PRISM: PROGRESSIVE PRESCRIPTION REASONING VIA INTEGRATED HIERARCHICAL LATENT AND ROLE-AWARE SYMPTOM-HERB MODELING FOR PRESCRIPTION RECOMMENDATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Traditional Chinese Medicine (TCM) prescription recommendation aims to generate effective herb combinations from complex symptom patterns. However, existing approaches suffer from three limitations: insufficient hierarchical modeling of symptom–herb relationships, failure to simulate hierarchical prescription construction, and lack of interpretable reasoning trajectories. In this paper, we propose PRISM (Progressive Prescription Reasoning via Integrated Hierarchical Latent and Role-aware Symptom-herb Modeling), a progressive reasoning framework for TCM prescription recommendation. PRISM introduces hierarchical symptom–herb graph reasoning to capture multi-level clinical dependencies, adaptive main–auxiliary prescription reasoning to model progressive prescription construction, and progressive latent prescription states to provide interpretable reasoning trajectories. Extensive experiments on widely used TCM prescription datasets demonstrate that PRISM consistently outperforms existing baselines. The source code and datasets are publicly available at `https://anonymous.4open.science/r/PRISM-A452/`.

1 INTRODUCTION

Traditional Chinese Medicine (TCM) prescription recommendation aims to automatically generate effective herb combinations from patients' symptoms. Unlike conventional drug recommendation, TCM prescription formulation follows the holistic principle of syndrome differentiation and treatment, where physicians interpret symptom patterns and select compatible herbs according to therapeutic considerations. Therefore, prescription recommendation can be viewed as a structured clinical reasoning problem that requires modeling complex interactions among symptoms and herbs.

Recent advances in artificial intelligence have promoted computational TCM prescription recommendation. Early studies mainly formulated the task as multi-label classification, directly predicting herbs from symptom representations (Yao et al., 2018; Li & Yang, 2019). Subsequent methods introduced topic models (Wang et al., 2019b; Jin et al., 2020b) and sequence generation approaches (Li & Yang, 2019; Wang et al., 2019c) to capture latent symptom–prescription patterns. More recently, graph-based methods (Jin et al., 2020b; Yang & Ding, 2023) have incorporated TCM knowledge and relational information among symptoms and herbs. Despite these advances, existing methods still face several limitations in modeling the actual prescription reasoning process, as illustrated in Figure 1.

Challenge 1: Insufficient hierarchical symptom–herb modeling. Existing approaches (Liu et al., 2023; Hu et al., 2025; Zhang & Tan, 2025) mainly capture direct symptom–herb associations or aggregate local graph neighborhoods. However, prescription reasoning involves multiple relational levels: symptom–symptom associations characterize clinical patterns, symptom–herb relations connect symptoms to candidate treatments, and herb–herb relations encode compatibility. Ignoring such hierarchical dependencies limits the modeling of clinical patterns and herb interactions.

Challenge 2: Failure to model hierarchical prescription construction. TCM prescriptions are not formed by selecting all herbs simultaneously. Guided by the "monarch, minister, assistant, and

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

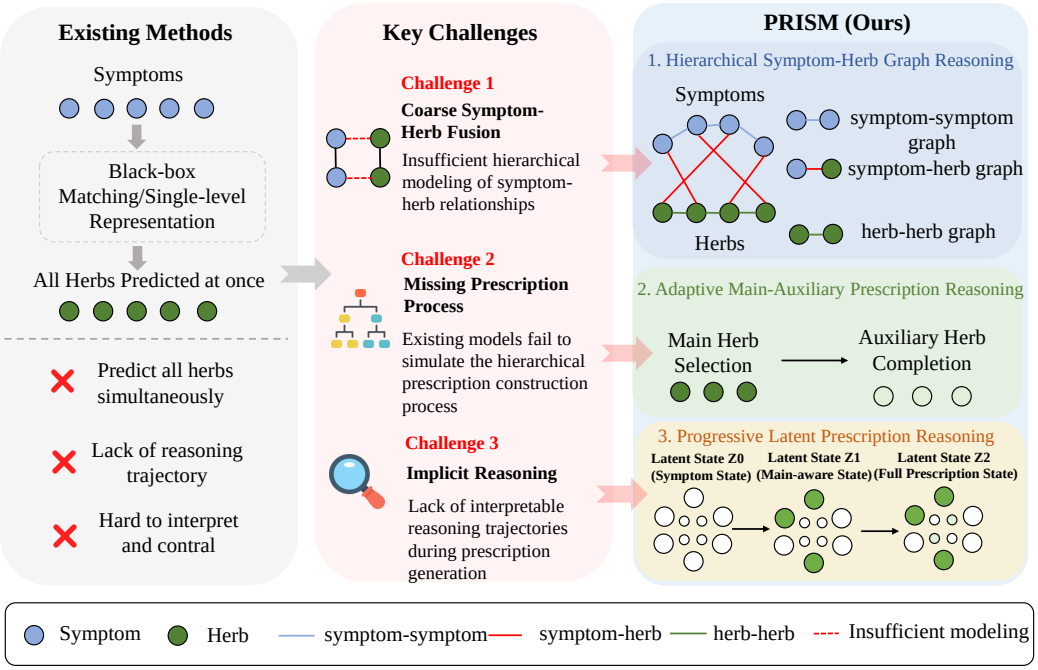

Figure 1: Overview of the key challenges in existing TCM prescription recommendation methods and the proposed PRISM framework.

envoy" principle, physicians typically identify core therapeutic herbs before adding auxiliary herbs for complementary effects. Existing methods (Zhao et al., 2022; Jin et al., 2020b; Dong et al., 2024) rarely model this main–auxiliary hierarchy, partly because large-scale datasets lack explicit herb-role annotations, making it difficult to reproduce the progressive prescription construction process.

Challenge 3: Lack of interpretable reasoning trajectories. Recent graph-based methods (Yang et al., 2022b; Qi et al., 2023; Yang & Ding, 2023) generally map symptoms directly to final herb predictions, providing limited insight into intermediate decision states (He et al., 2026; Zheng et al., 2026). A desirable model should not only recommend accurate prescriptions, but also reveal how symptom evidence progressively evolves into the final prescription (Wang et al., 2025; Zhang & Tan, 2025).

To address these challenges, we propose **PRISM**, a Progressive Prescription Reasoning framework via Hierarchical Latent and Role-aware Symptom-Herb Modeling. PRISM first performs **Hierarchical Symptom-Herb Graph Reasoning** over symptom–symptom (SS), symptom–herb (SH), and herb–herb (HH) relations to capture progressive clinical dependencies. It then introduces **Adaptive Main-Auxiliary Prescription Reasoning**, which automatically discovers herb roles from prescription statistics and performs main-herb selection followed by main-guided auxiliary completion. Finally, **Progressive Latent Prescription Reasoning** models three intermediate states, Z_0, Z_1, and Z_2, corresponding to symptom, main-aware, and full-prescription states, thereby connecting symptom evidence, hierarchical herb selection, and final prediction through an explicit reasoning trajectory.

Our main contributions are summarized as follows:

1. We propose **PRISM**, which reformulates TCM prescription recommendation as a progressive reasoning process and introduces hierarchical SS–SH–HH graph reasoning to capture multi-level clinical and herb interactions.

2. We develop an **adaptive main–auxiliary prescription reasoning** mechanism with data-driven role discovery, together with progressive latent states that explicitly model the transition from symptoms to main herbs and complete prescriptions.

3. Extensive experiments on real-world TCM prescription datasets demonstrate the effectiveness of PRISM, with ablation and reasoning analyses validating its key components.

