SEGCN-AIKC: Trustworthy Herbal Prescription Recommendation via Knowledge-Guided Verification and Correction

Quan Gan
Jiangsu Ocean University
School of Computer Engineering
Lianyungang, China
ganquan@jou.edu.cn

Jiaqing Shang
Jiangsu Ocean University
School of Computer Engineering
Lianyungang, China
2024220926@jou.edu.cn

Lulu Wei
Jiangsu Ocean University
School of Computer Engineering
Lianyungang, China
i2U0llovo@163.com

Chen Li
Zhejiang University School of Medicine
Department of Human Genetics and Women's Hospital
Hangzhou, China
chenli2012@zju.edu.cn

Chuanxia Liu
Jiangsu Ocean University
School of Foreign Languages
Lianyungang, China
liuchuanxia@jou.edu.cn

Li Jia*
Jiangsu Ocean University
School of Computer Engineering
Lianyungang, China
ljia0624work@gmail.com

***Abstract*—Traditional Chinese Medicine (TCM) herbal prescription recommendation involves selecting appropriate herb combinations from patient symptoms under symbolic constraints derived from TCM theory. While most data-driven approaches achieve strong predictive accuracy, they typically encode domain knowledge only implicitly and lack explicit mechanisms to verify and correct knowledge-inconsistent predictions, which limits their reliability for clinical decision support. In this paper, we propose SEGCN-AIKC, a Semantic-Enhanced Graph Convolutional Network with Adaptive Inconsistency-aware Knowledge Correction, which formulates TCM herbal recommendation as a constrained hypothesis inference problem. Specifically, a semantic-enhanced multi-view graph convolutional network captures heterogeneous symptom–herb relations and constructs a syndrome-level hypothesis, which is then verified and minimally corrected by AIKC under symbolic TCM constraints. By integrating neural perception with constraint-guided verification and repair, SEGCN-AIKC bridges data-driven learning with symbolic reasoning, yielding recommendations that are both accurate and theoretically consistent. Experiments on a public TCM prescription dataset demonstrate consistent improvements over baselines, while ablation and case studies further validate interpretability and clinical plausibility.**

***Index Terms*—Traditional Chinese medicine, Herbal prescription recommendation, Graph Convolutional Network, Neuro-symbolic reasoning, Inconsistency-aware Correction**

I. Introduction

Traditional Chinese Medicine (TCM) constitutes a structured medical reasoning paradigm grounded in formal theoretical systems [1]. In clinical practice, diagnosis and treatment are operationalized through syndrome differentiation, which serves as a latent inference layer linking observable symptoms to therapeutic decisions. Herbal prescription formulation therefore emerges as a reasoning outcome rather than a direct pattern-matching process. Clinicians map heterogeneous symptom evidence to latent pathological patterns and select herb combinations consistent with the inferred syndromes. However, this mapping is weakly constrained and highly experience-dependent, leading to substantial inter-practitioner variability even under similar clinical presentations [2]. Such variability motivates computational approaches that can model prescribing regularities while explicitly adhering to TCM principles.

Formalizing this reasoning process computationally is non-trivial. Unlike conventional recommender systems based on flat user–item interactions, TCM prescription recommendation is governed by multi-relational symbolic structures linking symptoms, syndromes, and herbs [3]. Moreover, clinical prescription data are often sparse, long-tailed, and weakly supervised, which challenges the generalization of purely data-driven models across diverse clinical scenarios [4]. Existing learning-based approaches typically embed domain knowledge only implicitly within model parameters, lacking explicit mechanisms for verification, interpretation, and correction [5]. As a result, their predictions remain difficult to audit and unreliable for safety-critical clinical decision support.

Therefore, TCM prescription recommendation requires more than implicit black-box prediction models, but rather an interpretable reasoning framework with explicit theoretical consistency. From a prescription reasoning perspective, symptoms can be viewed as perceptual evidence, prescriptions as explanatory hypotheses, and inference as being governed by symbolic constraints derived from TCM theory. Based on this perspective, we propose SEGCN-AIKC, a Semantic-Enhanced Graph Convolutional Network with Adaptive Inconsistency-aware Knowledge Correction. The frame-

*Corresponding author

work integrates semantic-enhanced multi-view graph representation learning with an AIKC-based verification-and-repair mechanism that explicitly detects and corrects knowledge-inconsistent predictions during inference.

The main contributions are: (i) We formulate TCM herbal recommendation as a constrained neuro-symbolic reasoning problem and introduce a verification-and-repair paradigm. (ii) We propose an SEGCN encoder to model heterogeneous symptom–herb relations under data sparsity. (iii) We design an AIKC module that performs explicit consistency evaluation and minimal-intervention refinement.

