https://doi.org/10.62836/jcmea.v4i1.040108

Global Science Publishing

JCMEA

Journal of Computational Methods in Engineering Applications

https://ojs.sgsci.org/journals/jcmea

Article

# Self-Reflective Retrieval-Augmented Framework for Reliable Pharmacological Recommendations

**Zhaoyan Zhang [1], Yu Qiao [2] and Peimin Lu [3,*]**

[1] *Zhongke Zhidao (Beijing) Technology Co., Ltd., Beijing 102627, China; zhaoyanzhang9394@gmail.com*

[2] *Meta Platforms, Inc., Bellevue, WA 98005, USA; jakeqiao@meta.com*

[3] *Reckitt Manon, Shanghai 200000, China*

**Abstract:** Pharmacological recommendations are critical for ensuring patient safety and treatment efficacy, yet traditional methods often struggle with inaccuracies and limited adaptability to new knowledge. To address these challenges, this paper proposes a novel *Self-Reflective Retrieval-Augmented Framework* for reliable pharmacological recommendations. The framework incorporates three key innovations: a self-reflective mechanism for dynamic error detection and correction, a pharmacological memory bank for long-term reasoning and knowledge accumulation, and a RAG-enhanced retrieval module to dynamically integrate up-to-date external knowledge during recommendation generation. Experiments on datasets from DrugBank and FDA adverse event reporting systems demonstrate that the proposed framework significantly improves recommendation accuracy, with the full model achieving a 92.3% accuracy and outperforming state-of-the-art methods across multiple evaluation metrics. This research provides a robust and adaptive solution for pharmacological recommendation tasks, paving the way for safer and more effective decision-making in healthcare.

**Keywords:** Retrieval-Augmented Generation; artificial intelligence; pharmacological recommendations; deep learning

## 1. Introduction

Pharmacological recommendations are critical for ensuring patient safety and optimizing treatment outcomes. With the increasing complexity of medical treatments and the proliferation of new drugs, providing accurate and timely drug recommendations has become a significant challenge. Traditional recommendation methods, including rule-based systems and collaborative filtering approaches, often fail to adapt to rapidly evolving pharmacological knowledge and struggle to handle complex drug interactions. Furthermore, errors in pharmacological recommendations, such as contraindicated drug combinations or inappropriate dosages, can lead to severe adverse effects, posing risks to patient safety.

Recent advances in deep learning and natural language processing have enabled significant improvements in recommendation systems, particularly in integrating structured and unstructured data sources. However, pre-trained models frequently rely on static knowledge, making themill-suited for tasks requiring up-to-date medical information or long-term reasoning. Additionally, the lack of mechanisms for dynamic error correction and self-improvement further limits their applicability in sensitive domains such as healthcare.

Received:4 November 2024; Accepted:9 December 2024.

* Corresponding: Peimin Lu (Shane.Lu@reckitt.com)

To address these challenges, this paper introduces a novel *Self-Reflective Retrieval-Augmented Framework* for reliable pharmacological recommendations. The proposed framework is designed to overcome key limitations of existing methods by:

- Incorporating a self-reflective mechanism to dynamically detect and correct errors in generated recommendations, ensuring higher reliability and safety.
- Utilizing a pharmacological memory bank that accumulates validated knowledge and feedback over time, enabling long-term reasoning and improved adaptability.
- Integrating a RAG-enhanced retrieval module to dynamically incorporate up-to-date external knowledge, addressing the limitations of static pre-trained models.

The remainder of this paper is organized as follows. Section 2 provides a detailed explanation of the methodology, including the self-reflective mechanism, pharmacological memory bank, and RAG-enhanced retrieval module. Section 3 describes the experimental setup, including datasets and evaluation metrics, followed by comprehensive ablation studies and comparisons with state-of-the-art methods. Finally, Section 4 concludes with key findings and discusses future research directions.

## 2. Related Work

### *2.1. Self-Reflective Mechanisms and Long-Term Reasoning*

The integration of self-reflective mechanisms and long-term reasoning capabilities has been a significant focus in advancing RAG frameworks. Works like ReAct [1] and Reflexion [2] have highlighted the importance of feedback loops in enhancing the reasoning and acting capabilities of LLMs. By incorporating self-reflection and external observations, these frameworks address challenges such as error propagation and the lack of dynamic adaptability. Similar self-reflective techniques have been explored in Retrieval-Augmented Generation (RAG) models, where iterative refinement can enhance factual accuracy [3–7]. In the context of pharmacological recommendations, self-reflective mechanisms can dynamically detect and correct errors in generated recommendations, ensuring higher reliability and safety.

Additionally, the use of a pharmacological memory bank, as proposed in our framework, enables long-term reasoning by accumulating validated knowledge and feedback over time. Memory-based retrieval and reasoning frameworks have been explored in various medical applications [8–11], where continuous learning from past cases helps in improving recommendation accuracy. By leveraging a structured memory mechanism, our framework avoids repeating past mistakes and adapts to complex or rare cases, enhancing both precision and generalization in drug recommendation.