2 Related Work

2.1 TCM Prescription Recommendation

Traditional Chinese Medicine (TCM) prescription recommendation aims to generate effective herb combinations from patients' observed symptoms. Early studies mainly formulated this task as multi-label classification, directly predicting herbs from symptom representations (Jin et al., 2020a). Subsequent approaches explored topic modeling (Yao et al., 2018; Ji et al., 2017; Wang et al., 2019b) and sequence-based generation (Li & Yang, 2019; Wang et al., 2019c; Hou et al., 2023) to capture latent patterns and dependencies between symptoms and prescriptions. More recently, graph-based methods have incorporated external TCM knowledge and relational information among symptoms, herbs, and prescriptions (Jin et al., 2020a; 2021; 2023; Yang et al., 2022a; Yang & Ding, 2023; Zhao et al., 2022; Dong et al., 2024). Despite their effectiveness, existing methods mainly focus on improving symptom–herb representations or relational modeling, while the structured reasoning process underlying prescription formation remains insufficiently explored. In particular, how multi-level clinical relations and functional herb roles progressively contribute to a complete prescription remains largely under-modeled.

2.2 Structured and Explainable Recommendation

Multi-label recommendation requires modeling dependencies among multiple prediction targets rather than treating them independently (Zhang & Zhou, 2013; Tsoumakas et al., 2010). Existing studies have explored structured prediction through classifier chains, label correlations, and ranking-based objectives (Read et al., 2011; Wang et al., 2016; Chen et al., 2019; Li et al., 2017). However, these approaches primarily capture statistical dependencies among labels, while the progressive construction of structured outputs remains less explored. This limitation is particularly relevant to TCM prescriptions, where herbs may play different functional roles and collectively form a structured therapeutic combination. Interpretability is also important in recommendation and healthcare systems (Zhang & Chen, 2020; Rajabi & Kafaie, 2022). Existing approaches improve transparency through attention mechanisms, knowledge integration, latent representations, and intermediate supervision (Wang et al., 2019a; 2018; Zhang et al., 2014; Wang et al., 2022). In TCM recommendation, domain knowledge and symptom–herb relations have also been exploited to enhance interpretability (Jin et al., 2021; 2023; Yang & Ding, 2023). Nevertheless, most approaches explain final predictions or learn implicit representations rather than explicitly modeling how intermediate clinical evidence evolves into a complete prescription (Xian et al., 2019; Geng et al., 2022). PRISM instead models prescription formation as a progressive reasoning process with explicit herb roles and intermediate prescription states.

3 Method

3.1 Overview

Given a symptom set, PRISM formulates prescription recommendation as a progressive reasoning process rather than direct symptom–herb matching. As illustrated in Figure 2, PRISM consists of four components: hierarchical symptom–herb graph reasoning, adaptive main–auxiliary prescription reasoning, progressive latent prescription reasoning, and role-aware ranking.

The overall reasoning trajectory is summarized as

$$\mathbf{S}^0 \rightarrow \mathbf{S}^{SS} \rightarrow (\mathbf{S}^{SH}, \mathbf{H}^{SH}) \rightarrow \mathbf{H}^{HH} \rightarrow \mathbf{z}_0 \rightarrow \mathbf{z}_1 \rightarrow \mathbf{z}_2 \rightarrow \mathbf{l}_{final}, \tag{1}$$

where $\mathbf{z}_0$, $\mathbf{z}_1$, and $\mathbf{z}_2$ represent the symptom-driven, main-aware, and completed prescription states, respectively.

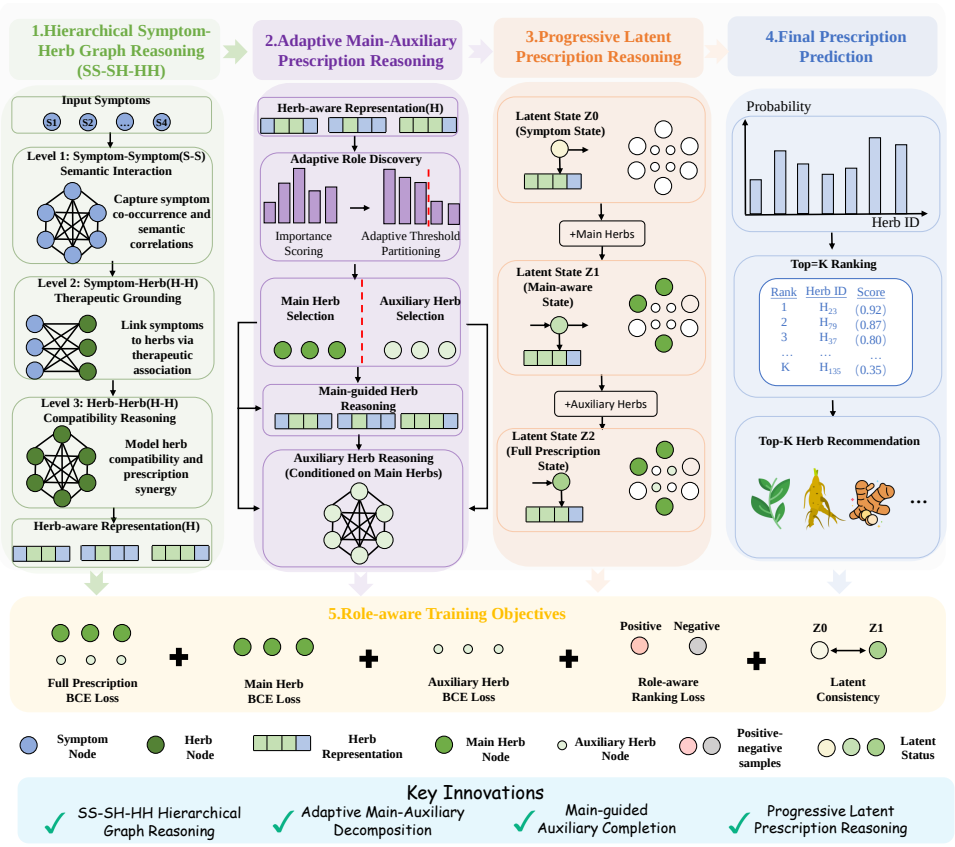

Figure 2: Overview of our framework.

3.2 Hierarchical Symptom-Herb Graph Reasoning

TCM prescription formation involves multiple levels of clinical dependencies. We therefore construct a hierarchical graph $\mathcal{G} = \{\mathcal{G}_{SS}, \mathcal{G}_{SH}, \mathcal{G}_{HH}\}$, corresponding to symptom–symptom, symptom–herb, and herb–herb relations, respectively. All graph structures and structural affinity matrices are constructed exclusively from the training set.

Let $\mathbf{S}^0$ and $\mathbf{H}^0$ denote the initial symptom and herb embeddings. We first propagate symptom information over $\mathcal{G}_{SS}$:

$$\mathbf{S}^{SS} = \mathrm{GAT}_{SS}(\mathbf{S}^0, \mathcal{G}_{SS}). \tag{2}$$

To preserve original symptom semantics, the SS-enhanced representations are gated with the initial embeddings as $\widetilde{\mathbf{S}}^{SS} = \mathbf{G} \odot \mathbf{S}^{SS} + (1 - \mathbf{G}) \odot \mathbf{S}^0$, where $\mathbf{G} = \sigma(\mathbf{W}_g[\mathbf{S}^0 \| \mathbf{S}^{SS}])$.

