

SeRTS: Self-Rewarding Tree Search for Biomedical Retrieval-Augmented Generation

Minda Hu^{1*} Licheng Zong^{1*} Hongru Wang² Jingyan Zhou² Jingjing Li¹
Yichen Gao¹ Kam-Fai Wong² Yu Li¹ Irwin King¹

¹Department of Computer Science and Engineering

²Department of Systems Engineering and Engineering Management

The Chinese University of Hong Kong

{mindahu21, lczong21, liyu, king}@cse.cuhk.edu.hk

Abstract

Large Language Models (LLMs) have shown great potential in the biomedical domain with the advancement of retrieval-augmented generation (RAG). However, existing retrieval-augmented approaches face challenges in addressing diverse queries and documents, particularly for medical knowledge queries, resulting in sub-optimal performance. To address these limitations, we propose a novel plug-and-play LLM-based retrieval method called **Self-Rewarding Tree Search (SeRTS)** based on Monte Carlo Tree Search (MCTS) and a self-rewarding paradigm. By combining the reasoning capabilities of LLMs with the effectiveness of tree search, SeRTS boosts the zero-shot performance of retrieving high-quality and informative results for RAG. We further enhance retrieval performance by fine-tuning LLMs with Proximal Policy Optimization (PPO) objectives using the trajectories collected by SeRTS as feedback. Controlled experiments using the BioASQ-QA dataset with GPT-3.5-Turbo and Llama2-7b demonstrate that our method significantly improves the performance of the BM25 retriever and surpasses the strong baseline of self-reflection in both efficiency and scalability. Moreover, SeRTS generates higher-quality feedback for PPO training than self-reflection. Our proposed method effectively adapts LLMs to document retrieval tasks, enhancing their ability to retrieve highly relevant documents for RAG in the context of medical knowledge queries. This work presents a significant step forward in leveraging LLMs for accurate and comprehensive biomedical question answering.

1 Introduction

Large Language Models (LLMs) have made significant advancements and gained widespread adoption in various natural language tasks, such as dialogue systems (Bang et al., 2023; Wang et al., 2023) and question-answering (Kamalloo et al.,

2023; Li et al., 2023; Wang et al., 2024b). However, their performance in specialized domains like medicine remains unsatisfactory due to the lack of publicly available training data, which limits the amount of domain-specific knowledge these LLMs can encode in their parameters (Xia et al., 2022; Shi et al., 2024). To address this limitation, recent studies have turned to retrieval-augmented generation (RAG) frameworks (Xiong et al., 2024; Wang et al., 2024a; Gao et al., 2024a). These frameworks can enhance LLMs' capabilities by retrieving relevant documents from large and private databases, thereby enabling the models to provide more accurate and informative answers.

Existing RAG framework mainly consists of two modules: an external retriever to retrieve documents according to semantic similarity between documents and queries, and the reader to generate the final answer given retrieved documents and query (Gao et al., 2024b). However, this kind of setting overlook the two significant challenges in medical domain: the complexity of the queries and the varying importance of different documents. For example, patients may describe the same disease in diverse ways, leading to variations in query formulation. Additionally, documents from different medical disciplines should be weighted differently based on both of their relevance and importance. Thus, effectively identifying the most pertinent and useful documents from a vast corpus of medical literature is crucial. This not only enhances performance but also improves the explainability and credibility of the system.

To address these issues, we propose a plug-and-play LLM-based retrieval method, **SELF-REWARDING TREE SEARCH (SERTS)**, combines the exceptional reasoning capabilities of LLMs with the effective searching capabilities of Monte Carlo Tree Search (MCTS). Specifically, we first formulate the task of document retrieval as a tree search problem, considering both the diversity of

*The first two authors have equal contributions.

queries and different relationship between multiple documents. Inspired by recent studies using MCTS to search the optimal reasoning paths of LLMs, we alternatively employ it to find the best query and most related documents from vast candidate spaces. To ensure the success of MCTS, we carefully design each operations in MCTS and provide detailed feedback from different levels of documents. Furthermore, to validate the effectiveness and quality of our tree search retriever, we then collect these feedbacks with all intermediate reasoning results to finetune LLMs using PPO method. To sum up, our contributions can be summarized as follows:

- We present a novel perspective on automating query optimization by combining LLM and MCTS for improved document retrieval in the medical domain.
- We propose a Self-Rewarding Tree Search framework, consisting of a tree search retriever and a PPO self-training process.
- Experimental results on two LLMs, GPT-3.5-TURBO and LLAMA2-7B-CHAT, demonstrate that our SeRTS method significantly improves the performance of BM25 retrievers and surpasses the strong baseline of self-reflection in both efficiency and scalability.

2 Related Works

Medical RAG. Retrieval-Augmented Generation represents a significant advancement to provide up-to-date and trustworthy collections and improves the transparency of LLMs by grounding their reasoning on the retrieved documents (Xiong et al., 2024), particularly in the medical question-answering (Tian et al., 2019; Xia et al., 2022) and dialogue tasks (Shi et al., 2024). Most of previous works focus on query rewriting (Zhang et al., 2022) and multi-step retrieval (Mrini et al., 2022). More recent studies explore how LLMs can improve literature information-seeking and clinical decision-making RAG (Zakka et al., 2023; Wang et al., 2024c; Jeong et al., 2024). In detail, Wang et al. (2024c) design a hybrid retriever and conduct complex filter operations to get high-quality documents while Jeong et al. (2024) propose Self-BioRAG to utilize retrieval and self-reflection together, resulting in better medical reasoning.

Applications of MCTS. Monte Carlo Tree Search (MCTS) has been validated as an effective

method for finding better solution paths among a wide range of potentially feasible options, enabling proficient exploration and enhancing decision-making (Silver et al., 2017). It has proven to be highly successful in various domains, including game-playing (Sironi et al., 2018; Ontanón, 2016), and planning (Zhou et al., 2023; Yu et al., 2023). By combining MCTS with reinforcement learning, models can learn from self-play and achieve human-level performance or even surpass it in complex tasks, such as the game of Go (Silver et al., 2016). Recent works mainly utilize MCTS to determine the most promising reasoning path of LLMs, leading to better performance without additional annotations (Feng et al., 2024; Zhou et al., 2024; Tian et al., 2024). Specifically, Feng et al. (2024) fix a search step to be either a token or a sentence while Tian et al. (2024) employ η MCTS to search across multiple levels of planning with carefully defined reward signals. Additionally, other studies use MCTS to collect high-quality reasoning traces to train the reward model, aiming to continuously enhance LLMs over multiple iterations (Zhang et al., 2024). Furthermore, Zhou et al. (2024) expands the search space to encompass all possible reasoning and acting steps, thereby unifying reasoning, acting, and planning in language models.

Learning from Feedback. LLMs can evolve and then refine its outputs by learning from self-assessed rewards or environmental signals, such as self-learning or reinforcement learning from human feedback (RLHF). On the one hand, self-learning methods leverage the LM’s ability to critique its outputs, collecting natural language feedback as signals, such as Reflection (Shinn et al., 2024), Self-Rewarding LMs (Yuan et al., 2024a). On the other hand, lots of works apply reinforcement learning to train the external reward model from human preferences (Ouyang et al., 2022). The reward model is then frozen and used to train the LLM using RL, e.g., Proximal Policy Optimization (PPO) (Schulman et al., 2017; Wang et al., 2022), which is a popular reinforcement learning algorithm that has been successfully applied to various domains, including natural language processing tasks (Han et al., 2023; Akyurek et al., 2023). However, most existing reinforced self-improvement approaches, e.g., STaR (Zelikman et al., 2022) and RFT (Yuan et al., 2023), overlooking the possibility that incorrect or low-quality reasoning traces may still lead to the correct solution due to hallucinations or

randomness, leading to an unreliable reward model (Lanham et al., 2023).