The integration of search algorithms and reasoning capabilities has been a significant focus in advancing large language models (LLMs). Methods like chain-of-thought (CoT) prompting [12] and Tree-of-Thought (ToT) frameworks [13] have demonstrated the importance of generating structured reasoning paths for solving complex tasks. These approaches align closely with retrieval-augmented frameworks like SR-RAG, where reasoning over external knowledge is critical for improving accuracy and adaptability. Furthermore, structured search processes in retrieval-based reasoning [14] have been shown to refine knowledge retrieval and decision-making, reinforcing the self-reflective mechanisms in SR-RAG.

### *2.2. Retrieval-Augmented Generation and Memory Bank for Pharmacological Recommendations*

The combination of retrieval and generative models in the Retrieval-Augmented Generation (RAG) framework has demonstrated significant potential, particularly in healthcare tasks. The MMED-RAG framework [15] highlights how multi-modal retrieval can enhance factual accuracy and context-awareness by integrating domain-specific retrieval mechanisms. This allows the model to retrieve relevant medical knowledge, improving reasoning capabilities and making it effective for tasks like medical question answering (VQA) and report generation. Similarly, in our pharmacological recommendation framework, retrieval plays a critical role in ensuring that the system utilizes the most up-to-date pharmacological data and clinical knowledge to provide reliable recommendations. This modular retrieval system interacts with a generative model to personalize drug

recommendations based on patient-specific data, improving the quality and safety of the recommendations.

Recent advancements in Retrieval-Augmented Generation (RAG) models have focused on fine-tuning and RAG techniques to improve the factual accuracy of medical models [3,5,16,17]. While fine-tuning is a commonly used method to enhance model performance, it faces challenges in the medical domain due to the scarcity of high-quality labeled data and the distribution gap between training datasets and real-world data [18]. In contrast, RAG methods help address these limitations by incorporating external references during the inference phase, thereby improving the factuality of Medical Large Vision-Language Models (Med-LVLMs) [9, 11]. These external references are particularly beneficial in domains like pharmacology, where the medical knowledge base is vast and continually evolving [11, 14]. However, current RAG implementations still face significant limitations, such as dataset-specific biases and misalignment issues, which can reduce generalizability and lead to factuality problems when applied across different medical domains [19–23].

Moreover, incorporating a Memory Bank for continuous learning, as demonstrated in MMED-RAG, allows models to store and retrieve past experiences, thus enhancing long-term reasoning. Memory driven RAG systems have been explored to reinforce factual consistency and ensure improved retrieval alignment [4, 6, 7]. This reflective learning mechanism enables the model to adapt over time by learning from past mistakes, which is crucial in the context of pharmacological recommendations where drug interactions and treatment protocols evolve. By leveraging a memory bank, our framework can ensure that the system does not propagate past errors and continually refines its decision-making process.

Additionally, modular and domain-specific frameworks have been shown to improve RAG's performance in healthcare. For example, [9] combined different retrieval methods and fine-tuning strategies for medical tasks, while [24] integrated knowledge graphs (KGs) to enhance RAG's performance in medical question answering. These approaches enable the framework to focus on domain-specific knowledge, such as drug interactions and clinical guidelines, improving the accuracy of drug recommendations in our framework. Multi-agent systems, like Clinfo. ai [25], have also demonstrated the power of combining specialized models for comprehensive healthcare solutions, a strategy we similarly adopt to integrate pharmacological and clinical knowledge in our recommendation system.

Furthermore, the importance of cross-modal alignment in RAG models has been emphasized in recent research. The RULE framework [14] addresses over-reliance on retrieved contexts through preference fine-tuning, while other studies [20, 26, 27] investigate alignment issues within medical retrieval tasks. Our framework builds upon these insights by enhancing both retrieval alignment and long-term factual consistency through self-reflective learning and memory augmentation.

## 3. Methodology

### *3.1. Framework Overview*

The workflow of our proposed framework is illustrated in Figure 1. The system integrates a *Self-Reflective Mechanism*, *RAG-Enhanced Retrieval*, and a *Pharmacological Memory Bank* to generate accurate and personalized pharmacological recommendations. The key steps are as follows:

- **Step 1:** The patient information is input into the system and first processed by the *Self-Reflective Mechanism*, which analyzes the patient's medical background and identifies the relevant context for recommendation generation.
- **Step 2:** The *Self-Reflective Mechanism* queries the *RAG-Enhanced Retrieval* module to obtain relevant knowledge from external sources.
- **Step 3:** The *RAG-Enhanced Retrieval* interacts with the *External Knowledge Base* to fetch domain-specific pharmacological information, such as drug interactions, guidelines, and treatment protocols.
- **Step 4:** The retrieved knowledge is sent back to the *RAG-Enhanced Retrieval* module for processing.
- **Step 5:** The retrieved knowledge is then passed to the *Self-Reflective Mechanism*, which validates and integrates this information into the patient-specific recommendation pipeline. This step ensures that retrieved data is correctly aligned with the patient's context before final decision-making.

● **Step 6:** The *Self-Reflective Mechanism* interacts with the *Pharmacological Memory Bank* to store and retrieve historical recommendations. This step allows the model to avoid redundant suggestions and improve long-term decision-making.