The refined symptom representations are subsequently propagated through the symptom–herb graph and then the herb–herb graph:

$$\begin{aligned} (\mathbf{S}^{SH}, \mathbf{H}^{SH}) &= \mathrm{HGAT}_{SH}(\widetilde{\mathbf{S}}^{SS}, \mathbf{H}^0, \mathcal{G}_{SH}), \\ \mathbf{H}^{HH} &= \mathrm{GAT}_{HH}(\mathbf{H}^{SH}, \mathcal{G}_{HH}). \end{aligned} \tag{3}$$

Finally, multi-level herb representations are fused as

$$\mathbf{H} = f_h([\mathbf{H}^0 \| \mathbf{H}^{SH} \| \mathbf{H}^{HH}]). \tag{4}$$

For each prescription, we mean-pool the active symptom nodes from $\mathbf{S}^{SS}$ and $\mathbf{S}^{SH}$ to obtain prescription-level representations $\mathbf{s}_{SS}$ and $\mathbf{s}_{SH}$, which are used for subsequent prescription reasoning.

3.3 Adaptive Main-Auxiliary Prescription Reasoning

Rather than treating a prescription as a flat herb set, PRISM decomposes the gold prescription into disjoint main and auxiliary subsets, $\mathcal{H} = \mathcal{H}_{main} \cup \mathcal{H}_{aux}$.

3.3.1 Adaptive Herb Role Discovery

Since explicit herb-role annotations are unavailable, we derive weak role supervision from training prescriptions. For symptom set $\mathcal{S}$, we combine its exact symptom-set herb probability with a single-symptom back-off estimate through support-dependent shrinkage:

$$\widetilde{p}(h_j|\mathcal{S}) = \rho_{\mathcal{S}} p_{\text{exact}}(h_j|\mathcal{S}) + (1-\rho_{\mathcal{S}}) p_{\text{sym}}(h_j|\mathcal{S}), \qquad \rho_{\mathcal{S}} = \frac{C(\mathcal{S})}{C(\mathcal{S})+\kappa}. \tag{5}$$

To reduce the dominance of globally frequent herbs, we reweight the conditional probability by its relative lift with respect to the global herb prior. Specifically, the specificity score is

$$r_j = \widetilde{p}(h_j|\mathcal{S}) \sqrt{\min\left(\frac{\widetilde{p}(h_j|\mathcal{S})}{p(h_j)+\epsilon}, c\right)}, \tag{6}$$

where c caps excessively large lift values caused by rare herbs. The positive herbs are sorted by r_j, and the adaptive split point is determined by the largest adjacent score gap, $k^* = \arg\max_k (r_{(k)} - r_{(k+1)})$. The top-$k^*$ herbs form $\mathcal{H}_{main}$ and the remaining positive herbs form $\mathcal{H}_{aux}$. This produces sample-specific role assignments without requiring a predefined number of main herbs.

3.3.2 Main-Herb Reasoning

We first construct a symptom-conditioned prescription query $\mathbf{q} = f_q([\mathbf{s}_{SS} \| \mathbf{s}_{SH}])$ and initialize the latent prescription state as $\mathbf{z}_0 = f_0([\mathbf{s}_{SS} \| \mathbf{s}_{SH} \| \mathbf{q}])$.

Given the fused herb representation $\mathbf{H}$, the main-herb logits are computed by

$$\mathbf{l}_{main} = \frac{\mathbf{z}_0 \mathbf{K}_m^\top}{\sqrt{d\tau}} + \mathbf{b}_m + \eta_s \log(1 + \mathbf{x}\mathbf{A}_{SH}), \tag{7}$$

where $\mathbf{K}_m$ is the projected herb representation and $\mathbf{A}_{SH}$ provides symptom–herb structural evidence. We obtain $\mathbf{p}_{main} = \text{softmax}(\mathbf{l}_{main})$ and aggregate the main-herb context as $\mathbf{c}_{main} = \mathbf{p}_{main}\mathbf{V}_m$.

3.3.3 Main-Guided Auxiliary Completion

The initial prescription state is first updated using the predicted main herbs:

$$\mathbf{z}_1 = \text{LN}\left(\mathbf{z}_0 + \text{GRU}_m(\mathbf{c}_{main}, \mathbf{z}_0)\right). \tag{8}$$

Auxiliary-herb reasoning is then conditioned on both the main-aware state and main-herb context. Specifically, $\mathbf{q}_{aux} = f_{aux}([\mathbf{z}_1 \| \mathbf{c}_{main}])$, and

$$\mathbf{l}_{aux} = \frac{\mathbf{q}_{aux}\mathbf{K}_a^\top}{\sqrt{d\tau}} + \eta_s \mathbf{r}_S + \eta_h \log\left(1 + \mathbf{p}_{main}\mathbf{A}_{HH}\right), \tag{9}$$

where $\mathbf{r}_S = \log(1 + \mathbf{x}\mathbf{A}_{SH})$ and $\mathbf{A}_{HH}$ represents herb compatibility.

The normalized auxiliary distribution is obtained from $\sigma(\mathbf{l}_{aux})$, and the auxiliary context is denoted by $\mathbf{c}_{aux}$.

3.4 Progressive Latent Prescription Reasoning

The main–auxiliary hierarchy determines the order of herb generation, while the latent states explicitly represent how the prescription evolves through this process. PRISM maintains $\mathbf{z}_0 \to \mathbf{z}_1 \to \mathbf{z}_2$, where $\mathbf{z}_0$ encodes symptom evidence and $\mathbf{z}_1$ incorporates main-herb information.

After auxiliary-herb reasoning, the final latent prescription state is updated as

$$\mathbf{z}_2 = \mathrm{LN}\left(\mathbf{z}_1 + \mathrm{GRU}_a(\mathbf{c}_{aux}, \mathbf{z}_1)\right). \tag{10}$$

The latent state is decoded into prescription scores by $\mathbf{l}_{latent} = \mathbf{z}_2\mathbf{K}_a^\top/(\sqrt{d}\tau)$. Meanwhile, a global prescription representation $[\mathbf{s}_{SS}\|\mathbf{s}_{SH}\|\mathbf{c}_{main}\|\mathbf{z}_2]$ is mapped to $\mathbf{l}_{global}$ through a learnable fusion network. The final prescription prediction combines complementary reasoning paths:

$$\boxed{\mathbf{l}_{final} = \mathbf{l}_{global} + \mathbf{l}_{latent} + \gamma_m\mathbf{l}_{main} + \gamma_a\mathbf{l}_{aux}} \tag{11}$$

where γ_m and γ_a are learnable non-negative coefficients implemented using softplus parameterization.

3.5 Role-Aware Adaptive Ranking

Multi-label prescription recommendation contains substantially more negative herbs than positive herbs, and classification loss alone may fail to sufficiently separate relevant herbs from competitive negatives. We therefore introduce a role-aware hard-negative ranking objective. For each prescription, let $\mathcal{P} = \{h_j \mid y_j = 1\}$ denote the positive herbs and let $\mathcal{N}$ contain the top-K highest-scoring incorrect herbs according to $\mathbf{l}_{final}$. The ranking objective is

$$\mathcal{L}_{rank} = \frac{1}{B}\sum_{i=1}^{B}\frac{\sum_{p\in\mathcal{P}_i}\sum_{n\in\mathcal{N}_i} w_p[m - l^{(i)}_{final,p} + l^{(i)}_{final,n}]_+}{|\mathcal{N}_i|\sum_{p\in\mathcal{P}_i} w_p}, \tag{12}$$

where m is the margin and $w_p = \alpha$ for main herbs and $w_p = \beta$ for auxiliary herbs, with $\alpha > \beta$. Thus, ranking errors involving main herbs receive stronger penalties.