II. Related Work

A. Rule-based and Data-driven Herbal Recommendation

Traditional prescription formulation in TCM follows syndrome differentiation and is inherently knowledge-intensive and experience-dependent [6]. Early computer-aided systems therefore adopted expert-system paradigms, in which syndrome rules and herb-compatibility knowledge were manually encoded to support decision-making [7]. Although interpretable, such systems suffer from limited coverage, high maintenance costs, and poor scalability [8].

To address these limitations, data-driven methods were introduced to learn prescribing regularities from clinical records. Network-based approaches extract co-occurrence structures to uncover herb compatibility patterns [9], while conventional machine-learning models learn mappings from symptoms to prescriptions [10]. Collaborative filtering and association-rule mining further identify frequently co-prescribed herb sets [11]. However, these methods rely mainly on correlation-based signals and lack explicit mechanisms to verify whether predicted prescriptions conform to normative TCM principles.

Recent advances in deep learning have significantly improved representation capacity for herbal recommendation. Graph neural network (GNN) is widely adopted to model relational structures among symptoms and herbs [12], including representative models such as MGCN [13], SMRGAT [14], and HPGCN [15], which incorporate multi-layer graphs, attention mechanisms, and pharmacological features.

Beyond GNN, reinforcement learning and subnetwork-based models have also been explored, such as PrescDRL [16] and TCMPR [17]. While these end-to-end models achieve strong predictive performance, most encode domain knowledge only implicitly within model parameters [18], and are not designed to explicitly assess or correct knowledge-inconsistent predictions during inference.

B. Neuro-Symbolic Reasoning and Verification

To overcome the limitations of both expert systems and purely neural models, recent studies integrate neural perception with symbolic constraints to ensure theoretical consistency. In such hybrid frameworks, neural components generate candidate hypotheses, while symbolic rules provide explicit specifications for verification and structured correction [19], enabling verifiability and interpretability.

In the TCM domain, existing work combines knowledge graph completion with quality evaluation to perform constraint-based consistency checking and flag contradictory triples [20]. Representation learning has also been adopted to construct TCM knowledge graphs for inference and decision support [21], and this research direction has been summarized in surveys covering rule-based, embedding-based, and neural paradigms [22]. These studies suggest that symbolic constraints offer effective mechanisms for both verifiable detection and structured repair, which naturally motivates the verification-and-repair formulation adopted in this work.

III. Approach

A. Problem Formulation and Overview of SEGCN-AIKC

TCM prescription recommendation aims to select a coherent set of herbs from observed symptoms while satisfying symbolic constraints derived from TCM theory. From a verification-and-repair perspective, we formulate the task as a *knowledge-constrained hypothesis inference* problem: a latent syndrome-level hypothesis is first induced from symptoms based on the learned symptom–herb relational space, and subsequently verified (and, if necessary, repaired) to ensure theoretical consistency.

Let $S = \{s_1, s_2, \ldots, s_N\}$ denote the set of symptom types and $H = \{h_1, h_2, \ldots, h_M\}$ denote the set of herb types. Each clinical prescription is represented as $p = \langle S_p, H_p \rangle$, where $S_p \subseteq S$ and $H_p \subseteq H$ denote the observed symptom set and the corresponding herb set, respectively. Given S_p, the goal is to learn a mapping $f : 2^S \rightarrow [0,1]^M$, which assigns a relevance score to each herb. The predicted score vector is $\hat{\mathbf{y}} = f(S_p; \theta, \mathcal{K})$, where θ denotes learnable parameters and $\mathcal{K}$ is an external TCM knowledge base encoding symbolic constraints. This defines a multi-label prediction task whose outputs must be both data-supported and theoretically consistent.

Fig. 1 illustrates SEGCN-AIKC, which separates neural representation learning from symbolic constraint enforcement at inference time. The framework consists of three stages: neural perception, syndrome-level hypothesis construction, and optional verification-and-repair via AIKC.

B. Unified Encoding and Multi-View Representation

We map symptoms and herbs into a unified latent space by encoding textual descriptions with a pre-trained BERT model and integrating semantic features with structured TCM attributes via linear projections:

$$\begin{aligned} \mathbf{e}_s^{(0)} &= f_{\text{sem}}(\mathrm{BERT}(\mathrm{desc}(s))), \\ \mathbf{e}_h^{(0)} &= f_{\text{sem}}(\mathrm{BERT}(\mathrm{desc}(h))) + f_{\text{attr}}(\mathbf{a}_h), \end{aligned} \tag{1}$$

All entities are projected into a shared embedding space of dimension $d = 256$, serving as the initial representations for subsequent learning.