● **Step 7:** The *Pharmacological Memory Bank* provides additional insights by retrieving past validated recommendations and their outcomes, further refining the current decision.

● **Step 8:** Finally, the refined recommendation is generated, combining patient-specific data, external medical knowledge, and insights from the memory bank to produce the most reliable and personalized prescription.

Solid arrows in Figure 1 represent the primary data flow. Dashed arrows indicate auxiliary feedback mechanisms, such as the iterative refinement using retrieved knowledge (*Step 5*) and long-term learning via the Pharmacological Memory Bank (*Step 6, Step 7*). This dual-process ensures that recommendations improve over time, leveraging both real-time knowledge and historical insights.

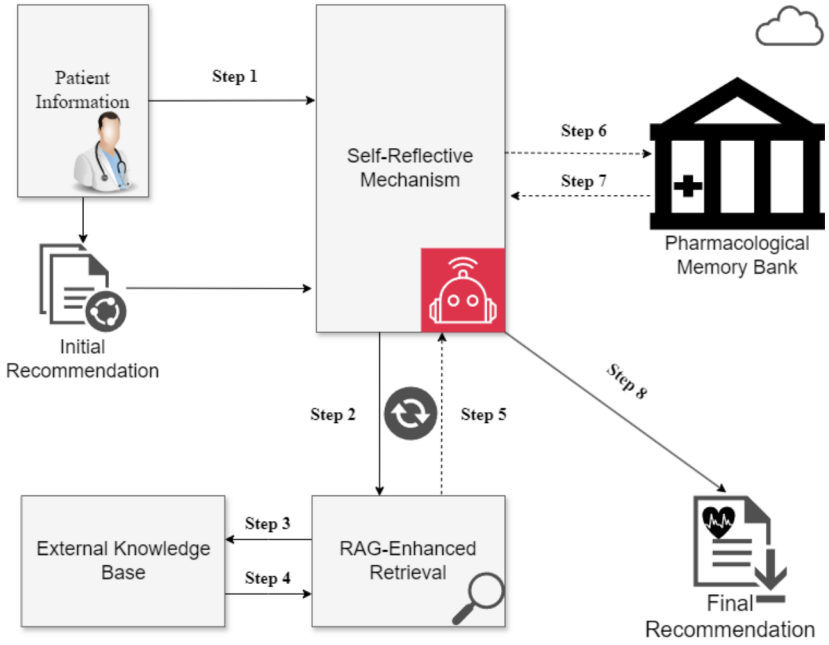


**Figure 1.** Overview of the Self-Reflective Retrieval-Augmented Framework.

### *3.2. Self-Reflective Mechanism for Pharmacological Recommendations*

To ensure the reliability and safety of pharmacological recommendations, we propose a self-reflective mechanism inspired by the human cognitive process of reflection and correction. This module evaluates the generated recommendations for potential errors, such as dosage conflicts or prohibited drug interactions, and provides corrections based on predefined rules and external knowledge sources.

The self-reflective mechanism operates in three phases: detection, correction, and optimization. It dynamically evaluates the generated recommendations and iterativelyrefines them to achieve higher accuracy and reliability.

**Detection Phase:** In the detection phase, the system performs multidimensional validation of the generated pharmacological recommendations. This includes:

$$Validation(x) = \{r_1, r_2, \ldots, r_n\},$$

where x represents the generated recommendation, and ri corresponds to validation rules such as dosage limits, known drug interaction risks, and adverse effects. These rules are retrieved from an external knowledge base K containing structured pharmacological information:

$$K = \{(d_1, d_2, r, p)\},$$

where $d_1$ and $d_2$ are interacting drugs,r is the interaction risk, and p is the probability or severity of the risk.

If a recommendation violates any ruleri , it is flagged for correction, and the system records the error type $E(x)$, such as dosage conflict or interaction risk, for further analysis.

**Correction Phase:** The correction phase involves refining the initial recommendation by incorporating feedback and external knowledge. A secondary module, guided by the error feedback $E(x)$ and external knowledge $K$, refines the recommendation $x$ as follows:

$$x' = Refine(x, E(x), K),$$

where the function Refine combines deterministic rules and adjustments via a neural model to correct the identified errors. Specifically, the model generates a new recommendation y by maximizing the probability of a corrected recommendation, conditioned on the original input $x$, error feedback $E(x)$, and the external knowledge $K$:

$$x' = argmax_y P(y|x, E(x), K),$$

where $P(y|x, E(x), K)$ represents the probability of generating a corrected recommendation y given the input context, error feedback, and external knowledge.

To further enhance the correction process, reinforcement learning (RL) is employed to optimize the generation of recommendations. In this setup, the model receives a reward R that reflects the quality of the recommendation, which is based on multiple factors such as accuracy, safety, and error rate. The reward function is defined as:

$$R = \alpha \cdot Accuracy + \beta \cdot Safety - \gamma \cdot ErrorRate,$$

where α, β, γ are hyperparameters that weight the importance of each evaluation metric. By maximizing this reward function, the model learns to generate recommendations that are both accurate and safe, while minimizing the occurrence of errors. This reinforcement learning approach helps to continually refine the recommendation process, ensuring that the system produces optimal and contextually appropriate recommendations over time.