3 Method

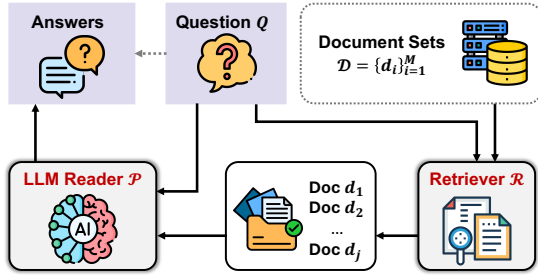


Figure 1: A straightforward Retrieval-Augmented Generation (RAG) pipeline

Figure 1 illustrates a straightforward Retrieval-Augmented Generation (RAG) pipeline we employed as the basis of our experiments. This process can be defined as follows: Give a literature document set $\mathcal{D} = \{d_i\}_{i=1}^M$, a Language Model based reader \mathcal{P} , an external retriever \mathcal{R} , and a question Q , use \mathcal{R} to retrieve several documents $\{d_{1..j}\}$ related to Q , so that the reader \mathcal{P} can generate appropriate answers using them. When \mathcal{P} is fixed, the quality of the generated answers mainly depends on the retrieved $\{d_{1..j}\}$.

To improve the relevance of $\{d_{1..j}\}$ to Q for higher quality of the answers, we propose a novel method for retrieval, called **Self-Rewarding Tree Search (SeRTS)** as shown in Figure 2 illustrating the Expansion & Evaluation process in the MCTS and its interaction with two LLM agents. The core of SeRTS is a search algorithm based on Monte Carlo Tree Search (MCTS) that learns from trial-and-error experiences based on Self-Rewarding Paradigms and explores better trajectories of self-reflection reasoning to retrieve more relevant documents for Retrieval Augmented Generation (RAG).

It is worth noting that our method is decoupled from the RAG pipeline. All components in the pipeline, including the document set $\mathcal{D} = \{d_i\}_{i=1}^M$, the LM reader \mathcal{P} and the retriever \mathcal{R} could be replaced by other types. Our method acts as a plugger that can be integrated into any RAG pipelines.

3.1 Self-Rewarding Tree Search

As a decision policy, MCTS effectively balances exploration and exploitation when searching for optimal trajectories, so it’s employed to find the best retrieval results in our proposed method.

We formulate the document retrieval process as a tree search problem. In the tree shown in Figure 2, each node $s_i = [Q, A_i, D_i, R_i, F_i, O_i]$ represents a state with the original question Q , query proposal A_i , retrieved documents F , result score R and feedback F , and previous process observation O_i (query proposals from all s_i ’s ancestor nodes $O_q^{anc} = \bigcup_{k \in \text{anc}(s_i)} A_k$, retrieved documents from all s_i ’s ancestor nodes $O_d^{anc} = \bigcup_{k \in \text{anc}(s_i)} D_k$, query proposals from all s_i ’s sibling nodes $O_q^{sib} = \bigcup_{k \in \text{sib}(s_i)} A_k$, feedback from all s_i ’s sibling nodes $O_f^{sib} = \bigcup_{k \in \text{sib}(s_i)} F_k$).

Starting from the root node (initial state s_0), SeRTS uses the MCTS algorithm that iteratively searches through the Query Space $A_{1..i}$ and uses $A_{1..i}$ to retrieve documents. The retriever employed here is the BM25, a strong retriever in the biomedical retrieval domain, supported by multiple evaluations (Xiong et al., 2024). The final target is to find node s_o that has retrieved documents from itself and its ancestors are evaluated to have the highest scores in helping generate answers for question Q .

3.1.1 Result Evaluator for Quantitatively Measuring

Quantitatively measuring the quality of retrieved documents by providing high-quality feedback with a specific score is essential for the MCTS. Here, we propose to utilize LMs to automatically do the evaluation and provide feedback, inspired by the Self-Rewarding paradigm (Yuan et al., 2024a). The quality of feedback plays a crucial role in producing accurate reasoning, and low-quality feedback can hinder the self-improving ability of LLMs during self-reflection.

We design a **Result Evaluator** P_θ^{eval} using the system prompt in Table 1 to provide descriptive feedback and quantitative measurement for retrieved documents.

It utilizes the 5-point rubrics to assess the relevance of retrieved documents to the query and their potential contribution to the reasoning process, as well as providing descriptive feedback to the search algorithm. This scoring mechanism enables MCTS to make informed decisions during the search process, guiding it toward retrieving the most informative and relevant documents for RAG.

3.1.2 Query Proposer for Efficiently Searching

The vast number of potential token combinations in natural language tasks results in an exponentially

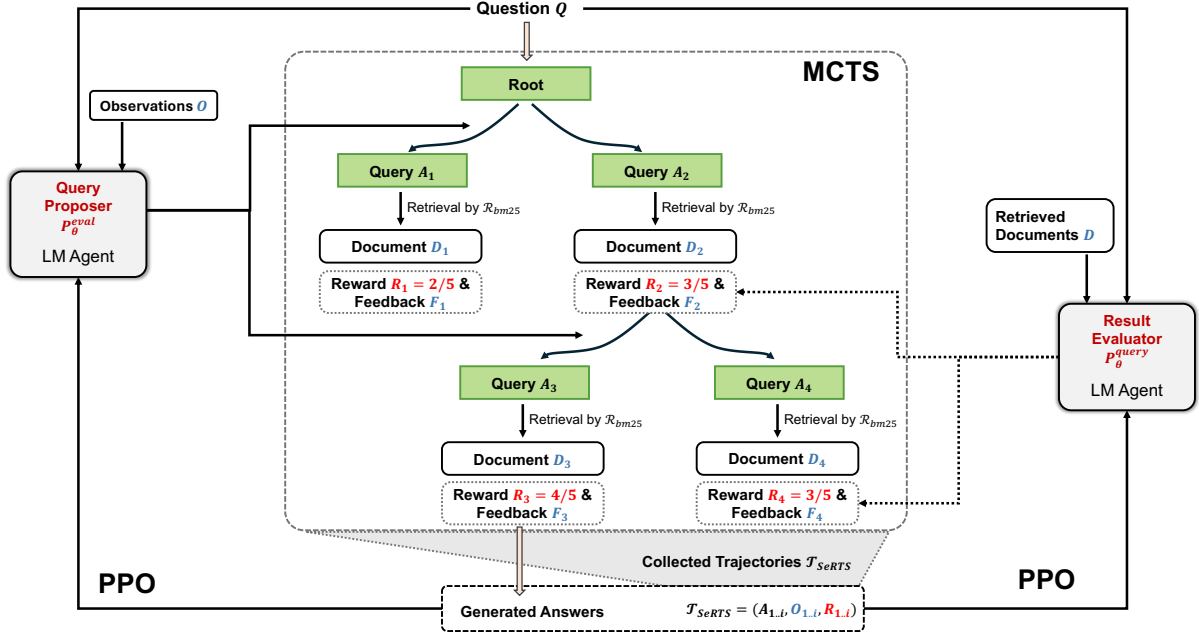


Figure 2: SeRTS method overview: **Query Proposer** P_{θ}^{query} generates query A_i . \mathcal{R}_{bm25} retrieves relevant documents D_i . **Result Evaluator** P_{θ}^{eval} assesses D_i , provides Reward R_i and Feedback F_i . Observations O_i (previous queries, retrieved documents, feedback) serve as input to P_{θ}^{query} for subsequent query proposals. The entire trajectory \mathcal{T}_{SeRTS} (initial question Q , $A_{1..i}$, $O_{1..i}$, $R_{1..i}$) is used to fine-tune the two language model agents via PPO, to improve their performance in query proposal and evaluation.