To capture heterogeneous relational structures, we construct four complementary graphs over the unified entity set $V = S \cup H$: (i) a Symptom–Herb Graph (SHG) based on co-occurrence, (ii) a Symptom–Symptom Graph (SSG), (iii) a Herb–Herb Graph (HHG), and (iv) a Herbal Knowledge Graph

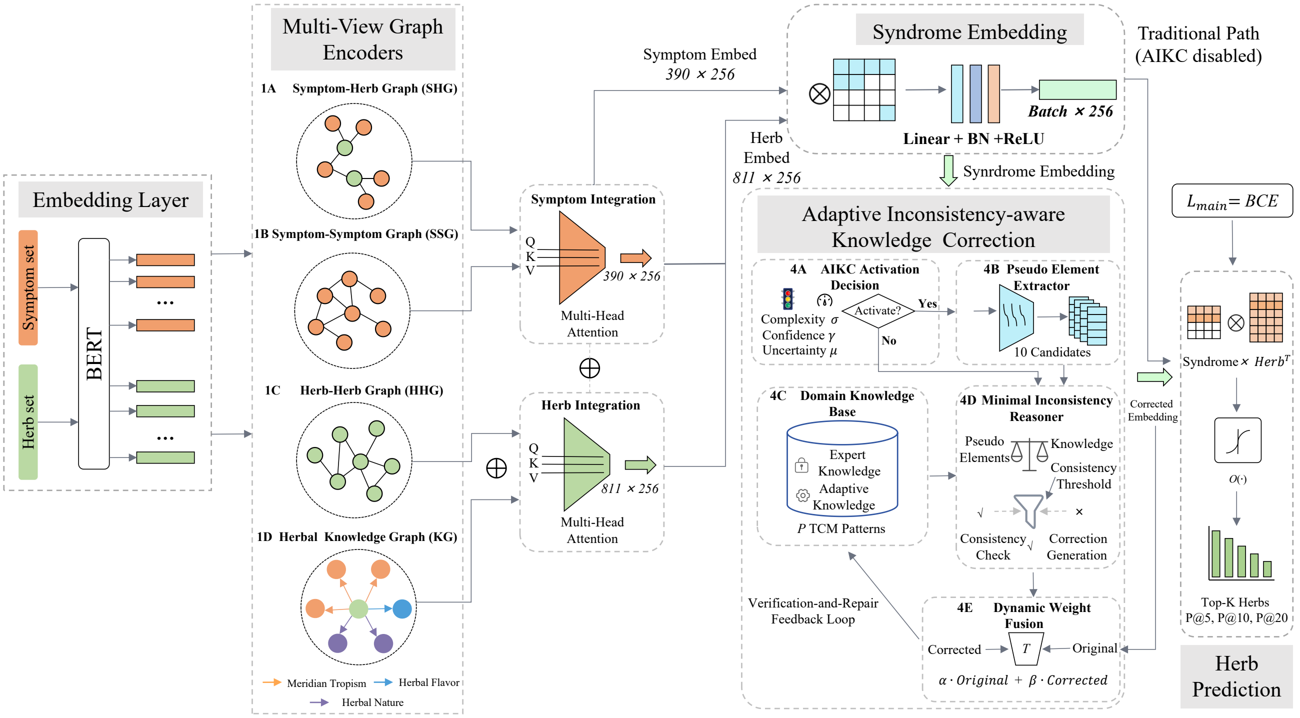

Fig. 1. Overview of SEGCN-AIKC. The framework follows a perception–hypothesis–verification–correction pipeline: (i) SEGCN encodes heterogeneous symptom–herb relations; (ii) representations are fused into a syndrome-level embedding; (iii) AIKC verifies and corrects predictions under TCM constraints.

(KG) derived from attribute similarity. Each graph is encoded using a two-layer GCN with residual connections, producing view-specific representations $\mathbf{Z}^{\text{sh}}, \mathbf{Z}^{\text{ss}}, \mathbf{Z}^{\text{hh}}, \mathbf{Z}^{\text{kg}}$. Multi-head attention is then applied to integrate information across views. Symptom representations aggregate SHG and SSG, while herb representations aggregate SHG, HHG, and KG, yielding fused embeddings $\mathbf{Z}_S \in \mathbb{R}^{N\times d}$ and $\mathbf{Z}_H \in \mathbb{R}^{M\times d}$.

Finally, mean-pooled symptom and herb embeddings are combined via outer-product fusion and nonlinear transformation to obtain a syndrome-level representation:

$$\mathbf{v}_{\text{syn}} = \text{ReLU}(\text{BN}(W_{\text{syn}}(\bar{\mathbf{z}}_S \otimes \bar{\mathbf{z}}_H))), \tag{2}$$

which serves as the latent hypothesis for subsequent constraint-guided verification and correction.

