic principles.

The second issue concerns treatment confounding. In observational clinical data, patients with more severe symptoms may receive more intensive interventions. If these patients subsequently show poorer outcomes, the model might associate treatment intensity with poor response even when treatment intensity was partly a consequence of disease severity. Such relationships should not be interpreted as evidence that more intensive treatment causes poorer outcomes.

A third limitation concerns outcome measurement. Subjective symptom scores, practitioner assessments, and patient-reported outcomes may introduce measurement variability. Even apparently standardized outcomes can be influenced by follow-up timing and patient expectations. Consequently, differences between predicted and observed responses may reflect limitations in outcome measurement rather than failures of the machine learning algorithm alone.

There is also a possibility that machine learning models will reproduce existing clinical biases. If certain patient groups have historically received less intensive treatment or have been documented less systematically, the model may encode these disparities. Additional validation across different institutions and patient populations is therefore necessary before broad clinical implementation can be considered.

These limitations also lead to alternative interpretations of model performance. A high predictive score may indicate genuine clinical signal, but it may also partly reflect data leakage, treatment-selection patterns, or highly specific institutional characteristics. Conversely, apparently modest performance does not necessarily imply that machine learning has little clinical value, because the underlying treatment-response relationship may genuinely be difficult to predict from the available variables.

4.3 Future Development and Theoretical Implications

Future development may move beyond static prediction toward models capable of representing changes during treatment. A patient's response at one stage may influence subsequent treatment decisions, which in turn may alter later outcomes. This creates a dynamic relationship between intervention and response that is not fully captured by a single baseline prediction.

Longitudinal and time-sensitive modeling could consequently complement the present framework. Repeated symptom measurements, physiological indicators, treatment adjustments, and irregular follow-up observations may provide additional information about how individual responses evolve. Such approaches may also help distinguish patients who respond rapidly from those whose improvement emerges more gradually.

Another important direction is the integration of multimodal information. Structured clinical variables could potentially be combined with narrative medical records, imaging information, physiological signals, or patient-reported data. However, increasing the number of information sources also increases the risk of overfitting, missing-data problems, and interpretational complexity. The expansion of data should therefore be guided by clinical relevance rather than technological availability alone.

At a theoretical level, machine learning may contribute to a gradual shift from population-level descriptions of TCM intervention toward more individualized representations of therapeutic response. Yet such a shift should not be interpreted as evidence that clinical individuality can be completely reduced to predictive variables. Some aspects of TCM practice may remain context-dependent and difficult to quantify. The more productive research direction may consequently be an iterative interaction between computational modeling and clinical interpretation, in which model outputs generate new questions for clinicians and clinical reasoning, in turn, informs the refinement of computational representations.

5. Transitional Discussion

The proposed machine learning framework provides a methodological basis for examining therapeutic-response heterogeneity in integrated TCM interventions while recognizing that prediction remains constrained by data quality, treatment variability, outcome definition, and the distinction between statistical association and clinical causation. Its broader significance may lie not simply in producing a numerical prediction, but in creating a structured analytical environment in which patient characteristics, intervention patterns, and observed responses can be examined together and subsequently interpreted within a clinical context. At the same time, the limitations identified above suggest that future work should move toward longitudinal and externally validated modeling, particularly where treatment decisions and patient responses evolve interactively over time. This perspective creates a natural transition from static machine learning prediction toward more dynamic approaches capable of examining how therapeutic responses develop across successive observations, treatment adjustments, and changing patient states.

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] *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.*

[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] *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.*

[4] *Li, D., Hu, J., Zhang, L., Li, L., Yin, Q., Shi, J., ... & Zhuang, P. (2022). Deep learning and machine intelligence: new computational modeling techniques for discovery of the combination rules and pharmacodynamic characteristics of Traditional Chinese Medicine. European journal of pharmacology, 933, 175260.*

[5] *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.*

[6] *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.*

[7] *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.*

[8] *Shahrezaei, A., Taherkhani, S., Dashti, L., Garmaroodi, G. A., & Nasirinezhad, F. (2025). Herbal medicine meets machine learning: a systematic review of AI-powered innovation in chronic inflammation management. Discover Applied Sciences, 8(2), 111.*

[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] *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] *Yan, D., Zheng, Q., Chang, K., Hua, R., Liu, Y., Xue, J., ... & Zhou, X. (2025). Artificial intelligence in traditional Chinese medicine: from systems biological mechanism discovery, real-world clinical evidence inference to personalized clinical decision support. Chinese journal of natural medicines, 23(11), 1310-1328.*