You are a medical expert. Review the user’s question and rate the corresponding retrieved knowledge using the additive 5-point scoring system described below. Points are accumulated based on the satisfaction of each criterion: **{Five-point Rubrics}**

```
<question> {Q} </question>
<retrieved_knowledge>
{D ∪ O_d^{anc}}
</retrieved_knowledge>
```

After examining the user’s instruction and the response:
- Briefly justify your total score, up to 100 words. Remember to assess from the AI Assistant perspective, utilizing web search knowledge as necessary. To evaluate the response in alignment with this additive scoring model, we’ll systematically attribute points based on the outlined criteria.
- Give suggestions for constructing better queries up to 100 words.

Conclude with a score between 0 and 5, strictly using the aforementioned additive 5-point scoring system and the format: “<score> Integer Score </score>”. For example, <score>3</score>.

Table 1: Prompt for Evaluator P_{θ}^{eval} . Details of **{Five-point Rubrics}** is shown in Appendix A.

large search space, posing a significant challenge to the efficiency of MCTS. Our approach tackles the search space issue by employing strategic expanding techniques and heuristics to navigate the

vast query space $A_{1..i}$ efficiently.

We design a **Query Proposer** P_{θ}^{query} using the system prompt in Table 2. Query Proposer efficiently explores new queries in a heuristic manner using previous results and feedback. At every node, P_{θ}^{query} is given the query proposals from its sibling nodes and the corresponding feedback produced by the Result Evaluator P_{θ}^{eval} , as well as previously retrieved documents from all its ancestor nodes. Considering the known information, P_{θ}^{query} is prompted to generate new query proposals efficiently for the next step.

3.1.3 Monte Carlo Tree Search Process

As shown in Figure 3, using the **Query Proposer** P_{θ}^{query} and the **Result Evaluator** P_{θ}^{eval} , SeRTS iteratively performs *Selection*, *Expansion*, *Evaluation*, and *Backpropagation* until it finds a good enough query result or a computational limit is reached. The full pseudocode is provided in Appendix D.2.

Selection. In the first operation, the algorithm selects the most promising leaf node for expansion, starting from the root node (initial state s_0). The Upper Confidence Bounds (UCB) is used to balance exploration and exploitation during node

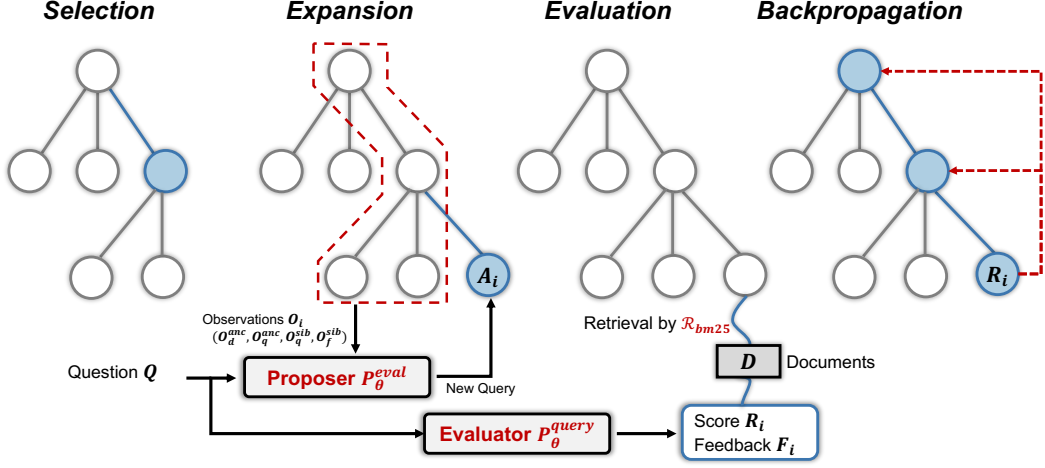


Figure 3: Overview of the four operations in MCTS

You are a medical expert, and you are tasked with searching for as much useful knowledge from a database of medical paper abstracts as possible to address the user’s question. The question is stated as follows: $\{Q\}$

You are given the BM25 retriever as the search tool. According to previous query proposals and proposals’ corresponding feedback and suggestions, give one single better query for the retriever to further search for extra information needed for answering the given question.

<query_proposal>

1. $\{O_{q,1}^{sib}\}$ Corresponding Feedback: $\{O_{f,1}^{sib}\}$
 -
 - i. $\{O_{q,i}^{sib}\}$ Corresponding Feedback: $\{O_{f,i}^{sib}\}$
- </query_proposal>

Information already known is as follows:<retrieved_info>
 Query and its reasoning for Abstracts: $\{O_q^{anc}\}$
 Retrieved Abstracts: $\{O_d^{anc}\}$
 </retrieved_info>

After examining the information above, justify your decision, in up to 100 words.

Hints:

1. Use simple keywords as queries if no information is retrieved.
2. You can take one step at a time, subdivide the original question, and search for information on the sub-questions.

Conclude with the query for the paper abstracts using the format: "<query> Query Here </query>".

Table 2: Prompt for Query Proposer P_θ^{query} .

selection:

$$UCB(s) = \frac{R(s)}{N(s)} + C * \sqrt{2 * \ln \frac{N(p)}{N(s)}} \quad (1)$$

, where $N(s)$ is the number of visits to a node s , $R(s)$ is the reward assigned to s , and p is the parent

node of s . At each level, the child node with the highest UCB value is selected. The C parameter balances exploration and exploitation.

Expansion. The selected node is expanded by sampling n query proposals from P_θ^{query} . The expansion is performed in a self-reflection manner to deal with the vast search space. All feedback O_f^{sib} from P_θ^{eval} , along with all the history of query proposals O_q^{sib} from sibling nodes, as well as retrieved documents O_d^{anc} and query proposals O_q^{anc} from all the ancestor nodes, are fed into P_θ^{query} as reflection material as shown in Table 2. This practice ensures diverse reasoning by expanding different and potentially better branches after each sampling.

Evaluation. A scalar value is assigned to each new child node using feedback give by P_θ^{eval} , quantifying the agent’s progress in task completion and steering the search algorithm towards promising regions. Following previous work Self-rewarding Language Models (Yuan et al., 2024b), we assess the currently retrieved document set $D \cup O_d^{anc}$ with LLM-as-a-Judge prompting with a 5-point rubric as shown in Table 1.

Backpropagation. The average value and visit counts of the newly generated node s' and its ancestors are updated using the scaled reward $R(s')$ from the Evaluation step.

By incorporating a self-rewarding paradigm and expanding strategy, SeRTS improves search efficiency and enhances the quality of retrieved documents. Details and pseudo-codes for SeRTS are illustrated in Appendix D.2.