C. AIKC via Constraint-Guided Reasoning

Although the neural perception and hypothesis construction modules capture rich empirical regularities from data, they do not explicitly enforce prescription-level theoretical constraints during representation learning. Consequently, the model may produce overconfident predictions that violate key TCM principles. To address this limitation without affecting model optimization, we introduce an AIKC mechanism that operates exclusively at inference time to perform explicit verification and targeted repair under symbolic constraints.

Since not all predictions require symbolic intervention, AIKC is selectively activated based on three indicators derived from the syndrome representation $\mathbf{v}_{\text{syn}}$: complexity σ, confidence γ, and uncertainty μ. It is triggered when any indicator suggests unreliable inference:

$$\mathbb{I}_{\text{AKC}} = \mathbf{1}[\sigma > \tau_\sigma \;\vee\; \gamma < \tau_\gamma \;\vee\; \mu > \tau_\mu], \tag{3}$$

where τ_σ, τ_γ, and τ_μ are validation-tuned thresholds. If $\mathbb{I}_{\text{AIKC}} = 0$, predictions are generated directly from $\mathbf{v}_{\text{syn}}$; otherwise, verification and correction are applied.

From a reasoning perspective, inferring a valid prescription from observed symptoms can be formulated as a constrained abductive inference problem. Given a symptom set S_p and background knowledge $\mathcal{K}$, the objective is to identify a hypothesis H that both explains the observations and remains consistent with domain knowledge:

$$H^* = \arg\max_H P(S_p \mid H, \mathcal{K})\, P(H \mid \mathcal{K}). \tag{4}$$

Direct symbolic search over the combinatorial herb space is intractable; therefore, AIKC performs inference in a continuous latent space and treats $\mathbf{v}_{\text{syn}}$ as an initial hypothesis.

To explore alternative explanations while maintaining efficiency, AIKC generates a small set of pseudo-hypotheses through stochastic projection. Specifically, the syndrome representation is first mapped to a continuous latent space, from which multiple candidate hypotheses are produced using lightweight projection functions with controlled perturbations:

$$\mathbf{p}_i = G_{((i-1) \bmod 3)+1}(\mathbf{v}_{\text{syn}}) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \sigma_i^2 I), \tag{5}$$

where $\{G_k\}$ denote parameter-sharing projection operators with different stochastic behaviors, and ϵ_i introduces controlled noise to encourage hypothesis diversity. In practice, we generate $L = 10$ pseudo-hypotheses.

Domain knowledge is represented by a dual-layer prototype memory that balances theoretical stability and data adaptivity. The expert layer $K_{\text{expert}} \in \mathbb{R}^{P\times d}$ encodes canonical TCM patterns curated from theory, while the adaptive layer $K_{\text{adaptive}} \in \mathbb{R}^{P\times d}$ is learned from data to capture implicit regularities beyond expert specification. Given a query $\mathbf{q} = W_q \mathbf{v}_{\text{syn}}$, knowledge retrieval is performed via temperature-scaled attention and dynamically fused:

$$\mathbf{k}_{\text{fused}} = \sum_{t\in\{\text{expert},\text{adaptive}\}} \beta_t \cdot \text{softmax}\left(\frac{\mathbf{q}K_t^\top}{\tau_k}\right) K_t, \tag{6}$$

where $[\beta_{\text{expert}}, \beta_{\text{adaptive}}] = \text{softmax}(W_{\text{bal}}\mathbf{q})$ adaptively balance expert-curated and data-driven knowledge. The adaptive layer is trained jointly with neural perception modules and remains fixed during inference-time correction.

Symbolic TCM principles are instantiated as a constraint set $\mathcal{R}$. For a candidate prescription H_i induced by pseudo-hypothesis $\mathbf{p}_i$, each constraint produces a soft violation score $\phi_m(H_i) \in [0,1]$. The aggregated violation vector $\phi(H_i)$ provides a differentiable measure of theoretical inconsistency. Each pseudo-hypothesis is then evaluated using a lightweight consistency estimator:

$$c_i = \sigma\big(\text{MLP}\big([\mathbf{p}_i; \mathbf{k}_{\text{fused}}; \phi(H_i)]\big)\big), \tag{7}$$

yielding a soft consistency score $c_i \in [0,1]$.

Following a minimal-intervention principle, pseudo-hypotheses with high consistency are preserved, while low-consistency hypotheses are selectively adjusted in latent space. For inconsistent candidates, a correction vector $\Delta\mathbf{p}_i$ is generated and applied with an adaptive weight:

$$\mathbf{p}_i' = \mathbf{p}_i + \omega_i \cdot \Delta\mathbf{p}_i, \quad \omega_i = \sigma\big(\gamma\big((1-c_i)-\tau\big)\big), \tag{8}$$

where τ denotes the correction threshold and γ controls the sharpness of correction activation. This design ensures that only conflicted components are modified while preserving semantically coherent structure.