Thus, the combination of error feedback, external knowledge, and reinforcement learning allows the model to iteratively improve the quality of its recommendations, addressing errors while ensuring that the final recommendation is aligned with the desired objectives of accuracy, safety, and reliability.

**Optimization Phase:** The optimization phase stores the correction process, including the original error and its resolution, in a pharmacological memory bank M for future reference:

$$M = M \cup \{(x, E(x), x')\}$$

This allows the system to dynamically learn from past mistakes and improve its performance in similar scenarios. During subsequent recommendations, the memory bank can be queried to preemptively avoid errors:

$$x' = Generate(x, M, K)$$

The overall process of the self-reflective mechanism is summarized in Algorithm 1.

This self-reflective mechanism enables dynamic detection and correction of errors, significantly improving the reliability and adaptability of pharmacological recommendations.

**Algorithm 1** Self-Reflective Mechanism for Pharmacological Recommendations

**Input:** Generated recommendation $x$, knowledge base $K$
**Output:** Corrected recommendation $x'$
Initialize error list $E(x) \leftarrow \varnothing$
Perform validation: $E(x) \leftarrow$ Validation$(x, K)$
**if** $E(x) \neq \varnothing$ **then**
    Refine recommendation: $x' \leftarrow$ Refine$(x, E(x), K)$
    Update memory bank: $M \leftarrow M \cup \{(x, E(x), x')\}$
**else**
    $x' \leftarrow x$ ▷ No errors detected
**end if**
**return** $x'$

### *3.3. Pharmacological Memory Bank for Long-Term Reasoning*

To enhance the adaptability and reliability of pharmacological recommendations, we propose the construction of a pharmacological memory bank. This memory bank stores historically validated drug

recommendations, interaction risks, and adverse reaction cases to support long-term reasoning and knowledge extension during the generation process. The memory bank is designed to address the limitations of short-term inference in generative models and enables continuous learning through dynamic updates.

The pharmacological memory bank consists of two primary components: static knowledge and dynamic updates. Static knowledge includes known drug interactions, dosage recommendations, and contraindication cases extracted from reliable databases. Dynamic updates involve integrating reflective feedback and corrections generated during recommendation tasks into the memory bank, ensuring continuous refinement of knowledge.

**Data Source:** The data incorporated into the memory bank is extracted from the following sources:

- **Public pharmacological databases:** Data from resources such as DrugBank and FDA adverse event reporting systems, containing drug names, interaction risks, and adverse reaction types.
- **Real-world clinical cases:** Verified drug recommendations and interaction reports from clinical practice.

All data is processed into structured formats, including tuples such as:

$$K = \{(d_1, d_2, r, p)\},$$

where d1 and d2 are interacting drugs, r represents the risk type (e.g., contraindication or adverse effect), and p denotes the severity or probability of the interaction.

**Memory Bank Structure:** The memory bank is structured into two components:

- **Static Knowledge:** Contains pre-verified information, such as drug interaction risks, dosage guidelines, and contraindications, providing a foundational knowledge base.
- **Dynamic Updates:** Reflective feedback from the self-reflective mechanism is dynamically incorporated. For example, if an unreasonable recommendation is corrected during the generation process, the corrected recommendation and the error type are stored in the memory bank as:

$$M_{dynamic} = M_{dynamic} \cup \{(x, E(x), x')\},$$

where x is the original recommendation, E(x) is the error feedback, and x′ is the corrected recommendation.

**Retrieval and Integration Mechanism:** During the generation process, the memory bank is queried to retrieve relevant knowledge based on the current task context. The retrieval process is formalized as:

$$Retrieve(q, K \cup M_{dynamic}),$$

where q represents the query derived from the current recommendation task. The retrieved knowledge is then integrated into the model's reasoning process to guide the generation of recommendations. Specifically:

$$x' = Generate\left(x, Retrieve\left(q, K \cup M_{dynamic}\right)\right),$$

where x is the initial input, and x′ is the final output informed by the memory bank.

The workflow for incorporating the pharmacological memory bank is summarized in Algorithm 2.

The pharmacological memory bank combines static and dynamically updated knowledge to enhance the reliability and adaptability of pharmacological recommendations. By leveraging historically validated knowledge, the model ensures accurate and context-aware recommendations. The dynamic update mechanism enables continuous learning from reflective feedback, helping the model avoid repeating past mistakes. Additionally, the memory bank facilitates adaptation to complex or rare cases by integrating domain-specific corrections into future recommendations. This iterative process allows the model to evolve over time, providing increasingly robust and reliable recommendations across diverse scenarios.