3.2 Proximal Policy Optimization of LLMs with Feedback

To further improve the reasoning abilities of the Query Proposer and the Result Evaluator, we utilize Proximal Policy Optimization (PPO) to fine-tune the two LLMs using the trajectories collected by SeRTS as feedback (Schulman et al., 2017). Additionally, through this process, we could validate the effectiveness of the SeRTS reasoning process by examining the trajectories it produces.

PPO aims to optimize the policy of an LLM agent by maximizing the expected cumulative reward while ensuring that the LLMs are updated conservatively in case of performance degradation. The objective function of PPO can be formulated as follows:

$$L^{ppo}(\theta) = \hat{\mathbb{E}} \left[\min \left(r(\theta) \hat{A}_t, \text{clip}(r(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_t \right) \right],$$

where $r_\theta = \frac{\pi_\theta(A_{1..i} | O_{1..i})}{\pi_{\theta_{old}}(A_{1..i} | O_{1..i})}$

(2)

Here θ represents the parameters of the policy network, r_θ indicates the ratio of the probability of selecting an action under the new policy to that under the old one, and $A_{1..i}$ and $O_{1..i}$ denotes sampled query proposals and previous observations. \hat{A}_t is the advantage function estimating the relative reward of taking action $A_{1..i}$ in the state of $O_{1..i}$ compared to the average value of the state, and ϵ is the clipping range hyperparameter.

During training, we sample trajectories $\mathcal{T}_{SeRTS} = (A_{1..i}, O_{1..i}, R_{1..i})$ from the SeRTS retrieval process and the generated answers to estimate the advantage function and update the policy network θ . The clipping function clip in $L^{ppo}(\theta)$ ensures that the policy update is conservative, preventing drastic changes that could lead to instability or performance degradation.

By updating the language model using PPO with the feedback \mathcal{T}_{SeRTS} , we aim to let LMs know how to make good query proposals given current observations. This training process allows the LLMs to adapt to the specific requirements of medical knowledge retrieval, self-improving their ability of RAG in the bio-medical domain. The finetuning details are listed in Appendix D.4.

4 Experiments

4.1 Datasets

The BioASQ-QA (Krithara et al., 2023) dataset is a comprehensive resource designed to meet the

authentic information needs of biomedical professionals. It comprises questions in English, accompanied by gold-enriched answers by biomedical experts and relevant supporting material. This dataset stands out from most existing biomedical QA datasets due to its realistic and challenging nature, reflecting the complexity of real-world biomedical inquiries. Most importantly, BioASQ provides ideal answers in the form of summaries rather than only exact answers. This unique feature makes the dataset particularly valuable for researchers working on multi-document summarization techniques in the biomedical domain. Furthermore, the BioASQ-QA dataset is updated annually through an open-challenge format, making it highly unlikely for LLMs to have accessed the data during training. To evaluate our approach, we utilize the BioASQ 2023 Task 11b test set (Nentidis et al., 2023), which consists of question-answer pairs along with their corresponding retrieved documents.

4.2 Baselines

We compare our proposed SeRTS method with Self-Reflection baselines. For a fair comparison, we implement Self-Reflection by making minor modifications to the SeRTS prompt. Specifically, we remove the sibling node information (O_f^{sib} and O_d^{sib}) from the query proposer agent P_θ^{query} prompt, as self-reflection does not involve sibling nodes during expansion. The evaluator agent P_θ^{eval} prompt remains unchanged. We measure the retrieval capability of all methods by indexing the ground truth documents. Appendix D.3 provides details and pseudo-code for the self-reflection process.

4.3 Metrics

We evaluate retriever performance using standard metrics: **Precision**, **Recall**, and **F1 score**, comparing retrieved document IDs with the ground truth. **Hit Rate** assesses robustness by measuring the ratio of questions with at least one ground-truth document retrieved.

To measure RAG quality, we use **ROUGE-2 (R-2)** and **ROUGE-SU4 (R-SU4)**. R-2 computes shared bigrams between predicted answers and ground truth, while R-SU4 calculates skip-bigram overlap with a skip distance of four. Higher ROUGE scores indicate greater content overlap and more informative answers, suggesting better retrieval methods.

Retrieval Method	Precision	Recall	F1 score	Hit Rate
BM25	72.48	47.16	45.87	83.00
GPT-3.5-TURBO + REFLECTION (SIM 6)	66.26	58.96	51.18	88.89
GPT-3.5-TURBO + SERTS (SIM 6, MAX UCB)	74.31	59.09	55.06	92.22
LLAMA2-7B + REFLECTION (SIM 6)	72.48	48.66	47.12	84.44
LLAMA2-7B + SERTS (SIM 6, MAX UCB)	67.85	50.95	47.58	83.33
GPT-3.5-TURBO + REFLECTION (SIM 12)	71.36	58.84	52.72	91.11
GPT-3.5-TURBO + SERTS (SIM 12, MAX UCB)	74.65	58.96	54.70	94.44
LLAMA2-7B + REFLECTION (SIM 12)	68.04	51.82	48.79	83.33
LLAMA2-7B + SERTS (SIM 12, MAX UCB)	70.62	52.51	49.39	85.55

Table 3: Retrieval Results on BioASQ Task 11b batch 1 Test Set. The highest values are **bolded**. All metrics are averaged across three runs with three random seeds {42, 43, 44}.

4.4 Main Results

4.4.1 Comparison of the Retrieval Capability

Table 3 shows that SERTS (MAX UCB) consistently improves **Recall** and **Hit Rate** compared to REFLECTION for both GPT-3.5-TURBO and LLAMA2-7B-CHAT models. This performance gap supports our DoT assumption that higher-quality feedback induces more accurate reasoning. The performance gap between GPT-3.5-TURBO and LLAMA2-7B-CHAT is attributed to GPT-3.5’s larger parameter count (175B) and stronger reasoning capacity compared to LLAMA2 (7B). In addition, increasing the maximal simulation number further boosts the performance of both SERTS and REFLECTION. SERTS (SIM 12, MAX UCB) improves **Recall** and **Hit Rate** by 5.5 and 2.5 percentage points compared to the vanilla **BM25** retriever.

4.4.2 Effect on RAG

Table 4 presents the RAG performance of LLAMA2-7B-CHAT and GPT-3.5-TURBO using different retrieval methods on the bio-medical QA task. The large gap between **BM25** and **GROUND TRUTH** highlights the importance of document retrieval quality in generating final answers for the difficult medical QA task. Variants of **SeRTS** perform best among all methods on all metrics, with SERTS (MAX UCB) achieving 0.5 **R-2 F1** and 0.5 **R-SU4 F1** score improvements over BM25 on BioASQ dataset. The improvement of SERTS (MAX UCB) over REFLECTION demonstrates the effectiveness of our proposed method.

Fine-tuning with the PPO objective further improves the performance of REFLECTION and SERTS, with SERTS (PPO) acquiring 0.4 **R-2 F1** and 0.6 **R-SU4 F1** uplift over SERTS (MAX UCB) on LLAMA2-7B-CHAT. This suggests that

PPO fine-tuning enables self-improvement and domain knowledge acquisition for bio-medical RAG. SERTS (MAX UCB) achieves 1.5 points **R-SU4 F1** increase over BM25 using GPT-3.5-TURBO model, while LLAMA2-7B-CHAT + SERTS (PPO) brings a total of 1.1 point improvement, narrowing the retrieval performance gap between the small and large language models.

The superiority of SERTS method over REFLECTION is evident in both PPO and zero-shot variants when computation amounts are controlled to be similar (SIM 6 & SIM 12), showing the high scalability of SERTS compared to REFLECTION. The detailed statistics of SeRTS’s computational overhead can be found in Appendix B.