The corrected pseudo-hypotheses are then aggregated with the original syndrome representation through dynamic weight fusion:

$$\mathbf{v}_{\text{syn}}' = \alpha \cdot \mathbf{v}_{\text{syn}} + \beta \cdot \sum_{i=1}^{L} w_i \cdot \mathbf{p}_i', \tag{9}$$

where fusion weights $[\alpha, \beta]$ are computed adaptively based on the average consistency score and the query representation, ensuring a smooth trade-off between original perception and corrected reasoning.

In the final stage, AIKC supports iterative refinement through an abductive feedback loop, where the corrected hypothesis is re-submitted for verification over a small number of iterations. The process terminates early when either the average consistency score exceeds a predefined threshold or the maximum number of iterations is reached, thereby improving theoretical compliance while maintaining computational efficiency.

D. Prescription Prediction and Model Learning

Given the (possibly corrected) syndrome representation $\mathbf{v}_{\text{syn}}'$, herb relevance scores are computed by matching it against the herb embedding matrix:

$$\hat{\mathbf{y}} = \sigma(\mathbf{v}_{\text{syn}}' \mathbf{Z}_H^\top), \tag{10}$$

and the Top-K herbs are returned as recommendations.

Model parameters are trained using the binary cross-entropy loss with ground-truth multi-hot labels. During training, predictions are computed from the uncorrected representation $\mathbf{v}_{\text{syn}}$, while at inference time, $\mathbf{v}_{\text{syn}}$ is optionally refined by AIKC to obtain $\mathbf{v}_{\text{syn}}'$.

IV. Experiments

A. Experimental Settings

We evaluate SEGCN-AIKC on the public LMA dataset [23], which contains 33,765 prescriptions with 390 symptom types and 811 herb types. The data are randomly split into training, validation, and test sets with a ratio of 6:2:2. The validation set is used for hyperparameter tuning, and all results are reported on the test set. The model is implemented in PyTorch 2.1 and trained with Adam for 200 epochs.

Following prior work, we adopt Precision@K, Recall@K, and F1@K with $K \in \{5, 10, 20\}$ to evaluate top-K recommendation performance in this multi-label setting. Precision@K measures the proportion of correct herbs among the top-K predictions, Recall@K measures the coverage of ground-truth herbs, and F1@K provides their harmonic mean.

B. Overall Performance Comparison

Table I reports the overall performance on the test set. SEGCN-AIKC consistently achieves the best results across all metrics and values of K, indicating improved recommendation accuracy and stable generalization.

Compared with the strongest baseline PresRecST, SEGCN-AIKC improves F1@5 from 0.2352 to 0.2557, F1@10 from 0.2575 to 0.2862, and F1@20 from 0.2365 to 0.2661. Similar gains are observed for Precision and Recall, showing a consistent improvement in the precision–recall trade-off. Moreover, SEGCN-AIKC maintains the highest F1 scores across $K = 5, 10, 20$, suggesting that the verification-and-repair mechanism provides robust and generalizable corrections.

C. Ablation Analysis

The ablation results are reported in Table II, where each variant removes one or more components while keeping all other settings unchanged.

(i) Removing the multi-view representation module leads to clear performance degradation across all metrics. For example, Recall@10 drops from 0.3436 to 0.3285 and F1@10 decreases from 0.2862 to 0.2817, indicating the importance of integrating heterogeneous relational views. (ii) Disabling the AIKC module causes consistent declines, especially in recall, suggesting that AIKC effectively refines candidate prescriptions by removing knowledge-inconsistent hypotheses. (iii) Removing the BERT encoder results in a moderate performance drop, showing that textual semantics provide useful but auxiliary contextual signals. (iv) Joint ablations yield larger performance losses than individual removals, indicating that the three components capture complementary and non-redundant information, and their combination is necessary for optimal performance.

TABLE I
PERFORMANCE COMPARISON ON THE LMA DATASET.