3.6 Training Objective

PRISM is jointly supervised by complete-prescription, main-herb, and auxiliary-herb prediction. We apply weighted binary cross-entropy to $\mathbf{l}_{final}$, $\mathbf{l}_{main}$, and $\mathbf{l}_{aux}$, resulting in $\mathcal{L}_{full}$, $\mathcal{L}_{main}$, and $\mathcal{L}_{aux}$, respectively. The overall objective is

$$\boxed{\mathcal{L} = \lambda_f\mathcal{L}_{full} + \lambda_m\mathcal{L}_{main} + \lambda_a\mathcal{L}_{aux} + \lambda_r\mathcal{L}_{rank}} \tag{13}$$

where the loss weights are set to $\lambda_f = 1.0$ and $\lambda_m = \lambda_a = \lambda_r = 0.1$.

4 Experiments

4.1 Datasets

We conduct experiments on the widely used TCM Prescription Dataset (**TCM-PD**) (Yang et al., 2022a), which contains 33,765 prescriptions, 390 symptoms, and 805 herbs. Following previous studies (Yang et al., 2022a; Hu et al., 2025), we use 20,709 prescriptions for training and 13,056 for testing, without an additional validation set. Symptoms and herbs are represented as multi-hot vectors. All SS, SH, and HH relations, as well as the statistics for adaptive main–auxiliary role discovery, are constructed exclusively from the training set to avoid information leakage.

4.2 Baselines

We compare PRISM with six representative TCM prescription recommendation methods: **SMGCN** (Jin et al., 2020a), which models SS, SH, and HH relations using graph convolution; **KDHR** (Yang et al., 2022a), which incorporates TCM knowledge into graph-based recommendation; **SMRGAT** (Yang & Ding, 2023), which employs multi-head graph attention and residual learning; **SCEIKG** (Liu et al., 2023), which models sequential condition evolution with knowledge graphs; **PresRecST** (Dong et al., 2024), which models progressive TCM diagnosis and treatment; and **TCMRGCL** (Hu et al., 2025), which combines graph contrastive learning with hierarchical structure modeling. All methods are evaluated under the same TCM-PD data split and evaluation protocol.

4.2.1 EVALUATION METRICS

Following previous studies (Yang et al., 2022a), we evaluate prescription recommendation using **Precision@K**, **Recall@K**, and **F1-score@K**, where $K \in \{5, 10, 20\}$. Detailed metric definitions are provided in the Appendix.

4.2.2 IMPLEMENTATION DETAILS

PRISM is implemented in PyTorch and optimized using Adam. We use an embedding dimension of 128 and a hidden dimension of 256, and train the model with a batch size of 512 and a learning rate of 3×10^{-4}. All experiments use a fixed random seed of 2021. Detailed implementation and hyperparameter settings are provided in Appendix A.4.

4.3 OVERALL PERFORMANCE

Table 1: Performance on TCM_PD compared with the baseline models.

Dataset	Model	Precision			Recall			F1-score		
		P@5	P@10	P@20	R@5	R@10	R@20	F1@5	F1@10	F1@20
TCM_PD	SMGCN	0.1933	0.1568	0.1153	0.1265	0.2031	0.3027	0.1529	0.1770	0.1670
	KDHR	0.2138	0.1660	0.1251	0.1510	0.2284	0.3414	0.1770	0.1922	0.1832
	SMRGAT	0.2461	0.1873	0.1343	0.1821	0.2777	0.3970	0.2093	0.2237	0.2007
	SCEIKG	0.2482	0.1942	0.1446	0.1722	0.2691	0.3973	0.2033	0.2256	0.2119
	PresRecST	0.2238	0.1749	0.1290	0.1512	0.2338	0.3465	0.1805	0.2001	0.1879
	TCMRGCL	0.2543	0.1984	0.1477	0.1852	0.2841	0.4192	0.2143	0.2337	0.2184
	Ours	**0.2726**	**0.2145**	**0.1577**	**0.2058**	**0.3239**	**0.4762**	**0.2346**	**0.2581**	**0.2369**

Table 1 presents the overall performance comparison between our method and representative baseline models on TCM_PD dataset. Our method consistently outperforms all existing approaches across different evaluation settings, demonstrating its effectiveness for multi-label prescription recommendation. Specifically, compared with the strongest baseline TCMRGCL, our method achieves improvements on all evaluation metrics. On Dataset1, our method obtains the best results with Precision@5, Recall@5, and F1@5 of 0.2726, 0.2058, and 0.2346, respectively. Moreover, the improvements become more significant under larger recommendation lists, where our method achieves Recall@20 of 0.4762 and F1@10 of 0.2581, indicating that it can retrieve more relevant herbs while maintaining effective ranking quality. These results demonstrate that our method is capable of capturing complex symptom–herb associations and generating more comprehensive prescription recommendations. Compared with existing graph-based approaches, the proposed model better exploits the underlying semantic and structural dependencies among symptoms and herbs, which is particularly important for traditional Chinese medicine prescription recommendation where multiple herbs are jointly selected.

4.4 ABLATION STUDY

Table 2: Ablation study with different components removed.

Dataset	Model	Precision			Recall			F1-score		
		P@5	P@10	P@20	R@5	R@10	R@20	F1@5	F1@10	F1@20
Dataset1	Full Model	**0.2726**	**0.2145**	**0.1577**	**0.2058**	**0.3239**	**0.4762**	**0.2346**	**0.2581**	**0.2369**
	w/o Auxiliary-herb BCE	0.2646	0.2102	0.1543	0.1998	0.3174	0.4660	0.2277	0.2529	0.2318
	w/o Prescription BCE	0.2366	0.1861	0.1386	0.1787	0.2810	0.4186	0.2036	0.2239	0.2082
	w/o Main-herb BCE	0.0764	0.0644	0.1117	0.0577	0.0973	0.3375	0.0657	0.0775	0.1679
	w/o FARL	0.2219	0.1817	0.1401	0.1676	0.2744	0.4231	0.1910	0.2186	0.2105
	w/o latent	0.2578	0.2041	0.1542	0.1946	0.3082	0.4658	0.2218	0.2456	0.2317

To investigate the contribution of each component in our model, we conduct ablation studies by removing different learning objectives and modules. Table 2 shows that the complete model consistently achieves the best performance across all evaluation metrics, demonstrating the effectiveness of the proposed components. First, removing the auxiliary-herb BCE loss leads to a performance degradation, with F1@5 decreasing from 0.2346 to 0.2277. This indicates that auxiliary herbs provide

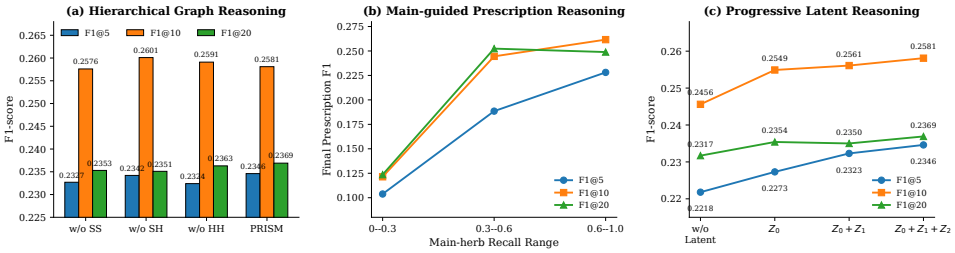

Figure 3: Analysis of progressive prescription reasoning in PRISM. (a) Effects of removing individual SS, SH, and HH reasoning components. (b) Final prescription performance under different main-herb recall ranges. (c) Performance evolution with progressively introduced latent states z_0, z_1, and z_2.

complementary information for prescription completion and help improve the overall recommendation coverage. Second, removing the prescription-level BCE loss causes a more noticeable decline, where F1@5 drops to 0.2036. This demonstrates that global prescription supervision is important for learning comprehensive symptom–herb correlations beyond individual herb roles. Third, removing the main-herb BCE loss results in the largest performance degradation, reducing F1@5 to 0.0651. This confirms that accurately identifying core herbs is crucial for prescription recommendation, and validates the necessity of explicitly modeling the hierarchical roles of herbs. Furthermore, removing FARL also causes consistent performance decreases across all metrics, with F1@10 dropping from 0.2581 to 0.2186. This verifies that the relational information captured by FARL effectively enhances the modeling of complex symptom–herb dependencies. Overall, the ablation results demonstrate that each component contributes from different perspectives, and their combination enables the model to generate accurate and comprehensive TCM prescription recommendations.