4.5 Ablation Study

4.5.1 Enhancement of PPO on SeRTS

We examine the effect of PPO fine-tuning on RAG performance by gradually switching the original weight θ of the proposer P_{θ}^{query} and evaluator P_{θ}^{eval} to θ_p fine-tuned by PPO. Table 5 shows that switching from P_{θ}^{query} to $P_{\theta_p}^{query}$ increases **R-SU4 F1** from 26.41 to 26.91 in REFLECTION and from 26.79 to 26.89 in SERTS. Further switching from P_{θ}^{eval} to $P_{\theta_p}^{eval}$ with $P_{\theta_p}^{query}$ as the query proposer continues to improve performance on all metrics. This suggests that $P_{\theta_p}^{eval}$ has learned to better evaluate retrieved documents and provide suggestions for improved queries, possibly due to the fact that the LM learns from feedback observations F and rewards R in trajectories \mathcal{T}_{SeRTS} about what feedback and suggestions can lead to better queries from the proposer P_{θ}^{query} .

Retrieval Method	R-2 Recall	R-2 F1	R-SU4 Recall	R-SU4 F1
GROUND TRUTH	57.02	56.64	58.26	57.42
BM25	25.58	22.74	29.63	26.22
GPT-3.5-TURBO + REFLECTION	26.92	24.08	30.86	27.41
GPT-3.5-TURBO + SERTS (MAX UCB)	27.25	24.35	31.25	27.77
LLAMA2-7B + REFLECTION	25.73	22.86	29.70	26.41
LLAMA2-7B + SERTS (MAX UCB)	26.18	23.20	30.30	26.79
LLAMA2-7B + REFLECTION (PPO)	26.20	23.41	30.24	26.93
LLAMA2-7B + SERTS (PPO)	26.51	23.63	30.73	27.34

Table 4: Answer Generation Results on BioASQ Task 11b batch 1 Test Set. The highest values are **bolded**. All metrics are averaged across three runs with three random seeds {42, 43, 44}.

Evaluating Agent	R-2 Recall	R-2 F1	R-SU4 Recall	R-SU4 F1
+ REFLECTION	25.73	22.86	29.70	26.41
+ REFLECTION (PPO) W. ORIGINAL EVALUATOR	26.15	23.36	30.11	26.83
+ REFLECTION (PPO) W. PPO EVALUATOR	26.20	23.41	30.24	26.93
+ SERTS	26.18	23.20	30.30	26.79
+ SERTS (PPO) W. ORIGINAL EVALUATOR	25.90	23.42	30.34	26.89
+ SERTS (PPO) W. PPO EVALUATOR	26.51	23.63	30.73	27.34

Table 5: Effect on Answer Generation Performance of PPO Finetuning on LLAMA2-7B-CHAT. The highest values are **bolded**. All metrics are averaged across three runs with three random seeds {42, 43, 44}.

Retrieval Method	R-2 Recall	R-2 F1	R-SU4 Recall	R-SU4 F1
BM25	25.58	22.74	29.63	26.22
REFLECTION	26.92	24.08	30.86	27.41
SERTS (RANDOM)	25.74	22.87	29.73	26.29
SERTS (MAX UCB, $C = 0.0$)	25.73	22.87	29.70	26.27
SERTS (MAX UCB, $C = 0.1$)	27.25	24.35	31.25	27.77

Table 6: Effect on Answer Generation Performance of Selection Policy in SERTS. The highest values are **bolded**. All metrics are averaged across three runs with three random seeds {42, 43, 44}.

4.5.2 Difference on Selection Policy

Table 6 examines how different selection policies in SERTS affect retrieval performance. Here we evaluate three policies: 1) RANDOM, which randomly explores nodes without exploiting assigned scores R ; 2) MAX UCB ($C = 0.0$), which only selects nodes with the highest R without exploring other options; and 3) MAX UCB ($C = 0.1$), which balances exploitation and exploration. The results empirically prove that striking a balance between exploitation and exploration is necessary for **SeRTS** to achieve ideal performance, while pure exploitation or exploration causes performance degradation.

5 Conclusion

We propose Self-Rewarding Tree Search (SeRTS), a plug-and-play retrieval method that combines LLM reasoning with conventional retrievers to address limitations in biomedical retrieval-augmented approaches. SeRTS introduces a reasoning process based on Monte Carlo Tree Search (MCTS) and a self-rewarding paradigm, boosting zero-shot performance in retrieving high-quality, informative results for Retrieval Augmented Generation (RAG). We further enhance retrieval by fine-tuning LLMs with PPO objectives using SeRTS retrieval trajectories as feedback.

Experiments using the BioASQ-QA dataset with GPT-3.5-Turbo and LLama2-7b show that

SeRTS significantly improves BM25 retriever performance, surpassing the self-reflection baseline in efficiency and scalability. SeRTS (PPO) effectively adapts LLMs to document retrieval, enhancing their ability to find highly relevant documents for RAG when answering medical queries. Our work presents an effective method for accurate, comprehensive medical knowledge retrieval using LLMs. While SeRTS outperforms Self-Reflection, both require numerous simulations, limiting widespread application. Future research should focus on developing more efficient reasoning strategies and leveraging the self-improvement capabilities of LLMs to further enhance retrieval efficiency and usability.

Limitations

PPO fine-tuning effectively improves SeRTS’s performance but is prone to Catastrophic Forgetting, degrading LLMs’ instruction-following capability. Table 7 shows error cases where P_{ϕ}^{query} fails to follow the designed instruction. Although manual templates can temporarily address these issues in our current experiments, the noisy trajectories impede further iteration and self-improvement using SeRTS sampling and PPO fine-tuning when fed back into the loop. As a result, we restrict our experiments to a single iteration of sampling and PPO fine-tuning. Future research should prioritize mitigating Catastrophic Forgetting and investigate the potential of multiple PPO fine-tuning iterations for further LLM self-improvement.

Acknowledgement

The work described in this paper was supported by the Chinese University of Hong Kong (CUHK; award numbers 4937025, 4937026, 5501517 and 5501329 to Y.L.) and the IdeaBooster Fund (IDBF23ENG05 and IDBF24ENG06 to Y.L.) and partially supported by grants from the Research Grants Council of the Hong Kong Special Administrative Region (Hong Kong SAR), China (project no. CUHK 14222922, RGC GRF 2151185 to I.K., and CUHK 24204023 to Y.L.) and a grant from the Innovation and Technology Commission of the Hong Kong SAR, China (project no. GHP/065/21SZ to Y.L.). This research was also funded by the Research Matching Grant Scheme at CUHK (award numbers 8601603 and 8601663 to Y.L.).

Based on the retrieved information, I would suggest the following query to search for more information:
`<query> Does losartan slow down brain atrophy in Alzheimer’s disease? A systematic review and meta-analysis. </query>`

Based on the information retrieved, I would suggest the following query to further search for relevant information:
`"Does losartan slow down brain atrophy in Alzheimer’s disease?"`

Based on the retrieved information, I would suggest the following query to retrieve more abstracts:
`<query> Does losartan slow down brain atrophy in Alzheimer’s disease? A systematic review and meta-analysis. Corresponding Feedback: Total Score: 40/50`

Based on the retrieved information, I would suggest the following query to further search for relevant information:
`"Losartan and brain atrophy in Alzheimer’s disease" Query Here`

Based on the retrieved information, I would suggest the following query to further search for additional information on dermatillomania:
`"<query> Causes of dermatillomania Query Here >"`

Table 7: Examples of P_{ϕ}^{query} generating query proposals. Cases where P_{ϕ}^{query} successfully follows the instruction are marked in red, while cases where P_{ϕ}^{query} fails to follow the instruction are marked in green.