Model	Precision			Recall			F1-Score		
	P@5	P@10	P@20	R@5	R@10	R@20	F1@5	F1@10	F1@20
LightGCN [24]	0.2121	0.1601	0.1156	0.1450	0.2157	0.3080	0.1723	0.1838	0.1681
SMGCN [25]	0.2472	0.1938	0.1419	0.1704	0.2644	0.3859	0.2017	0.2237	0.2075
PTM [23]	0.2526	0.2062	0.1523	0.1738	0.2818	0.4152	0.2059	0.2382	0.2229
TCMPR [17]	0.2688	0.2124	0.1579	0.1888	0.2989	0.4374	0.2218	0.2484	0.2320
KDHR [26]	0.2702	0.2108	0.1548	0.1960	0.2982	0.4311	0.2271	0.2470	0.2278
SMRGAT [14]	0.2819	0.2208	0.1611	0.2000	0.3058	0.4424	0.2324	0.2541	0.2344
PresRecST [27]	0.2775	0.2174	0.1594	0.2018	0.3086	0.4449	0.2352	0.2575	0.2365
SEGCN-AIKC	**0.3150**	**0.2502**	**0.1835**	**0.2237**	**0.3436**	**0.4841**	**0.2557**	**0.2862**	**0.2661**

TABLE II
ABLATION RESULTS ON THE LMA DATASET.

Model	Precision			Recall			F1-Score		
	P@5	P@10	P@20	R@5	R@10	R@20	F1@5	F1@10	F1@20
− Multi-View	0.3118	0.2466	0.1796	0.2112	0.3285	0.4739	0.2518	0.2817	0.2604
− AIKC	0.3134	0.2479	0.1818	0.2119	0.3300	0.4789	0.2529	0.2831	0.2635
− BERT	0.3139	0.2493	0.1839	0.2126	0.3328	0.4829	0.2533	0.2851	0.2657
− BERT & AIKC	0.3126	0.2481	0.1824	0.2123	0.3314	0.4809	0.2528	0.2837	0.2645
− Multi-View & AIKC	0.3102	0.2472	0.1800	0.2102	0.3291	0.4748	0.2506	0.2823	0.2610
− Multi-View & BERT	0.3117	0.2467	0.1767	0.2122	0.3301	0.4682	0.2525	0.2824	0.2566
SEGCN-AIKC (Full)	**0.3150**	**0.2502**	**0.1835**	**0.2237**	**0.3436**	**0.4841**	**0.2557**	**0.2862**	**0.2661**

TABLE III
CASE STUDIES OF AIKC. **BOLD** HERBS MATCH THE GROUND TRUTH, AND UNDERLINED HERBS ARE CORRECTED TO SATISFY CONSTRAINTS UNDER A MINIMAL-INTERVENTION PRINCIPLE.

Case	Symptom	Ground Truth Herbs	SEGCN Output	SEGCN-AIKC Output
Case 1	Headache	Licorice, Cimicifuga heracleifolia, Paeonia lactiflora, Gypsum, Ephedra sinica	**Licorice**, **Gypsum**, **Ephedra sinica**, Scutellaria baicalensis, Ligusticum chuanxiong, **Paeonia lactiflora**	**Licorice**, **Gypsum**, **Ephedra sinica**, **Paeonia lactiflora**, **Cimicifuga heracleifolia**, Angelica sinensis
Case 2	Vomiting	Atractylodes macrocephala, Fresh ginger, Licorice, Dried ginger, Clove, Agastache, Poria cocos, Ginseng, Pinellia, Magnolia officinalis	**Licorice**, **Ginseng**, **Atractylodes macrocephala**, **Fresh ginger**, Coptis chinensis, Amomum villosum, **Poria cocos**, **Magnolia officinalis**, **Pinellia**	**Licorice**, **Ginseng**, **Atractylodes macrocephala**, **Fresh ginger**, **Clove**, Aucklandia lappa decne, **Poria cocos**, **Magnolia officinalis**, **Pinellia**, Citrus reticulata

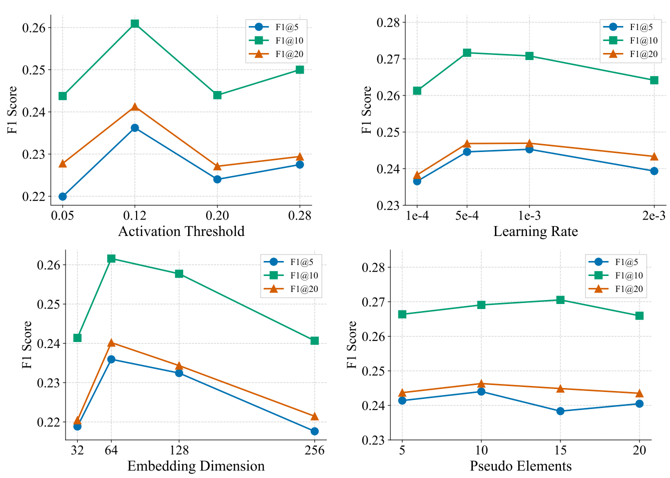

Fig. 2. Hyperparameter sensitivity analysis of SEGCN-AIKC.