4.5 Analysis of Progressive Prescription Reasoning

To further examine whether PRISM behaves as a progressive prescription reasoning framework, we analyze its three key reasoning stages from hierarchical graph modeling, main–auxiliary prescription construction, and latent state evolution. Figure 3 summarizes the results.

Hierarchical graph reasoning. Figure 3(a) evaluates the contribution of SS, SH, and HH reasoning by removing each relation individually. The results show that different relations contribute complementarily across recommendation cutoffs: the complete PRISM achieves the best F1@5 (0.2346) and F1@20 (0.2369), while removing SS, SH, or HH changes the performance to different degrees. This suggests that jointly modeling symptom correlations, symptom–herb associations, and herb compatibility provides complementary structural information for prescription recommendation.

Main–auxiliary prescription reasoning. Figure 3(b) groups test samples according to main-herb recall and examines their final prescription performance. Higher main-herb recall is strongly associated with better performance, particularly at smaller recommendation cutoffs: F1@10 increases from 0.1214 in the low-recall group to 0.2445 and 0.2616 in the medium- and high-recall groups, respectively. This association supports the use of main-herb predictions as informative intermediate guidance for subsequent prescription completion.

Progressive latent prescription reasoning. Figure 3(c) progressively introduces the latent states from no latent reasoning to z_0, $z_0 + z_1$, and $z_0 + z_1 + z_2$. F1@10 improves from 0.2456 to 0.2549, 0.2561, and 0.2581, respectively, while the complete latent trajectory achieves the best F1@5 and F1@20. These results demonstrate that progressively integrating symptom, main-herb, and auxiliary-herb information improves prescription representation and final recommendation.

4.6 Qualitative Analysis

Table 3 provides a qualitative comparison to illustrate the effect of the proposed main-guided auxiliary reasoning. Without the main–auxiliary mechanism, the model recovers only $3/8$ and $4/8$ ground-truth herbs in the two cases, respectively. In contrast, PRISM first identifies a compact set of

Table 3: Comparative case study of prescription recommendation with and without main-guided auxiliary reasoning. Herbs highlighted in red indicate correct predictions that overlap with the ground-truth prescription.

Symptom	w/o Main-Auxiliary	PRISM(w/ Main-Auxiliary)		
		Main Herb	Auxiliary Herb	Full Herb
皮肤甲错 (Scaly skin)	苍耳子 (Xanthium fruit) 郁李 (Bush cherry seed) 厚朴 (Magnolia bark) 龙胆 (Chinese gentian) 青葙子 (Celosia seed) 佩兰 (Eupatorium herb) 虻虫 (Gadfly) 桑螵蛸 (Mantis egg-case)	雷公藤 (Thunder god vine)	苍耳子 (Xanthium fruit) 皂荚 (Chinese honeylocust pod) 雷丸 (Omphalia) 蜗牛 (Snail) 紫草茸 (Lac) 郁李 (Bush cherry seed) 厚朴 (Magnolia bark)	苍耳子 (Xanthium fruit) 皂荚 (Chinese honeylocust pod) 雷公藤 (Thunder god vine) 雷丸 (Omphalia) 蜗牛 (Snail) 紫草茸 (Lac) 郁李 (Bush cherry seed) 厚朴 (Magnolia bark)
Prescription Coverage:	**3/8**	**1/8**	**7/8**	**8/8**
崩中 (Profuse uterine bleeding) 反胃 (Regurgitation)	青葙子 (Celosia seed) 郁李 (Bush cherry seed) 雷公藤 (Thunder god vine) 续断 (Dipsacus root) 龙胆 (Chinese gentian) 皂荚 (Chinese honeylocust pod) 苍耳子 (Xanthium fruit) 桑螵蛸 (Mantis egg-case)	青葙子 (Celosia seed)	续断 (Dipsacus root) 鹿茸 (Deer antler velvet) 雷公藤 (Thunder god vine) 郁李 (Bush cherry seed) 凤仙花 (Balsam flower) 矮地茶 (Japanese ardisia herb) 紫金龙 (Dactylicapnos root)	青葙子 (Celosia seed) 郁李 (Bush cherry seed) 雷公藤 (Thunder god vine) 续断 (Dipsacus root) 凤仙花 (Balsam flower) 矮地茶 (Japanese ardisia herb) 紫金龙 (Dactylicapnos root) 龙胆 (Chinese gentian)
Prescription Coverage:	**4/8**	**1/8**	**6/8**	**7/8**

main herbs and subsequently uses them to guide auxiliary herb inference, substantially improving prescription coverage. Specifically, the auxiliary reasoning stage recovers $7/8$ and $6/8$ ground-truth herbs, and the resulting full prescriptions achieve coverage of $8/8$ and $7/8$, respectively. These cases demonstrate that the main–auxiliary decomposition is not merely a role-wise partition of the output space; rather, it enables progressive prescription completion, where the inferred main herbs provide informative guidance for recovering complementary herbs that are otherwise difficult to identify.

5 Conclusion

In this paper, we presented **PRISM**, a progressive reasoning framework for TCM prescription recommendation that moves beyond direct symptom-to-herb matching. PRISM integrates three complementary processes: hierarchical graph reasoning over symptom–symptom, symptom–herb, and herb–herb relations; adaptive main–auxiliary reasoning that identifies main herbs before main-guided auxiliary completion; and progressive latent prescription states that model the transition from symptom evidence to complete prescriptions. Experiments on the TCM-PD dataset show that PRISM consistently outperforms representative baselines, while ablation and qualitative analyses validate the contributions of its key components. These results demonstrate the effectiveness of explicitly modeling the *process* of prescription formation, providing a structured and interpretable approach to TCM prescription recommendation.

REFERENCES

Zhao-Min Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo. Multi-label image recognition with graph convolutional networks. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5172–5181. IEEE, 2019.

Xin Dong, Chenxi Zhao, Xinpeng Song, Lei Zhang, Yu Liu, Jun Wu, Yiran Xu, Ning Xu, Jialing Liu, Haibin Yu, et al. 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*, 31(6):1268–1279, 2024.

Shijie Geng, Zuohui Fu, Juntao Tan, Yingqiang Ge, Gerard De Melo, and Yongfeng Zhang. Path language modeling over knowledge graphsfor explainable recommendation. In *Proceedings of the ACM web conference 2022*, pp. 946–955, 2022.

Yibo He, Shiyue Wu, Jiayang Li, Shuangyu Chen, Shiliang Chen, Zhezhong Zhang, Beihui He, Yaonan Hong, Chengtao Sun, and Guoyin Kai. Artificial intelligence in traditional chinese medicine: Unraveling herbal medicine's mechanisms. *Research*, 9:1224, 2026.