**Algorithm 2** Pharmacological Memory Bank for Long-term Reasoning

**Input:** Current recommendation task $q$, static knowledge $K$, dynamic updates $M_{dynamic}$
**Output:** Recommendation $x'$
Retrieve relevant knowledge: $R \leftarrow \text{Retrieve}(q, K \cup M_{dynamic})$
Generate recommendation: $x' \leftarrow \text{Generate}(q, R)$
Evaluate and update memory: $M_{dynamic} \leftarrow M_{dynamic} \cup \{(x, E(x), x')\}$
**return** $x'$

### *3.4. RAG-Enhanced Retrieval for Knowledge-Augmented Generation*

To address the limitations of pre-trained models relying on outdated knowledge, we integrate a Retrieval-Augmented Generation (RAG) framework intopharmacological recommendation tasks. By dynamically retrieving up-to-date external knowledge, such as newly released drugs and the latest clinical research, the RAG-enhanced approach ensures timely and accurate recommendations.

The integration of RAG bridges the gap between static internal memory and external knowledge, allowing the system to balance stability and adaptability during the recommendation process.

**Knowledge Retrieval Module:** The retrieval module identifies task-relevant information from external knowledge bases, such as drug databases or scientific literature. For each query derived from the recommendation context, the system retrieves a set of the most relevant knowledge entries based on semantic similarity. These entries are then ranked and filtered to ensure that only the most useful information is passed to the generation module.

**Knowledge-Augmented Generation:** The generation module dynamically combines retrieved external knowledge with the internal memory. An attention-based mechanism is employed to balance the contributions of the memory bank and the retrieved knowledge. This ensures that the generated output not only benefits from long-term, validated knowledge but also adapts to the latest medical advancements.

The generation process can be summarized as:

$$x' = Generate\left(q, M, K_{retrieved}\right),$$

where q represents the current query, M denotes the internal memory, and Kretrieved is the external knowledge retrieved for the task.

**Continuous Optimization:** After generating a recommendation, the self-reflective mechanism evaluates the output for potential inaccuracies. If issues such as outdated or irrelevant information are detected, the system adjusts its retrieval strategy or updates the relative weighting of internal and external knowledge. Over time, this iterative feedback process ensures that the model adapts to changing medical contexts and maintains high-quality recommendations.

By integrating dynamically retrieved knowledge, the RAG-enhanced framework significantly improves the timeliness and accuracy of pharmacological recommendations. It addresses the challenges posed by rapidly evolving pharmacological knowledge, ensuring that recommendations remain relevant and contextually accurate. Furthermore, the continuous optimization process allows the system to refine its approach iteratively, balancing stability from internal memory with adaptability from external knowledge sources.

## 4. Experiment

### *4.1. Dataset and Experimental Setup*

To validate the effectiveness of our proposed framework, we conduct experiments using both publicly available datasets and custom-built datasets related to pharmacological recommendations.

**Dataset Description** The primary dataset includes data from DrugBank and FDA adverse event reporting systems. It contains diverse pharmacological information such as drug names, dosages, adverse reactions, and drug interaction cases. To enhance usability, the data is preprocessed by removing redundancy, standardizing formats, and filling missing values.

The dataset provides detailed insights into several aspects of drug usage and interactions:

**Drug Categories:** Drugs are grouped into major therapeutic classes such as antibiotics, antihypertensives, and analgesics, with each category contributing to specific adverse reaction patterns.

**Dosage Ranges:** Includes dosage information for both common and rare medications, with a focus on identifying overdosing risks.

**Adverse Reactions:** Contains detailed reports of mild, moderate, and severe adverse reactions, categorized into physiological systems (e.g., gastrointestinal, cardiovascular, or neurological effects).

**Drug Interactions:** Documents known drug-drug interactions, including high-risk combinations and contraindications. Interaction severity levels are classified as low, moderate, or high.

Figure 2 provides a comprehensive visualization of the dataset distribution, including the proportions of drug categories, the frequency of adverse reactions, and the severity of drug interactions.

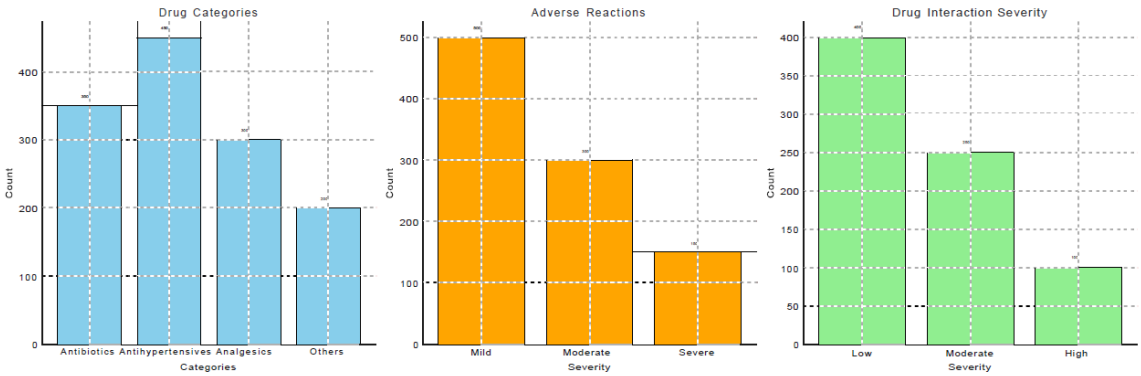


**Figure 2.** Comprehensive distribution of drug categories, adverse reactions, and drug interactions in the dataset.

**Experimental Environment** The experiments are conducted on a system with an NVIDIA Tesla V100 GPU (32 GB), 256 GB RAM, and an Intel Xeon processor. The implementation uses Python 3.9, PyTorch 1.10, and the Hugging Face Transformers library.