D. Hyperparameter Sensitivity

We analyze the sensitivity of SEGCN-AIKC to key hyperparameters that affect the verification-and-repair mechanism and representation capacity, including the AIKC activation threshold τ, the number of pseudo elements, the embedding dimension, and the learning rate. Figure 2 reports F1@5, F1@10, and F1@20 under different settings.

As shown in Figure 2(a), performance follows a unimodal trend with a peak at $\tau = 0.12$. Too small values lead to excessive correction, while too large values suppress verification, confirming the effectiveness of the minimal-intervention design. Figure 2(b) indicates that a learning rate of 5×10^{-4} yields the most stable performance across all metrics. Smaller values slow convergence, whereas larger values cause instability. Figure 2(c) shows that increasing the embedding dimension from 32 to 64 consistently improves performance, while larger dimensions provide only marginal gains, suggesting that 64 offers a good trade-off between capacity and efficiency. Figure 2(d) demonstrates that performance improves with the number of pseudo elements up to 10 and then saturates, indicating that a moderate number of candidates is sufficient for effective correction.

Overall, SEGCN-AIKC exhibits stable performance across a wide range of hyperparameter settings, suggesting that the proposed framework is robust.

E. Case Study and Reasoning Analysis

To illustrate the interpretability of the proposed verification-and-repair mechanism, Table III present two representative cases showing how AIKC detects knowledge-inconsistent predictions and performs targeted corrections during inference.

Case 1 (Headache). The base SEGCN predicts a prescription containing the bitter-cold herb *Scutellaria baicalensis*, which violates the warm-cold balance of the inferred syndrome. AIKC identifies the violated constraint and performs a localized correction on the conflicted herb, replacing it with a semantically compatible alternative, *Cimicifuga heracleifolia*.

Case 2 (Vomiting). The preliminary output includes the bitter-cold herb *Coptis chinensis*, which contradicts the inferred syndrome's warming requirement. AIKC detects the inconsistency and replaces it with *Aucklandia lappa Decne*, restoring theoretical coherence.

These cases show that AIKC explicitly verifies candidate prescriptions under symbolic constraints and performs minimal, localized corrections, providing auditable and theory-consistent recommendations beyond black-box prediction.

V. Conclusion

In this paper, we propose SEGCN-AIKC, which integrates multi-view graph learning with knowledge correction to couple neural perception with symbolic constraint-based verification. The model achieves consistent improvements over baselines while ensuring theoretical consistency and interpretability, as supported by ablation and case analyses.

Despite these results, limitations remain. Evaluation is currently restricted to a single dataset, and future work will extend to more diverse clinical data. In addition, the lack of feedback from AIKC to the encoder may limit correction-aware representation learning.

Acknowledgment

This work was supported by the National Natural Science Foundation of China (82272126), Jiangsu Philosophy and Social Science (24YYC002) and JOU Marine Pharmaceutical Industry College Research Fund (JOUMIRF013).

References

[1] Y. Wang, "The scientific nature of traditional Chinese medicine in the post-modern era," *Journal of Traditional Chinese Medical Sciences*, vol.6, no. 3, pp. 195–200, 2019.

[2] T. R. Mamidi, M. Yamini, V. S. P. Cherukuri, and S. R. Yarraguntla, "Alternative medicine: New ways to treat diseases and therapies," *Indian Journal of Pharmaceutical Sciences*, vol.83, no. 1, 2021.

[3] W. Zhao, W. Lu, Z. Li, C. Zhou, H. Fan, Z. Yang, and C. Li, "TCM herbal prescription recommendation model based on multi-graph convolutional network," *Journal of Ethnopharmacology*, vol.297, 115109, 2022.

[4] H. Hu, Y. Li, Z. Zheng, W. Hu, R. Lin, and Y. Kang, "A traditional Chinese medicine prescription recommendation model based on contrastive pre-training and hierarchical structure network," *Expert Systems With Applications*, vol. 268, 126318, 2025.

[5] M. Elhaddad and S. Hamam, "AI-driven clinical decision support systems: An ongoing pursuit of potential," *Cureus*, vol.16, e57728, 2024.

[6] F. Cheng, X. Wang, W. Song, Y. Lu, X. Li, H. Zhang, and Q. Wang, "Biologic basis of TCM syndromes and the standardization of syndrome classification," *Journal of Traditional Chinese Medical Sciences*, vol. 1, no. 2, pp. 92–97, 2014.

[7] X. Wang, H. Qu, P. Liu, and Y. Cheng, "A self-learning expert system for diagnosis in traditional Chinese medicine," *Expert Systems With Applications*, vol. 26, no. 4, pp. 557–566, 2003.

[8] T. Dash, S. Chitlangia, A. Ahuja, and A. Srinivasan, "A review of some techniques for inclusion of domain-knowledge into deep neural networks," *Scientific Reports*, vol. 12, no. 1, 1040, 2022.