Jiaxin Hou, Ping Song, Zijuan Zhao, Yan Qiang, Juanjuan Zhao, and Qianqian Yang. Tcm prescription generation via knowledge source guidance network combined with herbal candidate mechanism. *Computational and mathematical methods in medicine*, 2023(1):3301605, 2023.

Hailong Hu, Yaqian Li, Zeyu Zheng, Wenjun Hu, Riyang Lin, and Yanlei Kang. A traditional chinese medicine prescription recommendation model based on contrastive pre-training and hierarchical structure network. *Expert Systems with Applications*, 268:126318, 2025.

Wendi Ji, Ying Zhang, Xiaoling Wang, and Yiping Zhou. Latent semantic diagnosis in traditional chinese medicine. *World Wide Web*, 20(5):1071–1087, 2017.

Yuanyuan Jin, Wei Zhang, Xiangnan He, Xinyu Wang, and Xiaoling Wang. Syndrome-aware herb recommendation with multi-graph convolution network. In *2020 IEEE 36th international conference on data engineering (ICDE)*, pp. 145–156. IEEE, 2020a.

Yuanyuan Jin, Wei Zhang, Xiangnan He, Xinyu Wang, and Xiaoling Wang. Syndrome-aware herb recommendation with multi-graph convolution network. In *2020 IEEE 36th international conference on data engineering (ICDE)*, pp. 145–156. IEEE, 2020b.

Yuanyuan Jin, Wendi Ji, Wei Zhang, Xiangnan He, Xinyu Wang, and Xiaoling Wang. A kg-enhanced multi-graph neural network for attentive herb recommendation. *IEEE/ACM transactions on computational biology and bioinformatics*, 19(5):2560–2571, 2021.

Yuanyuan Jin, Wendi Ji, Yao Shi, Xiaoling Wang, and Xiaochun Yang. Meta-path guided graph attention network for explainable herb recommendation. *Health Information Science and Systems*, 11(1):5, 2023.

Wei Li and Zheng Yang. Exploration on generating traditional chinese medicine prescriptions from symptoms with an end-to-end approach. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pp. 486–498. Springer, 2019.

Yuncheng Li, Yale Song, and Jiebo Luo. Improving pairwise ranking for multi-label image classification. In *2017 IEEE conference on computer vision and pattern recognition (CVPR)*, pp. 1837–1845. IEEE, 2017.

Jingjin Liu, Hankz Hankui Zhuo, Kebing Jin, Jiamin Yuan, Zhimin Yang, and Zhengan Yao. Sequential condition evolved interaction knowledge graph for traditional chinese medicine recommendation. *ICLR*, 2023.

JuanZhi Qi, XinYu Wang, and Tao Yang. Traditional chinese medicine prescription recommendation model based on large language models and graph neural networks. In *2023 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 4623–4627. IEEE, 2023.

Enayat Rajabi and Somayeh Kafaie. Knowledge graphs and explainable ai in healthcare. *Information*, 13(10):459, 2022.

Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. Classifier chains for multi-label classification. *Machine learning*, 85(3):333–359, 2011.

Grigorios Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. Mining multi-label data. *Data mining and knowledge discovery handbook*, pp. 667–685, 2010.

Hongwei Wang, Fuzheng Zhang, Jialin Wang, Miao Zhao, Wenjie Li, Xing Xie, and Minyi Guo. Ripplenet: Propagating user preferences on the knowledge graph for recommender systems. In *Proceedings of the 27th ACM international conference on information and knowledge management*, pp. 417–426, 2018.

Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu. Cnn-rnn: A unified framework for multi-label image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2285–2294, 2016.

Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. Kgat: Knowledge graph attention network for recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 950–958, 2019a.

Xinyu Wang, Ying Zhang, Xiaoling Wang, and Jin Chen. A knowledge graph enhanced topic modeling approach for herb recommendation. In *International conference on database systems for advanced applications*, pp. 709–724. Springer, 2019b.

Xinyu Wang, Xiaohe Sun, Lei Yang, Yitong Zhang, Tao Yang, Jiadong Xie, and Kongfa Hu. Reinforcement learning for llm-based explainable tcm prescription recommendation with implicit preferences from small language models. *Chinese Medicine*, 20(1):193, 2025.

Xiting Wang, Kunpeng Liu, Dongjie Wang, Le Wu, Yanjie Fu, and Xing Xie. Multi-level recommendation reasoning over knowledge graphs with reinforcement learning. In *Proceedings of the ACM web conference 2022*, pp. 2098–2108, 2022.

Zeyuan Wang, Josiah Poon, and Simon Poon. Tcm translator: A sequence generation approach for prescribing herbal medicines. In *2019 IEEE international conference on bioinformatics and biomedicine (BIBM)*, pp. 2474–2480. IEEE, 2019c.

Yikun Xian, Zuohui Fu, Shan Muthukrishnan, Gerard De Melo, and Yongfeng Zhang. Reinforcement knowledge graph reasoning for explainable recommendation. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, pp. 285–294, 2019.

Xiaoyan Yang and Changsong Ding. Smrgat: A traditional chinese herb recommendation model based on a multi-graph residual attention network and semantic knowledge fusion. *Journal of ethnopharmacology*, 315:116693, 2023.

Yun Yang, Yulong Rao, Minghao Yu, and Yan Kang. Multi-layer information fusion based on graph convolutional network for knowledge-driven herb recommendation. *Neural Networks*, 146:1–10, 2022a.

Yun Yang, Yulong Rao, Minghao Yu, and Yan Kang. Multi-layer information fusion based on graph convolutional network for knowledge-driven herb recommendation. *Neural Networks*, 146:1–10, 2022b.

Liang Yao, Yin Zhang, Baogang Wei, Wenjin Zhang, and Zhe Jin. A topic modeling approach for traditional chinese medicine prescriptions. *IEEE Transactions on Knowledge and Data Engineering*, 30(6):1007–1021, 2018.

Chaobo Zhang and Long Tan. A hierarchical structure-enhanced personalized recommendation model for traditional chinese medicine formulas based on kg diffusion guidance. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, pp. 4066–4075, 2025.

Min-Ling Zhang and Zhi-Hua Zhou. A review on multi-label learning algorithms. *IEEE transactions on knowledge and data engineering*, 26(8):1819–1837, 2013.

Yongfeng Zhang and Xu Chen. Explainable recommendation: A survey and new perspectives. *Foundations and trends® in information retrieval*, 14(1):1–101, 2020.

Yongfeng Zhang, Guokun Lai, Min Zhang, Yi Zhang, Yiqun Liu, and Shaoping Ma. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*, pp. 83–92, 2014.

Wen Zhao, Weikai Lu, Zuoyong Li, Chang'en Zhou, Haoyi Fan, Zhaoyang Yang, Xuejuan Lin, and Candong Li. Tcm herbal prescription recommendation model based on multi-graph convolutional network. *Journal of Ethnopharmacology*, 297:115109, 2022.

Wanting Zheng, Yuanyuan Tong, Jinjian Huang, Ling Zhu, and Jiaqi Chai. Beyond transparency: why traditional chinese medicine (tcm) need explainable artificial intelligence (xai). *Chinese Medicine*, 21(1):187, 2026.

A Appendix

A.1 Additional Experiments Information

A.1.1 Random Role Assignment

Table 4: Comparison between adaptive role discovery and random role assignment.

Method	Main Consistency ↑	Auxiliary Consistency	Role Overlap ↓
Random Assignment	0.4971	0.6173	0.0000
Adaptive Discovery	**0.7197**	0.5620	0.0000

To verify that the discovered main/auxiliary roles are not obtained by arbitrary partitioning, we compare our adaptive role discovery strategy with a random role assignment baseline while preserving the same number of main herbs in each prescription.