References

- Afra Feyza Akyurek, Ekin Akyurek, Ashwin Kalyan, Peter Clark, Derry Tanti Wijaya, and Niket Tandon. 2023. [RL4F: Generating natural language feedback with reinforcement learning for repairing model outputs](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7716–7733, Toronto, Canada. Association for Computational Linguistics.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenhao Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. [A multi-task, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–718, Nusa Dua, Bali. Association for Computational Linguistics.
- Xidong Feng, Ziyu Wan, Muning Wen, Stephen Marcus McAleer, Ying Wen, Weinan Zhang, and Jun Wang. 2024. [Alphazero-like tree-search can guide large language model decoding and training](#).
- Yichen Gao, Licheng Zong, and Yu Li. 2024a. Enhancing biomedical question answering with parameter-

- efficient fine-tuning and hierarchical retrieval augmented generation. [CLEF Working Notes](#).
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024b. [Retrieval-augmented generation for large language models: A survey](#).
- Chengcheng Han, Xiaowei Du, Che Zhang, Yixin Lian, Xiang Li, Ming Gao, and Baoyuan Wang. 2023. [DialCoT meets PPO: Decomposing and exploring reasoning paths in smaller language models](#). In [Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing](#), pages 8055–8068, Singapore. Association for Computational Linguistics.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In [The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022](#). OpenReview.net.
- Minbyul Jeong, Jiwoong Sohn, Mujeen Sung, and Jaewoo Kang. 2024. [Improving medical reasoning through retrieval and self-reflection with retrieval-augmented large language models](#).
- Ehsan Kamaloo, Nouha Dziri, Charles L. A. Clarke, and Davood Rafiei. 2023. [Evaluating open-domain question answering in the era of large language models](#).
- Anastasia Krithara, Anastasios Nentidis, Konstantinos Bougiatiotis, and Georgios Paliouras. 2023. [Bioasqqa: A manually curated corpus for biomedical question answering](#). [Scientific Data](#), 10(1):170.
- Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, et al. 2023. [Measuring faithfulness in chain-of-thought reasoning](#). [ArXiv preprint, abs/2307.13702](#).
- Xibo Li, Bawei Zou, Yifan Fan, Yanling Li, Ai Ti Aw, and Yu Hong. 2023. [Interview evaluation: A novel approach for automatic evaluation of conversational question answering models](#). In [Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing](#), pages 3435–3446, Singapore. Association for Computational Linguistics.
- Khalil Mrini, Harpreet Singh, Franck Dernoncourt, Seunghyun Yoon, Trung Bui, Walter W. Chang, Emilia Farcas, and Ndapa Nakashole. 2022. [Medical question understanding and answering with knowledge grounding and semantic self-supervision](#). In [Proceedings of the 29th International Conference on Computational Linguistics](#), pages 2734–2747, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Anastasios Nentidis, Georgios Katsimpras, Anastasia Krithara, Salvador Lima López, Eulália Farré-Maduell, Luis Gasco, Martin Krallinger, and Georgios Paliouras. 2023. [Overview of bioasq 2023: The eleventh bioasq challenge on large-scale biomedical semantic indexing and question answering](#). In [International Conference of the Cross-Language Evaluation Forum for European Languages](#), pages 227–250. Springer.
- Santiago Ontanón. 2016. [Informed monte carlo tree search for real-time strategy games](#). In [2016 IEEE Conference on Computational Intelligence and Games \(CIG\)](#), pages 1–8. IEEE.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#).
- Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2022. [Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering](#). In [Proceedings of the Conference on Health, Inference, and Learning](#), volume 174 of [Proceedings of Machine Learning Research](#), pages 248–260. PMLR.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#).
- Xiaoming Shi, Zeming Liu, Li Du, Yuxuan Wang, Hongru Wang, Yuhang Guo, Tong Ruan, Jie Xu, and Shaoting Zhang. 2024. [Medical dialogue: A survey of categories, methods, evaluation and challenges](#).
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2024. [Reflection: Language agents with verbal reinforcement learning](#). [Advances in Neural Information Processing Systems](#), 36.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. [Mastering the game of go with deep neural networks and tree search](#). [nature](#), 529(7587):484–489.
- David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. 2017. [Mastering chess and shogi by self-play with a general reinforcement learning algorithm](#).
- Chiara F Sironi, Jialin Liu, Diego Perez-Liebana, Raluca D Gaina, Ivan Bravi, Simon M Lucas, and Mark HM Winands. 2018. [Self-adaptive mcts](#)

- for general video game playing. In International Conference on the Applications of Evolutionary Computation, pages 358–375. Springer.
- Ye Tian, Baolin Peng, Linfeng Song, Lifeng Jin, Dian Yu, Haitao Mi, and Dong Yu. 2024. [Toward self-improvement of llms via imagination, searching, and criticizing.](#)
- Yuanhe Tian, Weicheng Ma, Fei Xia, and Yan Song. 2019. [ChiMed: A Chinese medical corpus for question answering.](#) In Proceedings of the 18th BioNLP Workshop and Shared Task, pages 250–260, Florence, Italy. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models.](#) ArXiv preprint, abs/2307.09288.
- Hongru Wang, Minda Hu, Yang Deng, Rui Wang, Fei Mi, Weichao Wang, Yasheng Wang, Wai-Chung Kwan, Irwin King, and Kam-Fai Wong. 2023. [Large language models as source planner for personalized knowledge-grounded dialogues.](#) In Findings of the Association for Computational Linguistics: EMNLP 2023, pages 9556–9569, Singapore. Association for Computational Linguistics.
- Hongru Wang, Wenyu Huang, Yang Deng, Rui Wang, Zezhong Wang, Yufei Wang, Fei Mi, Jeff Z. Pan, and Kam-Fai Wong. 2024a. [Unims-rag: A unified multi-source retrieval-augmented generation for personalized dialogue systems.](#)
- Hongru Wang, Huimin Wang, Zezhong Wang, and Kam-Fai Wong. 2022. [Integrating pretrained language model for dialogue policy evaluation.](#) In ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 6692–6696. IEEE.
- Hongru Wang, Boyang Xue, Baohang Zhou, Tianhua Zhang, Cunxiang Wang, Guanhua Chen, Huimin Wang, and Kam fai Wong. 2024b. [Self-dc: When to retrieve and when to generate? self divide-and-conquer for compositional unknown questions.](#)
- Yubo Wang, Xueguang Ma, and Wenhui Chen. 2024c. [Augmenting black-box llms with medical textbooks for clinical question answering.](#)
- Fei Xia, Bin Li, Yixuan Weng, Shizhu He, Kang Liu, Bin Sun, Shutao Li, and Jun Zhao. 2022. [MedConQA: Medical conversational question answering system based on knowledge graphs.](#) In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pages 148–158, Abu Dhabi, UAE. Association for Computational Linguistics.
- Guangzhi Xiong, Qiao Jin, Zhiyong Lu, and Aidong Zhang. 2024. [Benchmarking retrieval-augmented generation for medicine.](#)
- Xiao Yu, Maximillian Chen, and Zhou Yu. 2023. [Prompt-based Monte-Carlo tree search for goal-oriented dialogue policy planning.](#) In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, pages 7101–7125, Singapore. Association for Computational Linguistics.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. 2024a. [Self-rewarding language models.](#)
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. 2024b. [Self-rewarding language models.](#) ArXiv preprint, abs/2401.10020.
- Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Chuanqi Tan, and Chang Zhou. 2023. [Scaling relationship on learning mathematical reasoning with large language models.](#) ArXiv preprint, abs/2308.01825.
- Cyril Zarka, Akash Chaurasia, Rohan Shad, Alex R. Dalal, Jennifer L. Kim, Michael Moor, Kevin Alexander, Euan Ashley, Jack Boyd, Kathleen Boyd, Karen Hirsch, Curt Langlotz, Joanna Nelson, and William Hiesinger. 2023. [Almanac: Retrieval-augmented language models for clinical medicine.](#)
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. 2022. [Star: Bootstrapping reasoning with reasoning.](#) Advances in Neural Information Processing Systems, 35:15476–15488.
- Dan Zhang, Sining Zhou, Yisong Yue, Yuxiao Dong, and Jie Tang. 2024. [Rest-mcts*: Llm self-training via process reward guided tree search.](#)
- Ming Zhang, Shuai Dou, Ziyang Wang, and Yunfang Wu. 2022. [Focus-driven contrastive learning for medical question summarization.](#) In Proceedings of the 29th International Conference on Computational Linguistics, pages 6176–6186, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. 2023. [Language agent tree search unifies reasoning acting and planning in language models.](#) ArXiv preprint, abs/2310.04406.
- Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. 2024. [Language agent tree search unifies reasoning acting and planning in language models.](#)