**Evaluation Metrics** We evaluate the framework performance using the following metrics: Accuracy (**Acc**) measures the proportion of correct recommendations. F1 Score evaluates the balance between precision and recall. Top-k Hit Rate assesses the likelihood of the correct recommendation appearing within the top-k results. For drug interaction detection, we report precision and recall (Prec/Rec) to capture the detection quality.

### *4.2. Ablation Study*

To demonstrate the effectiveness of each module in our framework, we conduct a comprehensive ablation study. Each module (Self-Reflective Mechanism, Pharmacological Memory Bank, RAG-enhanced Retrieval) is evaluated independently and in combination. Table 1 shows the results.

We evaluate the performance with and without the self-reflective mechanism. Without this module, the system generates unfiltered recommendations, resulting in higher error rates. Introducing the mechanism significantly improves the recommendation quality by detecting and correcting errors dynamically.

To validate the memory bank's impact, we conduct experiments with and without it. Without the memory bank, the model struggles with long-term reasoning tasks and often fails in scenarios requiring knowledge of rare or complex drug interactions. With the memory bank, the performance on these tasks improves substantially.

We test the model's ability to incorporate external knowledge dynamically through the RAG module. Without dynamic retrieval, the model fails to adapt to scenarios involving new drugs or updated clinical research. The RAG-enhanced version demonstrates a higher hit rate and accuracy in these cases.

**Table 1.** Ablation Study Results.

| Configuration | Accuracy (%) | F1 Score | Top-5 Hit Rate (%) | Drug Interaction Prec/Rec |
|---|---|---|---|---|
| Full Model | 92.3 | 0.89 | 94.7 | 0.91/0.88 |
| w/o Self-Reflective Mechanism | 85.6 | 0.82 | 88.1 | 0.85/0.81 |
| w/oPharmacological Memory Bank | 87.2 | 0.84 | 89.5 | 0.86/0.83 |
| w/o RAG-enhanced Retrieval | 88.4 | 0.85 | 90.8 | 0.87/0.84 |

### *4.3. Comparison with State-of-the-Art Methods*

In this section, we compare our proposed RAG framework with a range of state-of-the-art methods,

including traditional recommendation systems (rule-based systems and collaborative filtering), deep learning-based approaches (such as transformer models), and existing Retrieval-Augmented Generation (RAG) methods in the medical and pharmaceutical domains. Table 2 summarizes the quantitative results, demonstrating that our framework outperforms the traditional methods and the most recent RAG-based approaches in several key evaluation metrics.

We benchmark our system against the following categories:

- **Traditional Methods**: Rule-based systems and collaborative filtering models, which are widely used in recommendation systems but often lack the ability to incorporate external knowledge dynamically.
- **Deep Learning-based Models**: Transformer-based recommendation systems that utilize neural networks for learning complex patterns but may still struggle with domain-specific external knowledge integration.
- **Medical and Pharmaceutical RAG Methods**: Recent RAG frameworks applied in healthcare and drug recommendation tasks, such as those outlined in Section 2.2, which focus on integrating retrieval with generative models.

By comparing across these diverse approaches, we show that our RAG framework not only outperforms traditional methods but also provides superior performance over recent RAG-based medical recommendation systems. This demonstrates the effectiveness of incorporating a robust retrieval mechanism, error feedback, and continuous learning in enhancing the accuracy and reliability of pharmacological recommendations.

**Table 2.** Comparison with State-of-the-Art Methods.

| Method | Accuracy (%) | F1 Score | Top-5 Hit Rate (%) | Drug Interaction Prec/Rec |
|---|---|---|---|---|
| Rule-based System | 78.5 | 0.75 | 81.2 | 0.78/0.72 |
| Collaborative Filtering | 83.4 | 0.79 | 85.7 | 0.81/0.77 |
| Deep Learning Model (Transformer) | 88.1 | 0.85 | 91.4 | 0.87/0.83 |
| MMED-RAG [15] | 90.2 | 0.87 | 93.1 | 0.89/0.85 |
| Proposed Framework | **92.3** | **0.89** | **94.7** | **0.91/0.88** |

In conclusion, the results in Table 2 show that our proposed framework not only provides superior performance compared to traditional recommendation methods but also outperforms existing RAG-based solutions in the medical and pharmaceutical fields, reinforcing the effectiveness of our approach in generating reliable and accurate drug recommendations.

### *4.4. Case Study*

In this case study, we demonstrate how the proposed Self-Reflective Retrieval-Augmented Framework improves pharmacological recommendations (see Figure 3). The example involves a 65-year-old patient with hypertension currently prescribed Amlodipine (DrugA), a calcium channel blocker.