As shown in Table 4, adaptive role discovery achieves substantially higher main-herb consistency than random assignment (0.7197 vs. 0.4971), demonstrating that the identified main herbs exhibit stronger and more stable symptom-conditioned associations. In contrast, random assignment fails to capture such hierarchical structure despite maintaining identical role sizes.

Although random assignment obtains slightly higher auxiliary consistency, this does not indicate a better auxiliary role discovery. Random assignment incorrectly transfers some stable core herbs into the auxiliary category, artificially increasing auxiliary consistency. In contrast, our adaptive strategy concentrates highly stable symptom-related herbs into the main role, leaving auxiliary herbs as complementary components with naturally higher variability.

These results validate that the proposed specificity-based role discovery provides meaningful prescription hierarchy for subsequent main-guided auxiliary reasoning.

A.1.2 Statistical Validation of Role Discovery.

Table 5: Validation of adaptive main/auxiliary role discovery.

Metric	Main	Auxiliary
Symptom-group Consistency	**0.7197**	0.5620
Mean Frequency	0.0025	0.0058
Role Overlap	0.0000	
Repeated Symptom Groups	1231	

To verify whether the adaptive main/auxiliary decomposition captures meaningful prescription hierarchy, we analyze role consistency and separation. As shown in Table 5, main herbs achieve substantially higher symptom-group consistency than auxiliary herbs (0.7197 vs. 0.5620), indicating

that main herbs exhibit more stable associations with symptom patterns. Moreover, the discovered roles have zero overlap, demonstrating that the main and auxiliary branches capture complementary prescription components. These results suggest that the proposed data-driven role discovery provides a reasonable hierarchical decomposition for subsequent main-guided auxiliary reasoning.

A.1.3 Case Study 1: Latent Prescription Reasoning

Table 6: Case study of latent prescription reasoning. The similarity between latent states reflects the progressive incorporation of main-herb and auxiliary-herb information.

Symptoms	$\mathrm{Cos}(z_0, z_1)$	$\mathrm{Cos}(z_1, z_2)$	$\mathrm{Cos}(z_0, z_2)$
43	0.820	0.835	0.425
377	0.844	0.748	0.333
306,363,367	0.791	0.926	0.575

To further understand the latent reasoning process, we analyze the evolution of three latent prescription states: symptom-conditioned state z_0, main-herb conditioned state z_1, and full-prescription state z_2.

Table 6 reports the cosine similarity between different latent states. The transition from z_0 to z_1 preserves symptom-related therapeutic intent while incorporating main-herb information. Afterwards, z_2 further absorbs auxiliary-herb evidence and represents a more complete prescription state.

For example, in the third case, the similarity between z_1 and z_2 reaches 0.926, indicating that auxiliary reasoning is performed based on a stable main-herb representation rather than independently generating additional herbs.

These observations support the effectiveness of progressive latent prescription reasoning in PRISM.

A.1.4 Case Study 2: HH Graph Compatibility

Herb–Herb Compatibility Reasoning. To qualitatively examine the role of HH reasoning, we trace how main-herb information is propagated through herb–herb compatibility relations during prescription completion. As shown in Table 7, several herbs retrieved from the compatible neighborhoods of the main herbs are consistent with the ground-truth prescription and subsequently contribute to the final prediction. These cases illustrate that the HH graph provides an explicit pathway from identified main herbs to compatible prescription candidates, supporting the role of herb–herb relational reasoning in prescription completion.

A.1.5 Additional Hyperparameter Analysis

We further investigate the sensitivity of PRISM to several key hyperparameters, including the latent consistency weight $\lambda_{\mathrm{latent}}$, ranking loss weight λ_{rank}, the number of hard negatives K, and the role partition parameter c. The results demonstrate that PRISM remains relatively stable over reasonable hyperparameter ranges, while the configurations adopted in the main experiments generally achieve the best overall performance.

Table 8 studies the sensitivity of PRISM to the latent consistency weight $\lambda_{\mathrm{latent}}$. The model exhibits relatively stable performance across a wide range of values, suggesting that the latent reasoning mechanism is not overly sensitive to this hyperparameter. In particular, $\lambda_{\mathrm{latent}} = 0.01$ achieves the best F1@5 and F1@10, while a smaller weight of 10^{-5} yields a marginal improvement on F1@20. Considering the overall performance across different cutoffs, we adopt $\lambda_{\mathrm{latent}} = 0.01$ in all main experiments.

Table 9 evaluates the effect of the ranking loss weight λ_{rank}. PRISM achieves the best performance across all three cutoffs when $\lambda_{\mathrm{rank}} = 0.1$. Reducing the weight substantially degrades performance, indicating that an insufficient ranking signal limits the model's ability to distinguish relevant herbs from hard negatives. Meanwhile, assigning an excessively large weight ($\lambda_{\mathrm{rank}} = 1$) also leads to inferior results, suggesting that the ranking objective should be appropriately balanced with the other training objectives. We therefore set $\lambda_{\mathrm{rank}} = 0.1$ throughout our experiments.

Table 7: Case study of herb–herb compatibility reasoning. The predicted compatible herbs are supported by compatibility relations in the HH graph. Herbs highlighted in red indicate correct predictions that overlap with the ground-truth prescription.

Symptom	Main Herb	HH Compatible Herb	Final Prediction
善疑 (Suspiciousness)	椿皮 (Tree-of-heaven bark) 鸭跖草 (Dayflower) 续断 (Dipsacus root)	鸭跖草 (Dayflower) 金果榄 (Tinospora root) 厚朴 (Magnolia bark) 苍耳子 (Xanthium fruit) 浙贝母 (Fritillaria bulb)	鸭跖草 (Dayflower) 椿皮 (Tree-of-heaven bark) 山慈菇 (Cremastra tuber) 卷柏 (Spikemoss)
耳痒 (Ear itching)	青葙子 (Celosia seed) 马勃 (Puffball) 厚朴 (Magnolia bark)	雷公藤 (Thunder god vine) 郁李 (Bush cherry) 橘红 (Red tangerine peel) 芜菁 (Turnip)	半夏 (Pinellia tuber) 龟甲 (Tortoise shell) 蓖麻子 (Castor bean) 仙鹤草 (Agrimony)
声音嘶哑 (Hoarseness) 口中生疮 (Oral ulcer) 猫眼疮 (Erythema multiforme)	雷公藤 (Thunder god vine) 雷丸 (Omphalia) 郁李 (Bush cherry)	青葙子 (Celosia seed) 厚朴 (Magnolia bark) 续断 (Dipsacus root) 黄芪 (Astragalus root)	降香 (Dalbergia wood) 鼠妇 (Woodlouse) 石楠藤 (Photinia vine) 姜黄 (Turmeric)

Table 8: Effect of the latent consistency weight λ_{latent}. The adopted setting is $\lambda_{\text{latent}} = 0.01$. Best results are shown in bold and second-best results are underlined.

λ_{latent}	F1@5	F1@10	F1@20
1	0.2351	0.2566	0.2361
10^{-1}	0.2316	0.2540	0.2362
$\mathbf{10^{-2}}$	**0.2358**	**0.2604**	<u>0.2368</u>
10^{-3}	0.2338	0.2580	0.2369
10^{-4}	0.2343	0.2585	0.2360
10^{-5}	0.2334	<u>0.2585</u>	**0.2374**

Table 10 investigates the influence of the number of hard negatives K in the ranking objective. Increasing K from 10 to 30 consistently improves performance, indicating that a sufficient number of challenging negative herbs provides a more informative ranking signal. The best overall results are obtained at $K = 30$, whereas further increasing K to 40 or 50 provides no additional benefit. This suggests that excessively enlarging the hard-negative set introduces limited additional supervision once sufficiently informative negatives have been included. Accordingly, we use $K = 30$ as the default setting.