A Five-Point Rubrics of Evaluator

Details of evaluation rubrics are illustrated in Table 8.

- Add 1 point if the knowledge is relevant and provides some information related to the user’s inquiry, even if it is incomplete or contains some irrelevant content.
- Add another point if the knowledge addresses a substantial portion of the user’s question, but does not completely resolve the query or provide a direct answer.
- Award a third point if the response answers the basic elements of the user’s question in a useful way, regardless of whether it seems to have been written by an AI Assistant or if it has elements typically found in blogs or search results.
- Grant a fourth point if the knowledge addresses the user’s question directly and comprehensively, even if there is slight room for improvement in clarity, conciseness, or focus.
- Bestow a fifth point for knowledge that is impeccably tailored to the user’s question, without extraneous information, reflecting expert knowledge, and can be used to produce a high-quality, engaging, and insightful answer.

Table 8: Details of {Five-point Rubrics}

B Statistics of Computational Overhead

To show how much overhead is introduced, we measured the mean number of tokens per question generated by the LLAMA2-7B-CHAT model for both our method (SERTS) and the Self-Reflection baseline during the retrieval and reasoning process. The results are Table 9.

Retrieval Method	Avg. #Token
SELF-REFLECTION (SIM 6)	4691.0
SERTS (SIM 6)	4758.3
SELF-REFLECTION (SIM 12)	13,330.3
SERTS (SIM 12)	9,309.4
SERTS (SIM 12, PPO)	5,175.9

Table 9: Average number of tokens per question generated by LLAMA2-7B-CHAT during the experiment in Table 4.

SERTS significantly reduces the number of generated tokens compared to SELF-REFLECTION. With a simulation round of 12, SERTS generates 30.16% fewer tokens than SELF-REFLECTION. When incorporating PPO into SERTS, the token reduction achieves 57.12%.

C Experiments on MedMCQA dataset

We conducted experiments on the MedMCQA (Pal et al., 2022) dataset to demonstrate the robustness of SERTS. Using GPT-3.5-TURBO as the answer generator, we sample 100 questions from the development set and compare four configurations: GPT-3.5 without retrieval, GPT-3.5 with

BM25 retriever, SELF-REFLECTION (SIM 12), and SERTS (SIM 12). SERTS achieves an accuracy of 0.57, outperforming the no-retrieval baseline by 7.55% and SELF-REFLECTION by 3.64%, demonstrating its effective generalization to other biomedical QA datasets.

Method	Accuracy
GPT-3.5 WO. RETRIEVAL	0.53
GPT-3.5 W. RETRIEVAL	0.55
GPT-3.5 W. REFLECTION	0.55
GPT-3.5 W. SERTS	0.57

Table 10: Performance of multi-choice QA on MedMCQA dataset.

D Implementation Details

D.1 RAG Pipeline

It searches for relevant text from the document test set using the questions as queries and supplies these texts as references to the language model. The language model we used is the LLAMA2-7B-CHAT (Touvron et al., 2023), which has been fine-tuned on the BioASQ corpus using LoRA (Hu et al., 2022). This corpus comprises 4,719 question-answer pairs from the BioASQ task 11b training set. Importantly, the documents in the training set are distinct from those in the test set, ensuring no overlap. The BM25 retrieval function can be expressed as follows:

$$\text{score}(q, d) = \frac{\sum_{i=1}^n IDF(q_i) \cdot f(q_i, d) \cdot (k_1 + 1)}{f(q_i, d) + k_1 \cdot \left(1 - b + b \cdot \frac{|d|}{\text{avgdl}}\right)} \quad (3)$$

, where q is the query, and d is the document. q_i is the i -th term in the query. $f(q_i, d)$ is the term frequency of q_i in the document d . $|d|$ is the length of the document d in terms of the number of words. avgdl is the average document length in the collection. k_1 and b are free parameters, usually chosen, for example, $k_1 = 1.2$ and $b = 0.75$. $IDF(q_i)$ is the inverse document frequency weight of the term q_i , calculated as:

$$IDF(q_i) = \log \left(\frac{N - n(q_i) + 0.5}{n(q_i) + 0.5} + 1 \right) \quad (4)$$

, where N is the total number of documents in the collection. $n(q_i)$ is the number of documents containing the term q_i .

To minimize the influence of the prompt, we employed a very simple prompt to instruct the language model to generate appropriate answers based on the provided text. The structure of this prompt is illustrated as follows.

```
You are an expert in the field of biomedical science.

Below is an instruction that describes a task, paired with
an input that provides further context. Write a response
that appropriately completes the request.

### Instruction: {question}

### Input: {context}

### Response:
```

Table 11: Prompt for Question Answering.

D.2 Self-Rewarding Tree Search (SeRTS)

We use the Llama2-7b-chat model published in the Huggingface¹ for experiments in this work. The Llama2-7b-chat model is loaded in 4 bits, and it uses flash attention 2 as attention implementation for efficiency. For sequence generation, we keep the same configuration of setting the temperature to 0.7 and maximal length to 4096 tokens. For all experiments, we report the averaged metrics from three runs with the random seed fixed to {42,43,44}. For the SeRTS process illustrated in Algorithm 1, we set the maximal simulation number Sim to 12, maximal branch number b to 3, and maximal depth D to 3.

D.3 Self-Reflection

Compared with prompt for P_{θ}^{query} in Table 2, prompt for $P_{\theta}^{reflect}$ in Table 12 removes all feedback from sibling nodes. Algorithm 2 shows a detailed implementation of the self-reflection process, node selection, and backpropagation operation is removed compared to SeRTS. All other configurations of the self-reflection process remain unchanged.