**Input Data**

- **Patient Information:**

- Age: 65
- Condition: Hypertension (High Blood Pressure)
- Current Medication: Amlodipine (DrugA)–calcium channel blocker

**Initial Recommendation:**

- **Recommended Medication:** Amlodipine (DrugA) + Lisinopril (Drug B)–ACE inhibitor
- **Reason:** Based on clinical guidelines, the combination of Amlodipine (acalcium channel blocker) and Lisinopril (an ACE inhibitor) is commonly used to manage hypertension in elderly patients.

Error Detected (During Self-Reflective Mechanism)

- **Detected Issue:** The combination of Amlodipine and Lisinopril increases the risk of hyperkalemia (elevated potassium levels) in elderly patients, which could lead to heart arrhythmias or kidney problems.
- **Problem:** This drug combination could potentially cause dangerous side effects, especially in elderly patients with pre-existing kidney conditions.

Correction by Self-Reflective Mechanism:

- The system detects this risk and recommends substituting Lisinopril with Losartan (a different class of ACE inhibitor that poses a lower risk of hyperkalemia).

**External Knowledge Retrieval (RAG-Enhanced Retrieval)**

- **New Knowledge:** Recent clinical studies and guidelines suggest that Losartan is equally effective in managing hypertension and poses a lower risk of hyperkalemia compared to Lisinopril, especially for elderly patients.

**Output Recommendation**

- **Recommended Medication:** Amlodipine (DrugA) + Losartan (Drug C)
- **Reason:** After correcting the initial recommendation, Losartan is chosen to avoid the risk of hyperkalemia while maintaining effective blood pressure control. The combination with Amlodipine is safe and effective for elderly patients.

This case study highlights the effectiveness of our framework in dynamically correcting drug combinations to improve safety and efficacy.

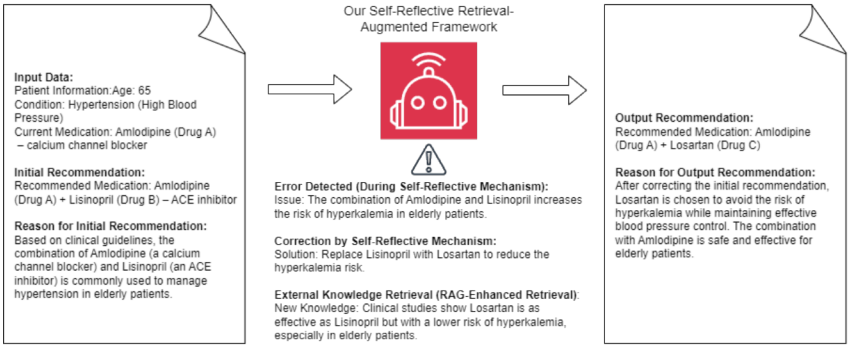


**Figure 3.** Improving Pharmacological Recommendation for Hypertension.

## 5. Conclusions and Future Work

In this paper, we proposed a novel Self-Reflective Retrieval-Augmented Framework for reliable pharmacological recommendations. By integrating a self-reflective mechanism for error detection and correction, a pharmacological memory bank for long-term reasoning and knowledge accumulation, and a RAG-enhanced retrieval module for incorporating up-to-date external knowledge, the framework addresses critical challenges in pharmacological recommendation systems. The proposed framework demonstrated superior performance compared to state-of-the-art methods, achieving a 92.3% accuracy and consistently outperforming baselines across multiple evaluation metrics. These results validate the efficacy of the framework in improving recommendation accuracy, safety, and adaptability.

However, there are still limitations in our approach. First, the reliance on external knowledge bases, while beneficial for adaptability, may introduce latency issues when retrieving large-scale data. Second, the pharmacological memory bank, though effective for long-term reasoning, may require optimization to handle the scalability of rapidly growing datasets. Third, the framework has not yet been extensively validated in real-world clinical environments, which may present unforeseen challenges.

Future work will focus on addressing these limitations. We plan to optimize the retrieval module for faster and more efficient integration of external knowledge while maintaining accuracy. Additionally, we aim to enhance the scalability of the pharmacological memory bank by employing advanced indexing and compression techniques. Finally, we will collaborate with healthcare professionals to test the framework in real-world clinical

settings, incorporating feedback to further refine its performance and usability. This future research will contribute to advancing safe and reliable pharmacological decision-making in healthcare.

**Funding**

This research has been partly funded by the National Natural Science Foundation of China (NSFC) through awards 61872364 and 61972221. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the NSFC.

**Author Contributions**

Writing—original draft preparation and writing—review and editing, Z.Z., Y.Q., P.L. All authors have read and agreed to the published version of the manuscript.

**Institutional Review Board Statement**

Not applicable.

**Informed Consent Statement**

Not applicable.

**Data Availability Statement**

Not applicable.

**Conflicts of Interest**

The authors declare no conflict of interest.

**References**

1 Yao S, Yu Z, Du N, *et al*. React: Synergizing Reasoning and Acting in Language Models. *arXiv* 2023; arXiv: 2210.03629.

2 Shinn N, Labash J, Savinov N. Reflexion: Language Agents with Verbal Reinforcement Learning. *arXiv* 2023; arXiv:2303.11366.

3 He S, Nie Y, Chen Z, *et al*. Meddr: Diagnosis-Guided Bootstrapping for Large-Scale Medical Vision-Language Learning. *arXiv* 2024; arXiv:2404.15127.