Table 11 examines the robustness of PRISM with respect to the role partition parameter c. The performance remains relatively stable across a broad range of values from 10 to 40, demonstrating that the proposed main–auxiliary role modeling is not highly sensitive to the specific choice of c. Among the tested settings, $c = 25$ achieves the best F1@5 and F1@10, while $c = 30$ yields only a marginally higher F1@20. We therefore adopt $c = 25$ as the default setting due to its strong and balanced performance across different recommendation cutoffs.

Table 9: Effect of the ranking loss weight λ_{rank}. The adopted setting $\lambda_{\text{rank}} = 0.1$ consistently achieves the best performance across all evaluation cutoffs.

λ_{rank}	F1@5	F1@10	F1@20
1	<u>0.2317</u>	<u>0.2447</u>	<u>0.2320</u>
10^{-1}	**0.2358**	**0.2604**	**0.2368**
10^{-2}	0.2060	0.2380	0.2242
10^{-3}	0.1913	0.2228	0.2113
10^{-4}	0.1871	0.2177	0.2085
10^{-5}	0.1870	0.2166	0.2082

Table 10: Effect of the number of hard negatives K used in the role-aware ranking objective. PRISM achieves the best overall performance with $K = 30$.

Hard negatives K	F1@5	F1@10	F1@20
10	0.2280	0.2477	0.2328
20	0.2312	0.2542	0.2355
30	**0.2358**	**0.2604**	**0.2368**
40	<u>0.2340</u>	<u>0.2599</u>	0.2366
50	0.2282	0.2591	<u>0.2368</u>

A.2 Evaluation Metrics

Following KDHR (Yang et al., 2022a) and TCMRGCL (Hu et al., 2025), we adopt **Precision@K**, **Recall@K**, and **F1-score@K** with $K \in \{5, 10, 20\}$ for evaluation.

Let $\mathcal{P}_i^K$ denote the top-K predicted herbs for the i-th prescription and $\mathcal{Y}_i$ denote its ground-truth herb set. The metrics are computed as

$$\text{Precision@K} = \frac{\sum_i |\mathcal{P}_i^K \cap \mathcal{Y}_i|}{\sum_i |\mathcal{P}_i^K|}, \tag{14}$$

$$\text{Recall@K} = \frac{\sum_i |\mathcal{P}_i^K \cap \mathcal{Y}_i|}{\sum_i |\mathcal{Y}_i|}, \tag{15}$$

$$\text{F1@K} = \frac{2 \cdot \text{Precision@K} \cdot \text{Recall@K}}{\text{Precision@K} + \text{Recall@K}}. \tag{16}$$

A.3 Baseline Details

We compare PRISM with representative TCM prescription recommendation methods covering multi-graph representation learning, knowledge-enhanced recommendation, attention-based graph modeling, sequential condition modeling, progressive clinical reasoning, and contrastive graph learning. For all baselines, we follow the implementations and parameter settings reported in their original papers whenever available. The details of the compared methods are summarized as follows.

SMGCN (Jin et al., 2020a) models prescription recommendation by incorporating an implicit syndrome induction process. It constructs symptom–symptom (S–S), symptom–herb (S–H), and herb–herb (H–H) graphs from prescription co-occurrences and applies graph convolution to learn relational representations of symptoms and herbs. The embeddings of symptoms in a prescription are further fused through an MLP to obtain an overall implicit syndrome representation, which is matched with herb representations for recommendation. This method represents an early multi-graph formulation of TCM prescription recommendation and explicitly exploits both symptom co-occurrence and herb compatibility information.

KDHR (Yang et al., 2022a) extends graph-based herb recommendation by introducing external herb knowledge as auxiliary information. In addition to modeling symptom–herb interactions, KDHR

Table 11: Robustness analysis with respect to the role partition parameter c. The adopted setting $c = 25$ achieves the best F1@5 and F1@10, while performance remains stable over a broad range of values.

c	F1@5	F1@10	F1@20
10	<u>0.2348</u>	0.2601	0.2359
15	0.2342	0.2601	0.2362
20	0.2352	<u>0.2602</u>	0.2361
25	**0.2358**	**0.2604**	<u>0.2368</u>
30	0.2328	0.2586	**0.2371**
35	0.2330	0.2577	0.2368
40	0.2338	0.2561	0.2360

constructs a herb knowledge graph containing multiple herbal properties to enrich herb representations. It further introduces a multi-layer information fusion mechanism that combines representations obtained from different GCN layers, allowing the model to preserve information from different neighborhood ranges while reducing the influence of noise introduced by deep graph propagation.

SMRGAT (Yang & Ding, 2023) enhances multi-graph prescription modeling with attention-based propagation and semantic knowledge. It constructs S–S, S–H, and H–H graphs and employs multi-head graph attention with residual connections to capture heterogeneous interactions among symptoms and herbs. In addition, symptom representations are initialized with semantic information, while herbal attributes such as nature, flavor, and meridian tropism are incorporated to enrich herb representations. The representations learned from different graphs are subsequently fused for herb prediction.

SCEIKG (Liu et al., 2023) formulates TCM recommendation as a sequential prescription-making problem. Unlike methods that rely only on the current symptom set, SCEIKG explicitly models the evolution of a patient's condition across consecutive diagnoses and uses the learned condition state to guide subsequent herb recommendations. It additionally constructs an interaction knowledge graph to incorporate TCM domain knowledge and capture high-order associations among symptoms, herbs, and related entities. Therefore, SCEIKG provides a strong sequential and knowledge-enhanced baseline for evaluating prescription reasoning models.

PresRecST (Dong et al., 2024) models herbal prescription recommendation according to the progressive clinical process of syndrome differentiation and treatment planning. Instead of directly mapping symptoms to herbs, it organizes TCM decision-making into sequential stages involving syndrome differentiation, treatment method determination, and herb recommendation. A TCM knowledge graph is used to learn representations of biomedical entities, while residual-like neural networks propagate information across the progressive prediction stages. This design explicitly incorporates intermediate clinical concepts into the prescription recommendation process.

TCMRGCL (Hu et al., 2025) improves symptom and herb representation learning through contrastive pre-training and hierarchical graph modeling. It first constructs augmented graph views and performs multi-view contrastive pre-training to obtain robust initial symptom and herb representations. It then incorporates similarity information into the S–S and H–H graphs to capture fine-grained homogeneous relationships. A hierarchical structure network further performs differentiated neighborhood aggregation according to node characteristics, with particular attention to sparse and less-popular nodes. These components jointly provide richer representations for subsequent prescription recommendation.

A.4 Implementation Details

PRISM is implemented in PyTorch and optimized using Adam. The complete hyperparameter settings used in our experiments are summarized in Table 12.

Table 12: Implementation details of PRISM.

Hyperparameter	**Value**
Embedding dimension	128
Hidden dimension	256
Attention heads	4
Dropout	0.1
Batch size	512
Learning rate	3×10^{-4}
Weight decay	1×10^{-3}
Maximum epochs	200
Early-stopping patience	15
Random seed	2021
Hard negatives	30
Ranking margin	0.1
Main-herb ranking weight	2.0
Auxiliary-herb ranking weight	1.0
Prescription BCE weight	1.0
Main-herb BCE weight	0.1
Auxiliary-herb BCE weight	0.1
Ranking loss weight	0.1
Latent consistency weight	0.01