D.4 LoRA PPO Finetuning

For PPO fine-tuning, we utilize the LoRA and PPO implementations from the peft and trl Python packages, following the official training script². The fine-tuning process is completed

¹<https://huggingface.co/meta-llama/Llama-2-7b-chat-hf>

²https://github.com/huggingface/trl/blob/main/examples/research_projects/stack_llama/scripts/rl_training.py

```
You are a medical expert, and you are tasked with searching
for as much useful knowledge from a database of
medical paper abstracts as possible to address the user's
question. The question is stated as follows: {Q}
```

```
You are given the BM25 retriever as the search tool. According
to previous query proposals and proposals' corresponding
feedback and suggestions, give one single better
query for the retriever to further search for extra information
needed for answering the given question.
```

```
Information already known is as follows:<retrieved_info>
Query and its reasoning for Abstracts: {O_q^{anc}}
Retrieved Abstracts: {O_d^{anc}}
</retrieved_info>
```

```
After examining the information above, justify your decision,
in up to 100 words.
```

Hints:

1. Use simple keywords as queries if no information is retrieved.
2. You can take one step at a time, subdivide the original question, and search for information on the sub-questions.

```
Conclude with the query for the paper abstracts using the
format: "<query> Query Here </query>".
```

Table 12: Prompt for Self-Reflect Query Proposer $P_{\theta}^{reflect}$.

on a single Nvidia A6000 GPU within 2 hours. The `lora_target_modules` is set to `["q_proj", "v_proj"]`, with `lora_r` as 16, `lora_alpha` as 32, and `lora_dropout` as 0.05. The `micro_batch_size` is 1, `batch_size` is 8, and `num_epochs` is 1. For optimization, we set the `learning_rate` to 1.41e-5, the `reward baseline` to 3.75, and the `random seed` to 0.

Algorithm 1 The Self-Rewarding Tree Search (SeRTS) Algorithm.

```
1: Input: question  $Q$ , query proposer  $P_{\theta}^{query}$ , evaluator  $P_{\theta}^{eval}$ , max_simulation_number  $Sim$ ,  
max_branch_number  $b$ , max_depth  $Depth$ , BM25_retriever  $\mathcal{R}_{bm25}$ .  
2:  $T_Q \leftarrow \text{Initialize\_tree}(Q)$   
3: for  $i$  in range( $Sim$ ) do  
4:    $C \leftarrow \text{root}(T_Q)$   
5:   Selection  
6:   while  $C$  is not leaf node do do  
7:      $C \leftarrow \arg \max_s UCB(s) = \frac{R(s)}{N(s)} + C * \sqrt{2 * \ln \frac{N(p)}{N(s)}}$   $\triangleright$  Select child node based on UCB  
8:     if  $C$ 's depth  $\geq Depth$  then  $\triangleright$  Break if  $C$  reaches maximal depth  
9:       break  
10:    end if  
11:  
12:    if  $C$  has less than  $b$  children nodes then  $\triangleright$  Select  $C$  if it is not fully expanded  
13:      break  
14:    end if  
15:  end while  
16:  if  $C$ 's depth  $\geq Depth$  then  $\triangleright$  Continue to select other nodes if  $C$  reaches maximal depth  
17:    continue  
18:  end if  
19:  
20:  Expansion & Evaluation  
21:   $O_q^{anc} \leftarrow \bigcup_{k \in \text{ancestors}(C)} A_k$  : query proposals from all  $C$ 's ancestor nodes  
22:   $O_d^{anc} \leftarrow \bigcup_{k \in \text{ancestors}(C)} D_k$  : retrieved documents from all  $C$ 's ancestor nodes  
23:   $O_q^{sib} \leftarrow \bigcup_{k \in \text{siblings}(C)} A_k$  : query proposals from all  $C$ 's sibling nodes  
24:   $O_f^{sib} \leftarrow \bigcup_{k \in \text{siblings}(C)} F_k$  : feedback from all  $C$ 's sibling nodes  
25:   $O \leftarrow (O_q^{anc}, O_d^{anc}, O_q^{sib}, O_f^{sib})$   
26:   $A \leftarrow P_{\theta}^{query}(O, Q)$   $\triangleright$  Make new query proposal  $A$  based on  $O$  and  $Q$   
27:   $D \leftarrow \mathcal{R}_{bm25}(A)$   $\triangleright$  Invoke  $\mathcal{R}_{bm25}$  to retrieve documents according to  $A$   
28:   $R, F \leftarrow P_{\theta}^{eval}(D \cup O_d^{anc}, Q)$   $\triangleright$  Get reward score  $R$  and descriptive feedback  $O_f$  for currently  
retrieved documents  $D \cup O_d^{anc}$  and  $Q$   
29:  if  $R == 5$  then  $\triangleright$  Stop early if  $D \cup O_d^{anc}$  receive highest score in 5-point rubrics  
30:    break  
31:  end if  
32:  
33:   $C' = \text{create\_new\_child}(Q, A, R, O)$   $\triangleright$  create a new child node for  $C$   
34:  Add  $C'$  to children nodes of  $C$   
35:  Backpropagation  
36:   $\text{back\_propagate}(C)$   $\triangleright$  Update value of parent nodes using new score from  $C'$   
37: end for  
38:  $C = \text{get\_best\_node}(T_Q)$   $\triangleright$  Fetch the node with the highest  $R$  in  $T_Q$   
39: return  $D \cup O_d^{anc}$  from  $C$   $\triangleright$  Return all retrieved documents in  $C$ 
```

Algorithm 2 The Self-Reflection Algorithm.

```
1: Input: question  $Q$ , self-reflect query proposer  $P_\theta^{reflect}$ , evaluator  $P_\theta^{eval}$ , max_simulation number  $Sim$ , max_branch_number  $b$ , max_depth  $Depth$ , BM25_retriever  $\mathcal{R}_{bm25}$ .
2:  $T_Q \leftarrow \text{Initialize\_tree}(Q)$ 
3: for  $i$  in range( $Sim$ ) do
4:    $C \leftarrow \text{root}(T_Q)$ 
5:   Expansion & Evaluation
6:    $O_q^{anc} \leftarrow \bigcup_{k \in \text{ancestors}(C)} A_k$  : query proposals from all  $C$ 's ancestor nodes
7:    $O_d^{anc} \leftarrow \bigcup_{k \in \text{ancestors}(C)} D_k$  : retrieved documents from all  $C$ 's ancestor nodes
8:    $O \leftarrow (O_q^{anc}, O_d^{anc})$ 
9:    $A \leftarrow P_\theta^{reflect}(O, Q)$  ▷ Make new query proposal  $A$  based on  $O$  and  $Q$ 
10:   $D \leftarrow \mathcal{R}_{bm25}(A)$  ▷ Invoke  $\mathcal{R}_{bm25}$  to retrieve documents according to  $A$ 
11:   $R, F \leftarrow P_\theta^{eval}(D \cup O_d^{anc}, Q)$  ▷ Get reward score  $R$  and descriptive feedback  $F$  for currently retrieved documents  $D \cup O_d^{anc}$  and  $Q$ 
12:  if  $R == 5$  then ▷ Stop early if  $D \cup O_d^{anc}$  receive highest score in 5-point rubrics
13:    break
14:  end if
15:
16:   $C' = \text{create\_new\_child}(Q, A, R, O)$  ▷ create a new child node for  $C$ 
17:  Add  $C'$  to children nodes of  $C$ 
18: end for
19:  $C = \text{get\_best\_node}(T_Q)$  ▷ Fetch the node with the highest  $R$  in  $T_Q$ 
20: return  $D \cup O_d^{anc}$  from  $C$  ▷ Return all retrieved documents in  $C$ 
```