4 Kumar Y, Marttinen P. Improving Medical Multi-Modal Contrastive Learning with Expert Annotations. *arXiv* 2024; arXiv:2403.10153.

5 Sun L, Zhao J, Han M, *et al*. Fact-Aware Multimodal Retrieval Augmentation for Accurate Medical Radiology Report Generation. *arXiv* 2024; arXiv:2407.15268.

6 Tao Y, Ma L, Yu J, *et al*. Memory-Based Cross-Modal Semantic Alignment Network for Radiology Report Generation. *IEEE Journal of Biomedical and Health Informatics* 2024; **28**: 4145–4156.

7 Yuan Z, Jin Q, Tan C, *et al.* Ramm: Retrieval-Augmented Biomedical Visual Question Answering with Multi-Modal Pre-Training. In Proceedings of the 31st ACM International Conference on Multimedia, Ottawa, ON, Canada, 29 October–3 November 2023; pp. 547–556.

8 Chen Z, Xu C, Qi Y, *et al*. Mllm Is a Strong Reranker: Advancing Multimodal Retrieval-Augmented Generation via Knowledge-Enhanced Reranking and Noise-Injected Training. *arXiv* 2024; arXiv:2407.21439.

9 Gao Y, Xiong Y, Gao X, *et al.* Retrieval-Augmented Generation for Large Language Models: A Survey. *arXiv* 2023;arXiv:2312.10997.

10 Qu X, Chen Q, Wei W, *et al*. Alleviating Hallucination in Large Vision-Language Models with Active Retrieval Augmentation. *arXiv* 2024; arXiv:2408.00555.

11 Wu Y, Zhu J, Xu S, *et al.* Ragtruth: A Hallucination Corpus for Developing Trustworthy Retrieval-Augmented Language Models. *arXiv* 2023; arXiv:2401.00396.

12 Wei J, Wang X, Schuurmans D, *et al*. Chain of thought Prompting Elicits Reasoning in Large Language Models. *arXiv* 2022; arXiv:2201.11903.

13 Yao S, Zhao J, Yu Z, *et al.* Tree of thoughts: Deliberate Problem Solving with Large Language Models. *arXiv* 2023; arXiv:2305.10601.

14 Xia P, Zhu K, Li H, *et al.* Rule: Reliable Multimodal Rag for Factuality in Medical Vision Language Models. *arXiv* 2024; arXiv:2407.05131.

15 Xia P, Zhu K, Li H, *et al*. Mmed-rag: Versatile Multimodal Rag System for Medical Vision Language Models. *arXiv* 2024; arXiv:2410.13085.

16 Li C, Wong C, Zhang S, *et al*. Llava-Med: Training a Large Language-and-Vision Assistant for Biomedicine in One Day. *Advances in Neural Information Processing Systems* 2023; **36**: 28541–28564.

17 Moor M, Huang Q, Wu S, *et al*. Med-Flamingo: A Multimodal Medical Few-Shot Learner. *PMLR* 2023; **225**: 353–367.

18 Schrouff J, Harris N, Koyejo S, *et al*. Diagnosing failures of fairness transfer across distribution shift in real-world medical settings. *Advances in Neural Information Processing Systems* 2022; **35**: 19304–19318.

19 Chen Z, Du Y, Wen Z, *et al*. MJ-Bench: Is Your Multimodal Reward Model Really a Good Judge for Text-to-Image Generation? *arXiv* 2024; arXiv:2407.04842.

20 Liu H, Li C, Li Y, *et al*. Improved Baselines with Visual Instruction Tuning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 26296–26306.

21 Liu H, Li C, Wu Q, *et al*. Visual Instruction Tuning. *Advances in Neural Information Processing Systems* 2023; **36**: 34892–34916.

22 Xia P, Yu X, Hu M, *et al*. Hgclip: Exploring Vision-Language Models with Graph Representations for Hierarchical Understanding. *arXiv* 2023; arXiv:2311.14064.

23 Zhou Y, Cui C, Yoon J, *et al.* Analyzing and Mitigating Object Hallucination in Large Vision-Language Models. *arXiv* 2023; arXiv:2310.00754.

24 Jiang X, Zhang R, Xu Y, *et al*. Think and Retrieval: A Hypothesis Knowledge Graph Enhanced Medical Large Language Models. *arXiv* 2023; arXiv:2312.15883.

25 Lozano A, Fleming SL, Chiang CC, *et al*. Clinfo. ai: An Open-Source Retrieval-Augmented Large Language Model System for Answering Medical Questions Using Scientific Literature. In P*acific Symposium on Biocomputing, 2024*; World Scientific: Singapore, 2023; pp. 8–23.

26 Royer C, Menze B, Sekuboyina A. Multimedeval: A benchmark and a toolkit for evaluating medical vision-language models. *arXiv* 2024; arXiv:2402.09262.

27 Xia P, Chen Z, Tian J, *et al*. Cares: A Comprehensive Benchmark of Trustworthiness in Medical Vision Language Models. *arXiv* 2024; arXiv:2406.06007.