[9] N. Wang, N. Du, Y. Peng, K. Yang, Z. Shu, K. Chang, and X. Zhou, "Network patterns of herbal combinations in traditional Chinese clinical prescriptions," *Frontiers in Pharmacology*, vol.11, 590824, 2021.

[10] W. Lee, Y. Lee, S. Lee, Y. K. Kim, and J. Kim, "A machine learning approach for recommending herbal formulae with enhanced interpretability and applicability," *Biomolecules*, vol. 12, no. 11,1604, 2022.

[11] J. Yang, Q. Wang, S. Zhou, F. Zhang, and Z. Lv, "Construction of Chinese medicine prescription data mining based on association rules of the system," *Exploratory Research in Hypothesis Medicine*, vol. 8, no. 4, pp. 307–318, 2023.

[12] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *CoRR*, 2016.

[13] S. Wang, Y. Zhu, and J. Liu, "MGCN: Graph convolutional network for drug molecular screening," *Journal of Physics: Conference Series*, vol. 1744, 042241, 2021.

[14] X. Yang and C. Ding, "SMRGAT: A traditional Chinese herb recommendation model based on a multi-graph residual attention network and semantic knowledge fusion," *Journal of Ethnopharmacology*, vol. 315,116693, 2023.

[15] Q. Niu, J. Wang, H. Li, L. Tong, H. Xu, W. Zhang, and B. Li, "HPGCN: A graph convolutional network-based prediction model for herbal heat/cold properties," *Current Plant Biology*, vol.41,100448,2025.

[16] K. Yang, Z. Yu, X. Su, F. Zhang, X. He, N. Wang, and X. Zhou, "Presc-DRL: Deep reinforcement learning for herbal prescription planning in treatment of chronic diseases," *Chinese Medicine*, vol.19, 144, 2024.

[17] X. Dong, Y. Zheng, Z. Shu, K. Chang, J. Xia, Q. Zhu, and X. Zhou, "TCMPR: TCM prescription recommendation based on subnetwork term mapping and deep learning," *BioMed Research International*, vol. 2022, 4845726, 2022.

[18] Q. Hu, T. Yu, J. Li, Q. Yu, L. Zhu, and Y. Gu, "End-to-end syndrome differentiation of Yin deficiency and Yang deficiency in traditional Chinese medicine," *Computer Methods and Programs in Biomedicine*, vol. 174, pp. 9–15, 2019.

[19] Z.-H. Zhou, "Abductive learning: Towards bridging machine learning and logical reasoning," *Science China Information Sciences*, vol. 62, no. 7, 76101, 2019.

[20] C. Liu, Z. Li, J. Li, Y. Qu, Y. Chang, Q. Han, and S. Lin, "Research on traditional Chinese medicine: Domain knowledge graph completion and quality evaluation," *JMIR Medical Informatics*, vol. 12,e55090, 2024.

[21] H. Weng, J. Chen, A. Ou, and Y. Lao, "Leveraging representation learning for the construction and application of a knowledge graph for traditional Chinese medicine: Framework development study," *JMIR Medical Informatics*, vol. 10, no. 9,e38414, 2022.

[22] Z. Guo, Q. Liu, and B. Zou, "Research on knowledge reasoning of TCM based on knowledge graphs," *Digital Chinese Medicine*, vol. 5, no. 4, pp. 386–393, 2022.

[23] L. Yao, Y. Zhang, B. Wei, W. Zhang, and Z. Jin, "A topic modeling approach for traditional Chinese medicine prescriptions," *IEEE Transactions on Knowledge and Data Engineering*, vol. 30, no. 6, pp. 1007–1021, 2018.

[24] X. He, K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang, "LightGCN: Simplifying and powering graph convolution network for recommendation," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval(ICDE)*,Dallas, 2020, pp. 639–648.

[25] Y. Jin, W. Zhang, X. He, X. Wang, and X. Wang, "Syndrome-aware herb recommendation with multi-graph convolution network," in *2020 IEEE 36th International Conference on Data Engineering*, 2020, pp. 145–156.

[26] Y. Yang, Y. Rao, M. Yu, and Y. Kang, "Multi-layer information fusion based on graph convolutional network for knowledge-driven herb recommendation," *Neural Networks*, vol. 146, pp. 1–10, 2022.

[27] X. Dong, C. Zhao, X. Song, L. Zhang, Y. Liu, J. Wu, and X. Zhou, "PresRecST: A novel herbal prescription recommendation algorithm for real-world patients with integration of syndrome differentiation and treatment planning," *Journal of the American Medical Informatics Association*, 2024.