

Structured Outputs Enable General-Purpose LLMs to be Medical Experts

Anonymous ACL submission

Abstract

Medical question-answering (QA) is a critical task for evaluating how effectively large language models (LLMs) encode clinical knowledge and assessing their potential applications in medicine. Despite showing promise on multiple-choice tests, LLMs frequently struggle with open-ended medical questions, producing responses with dangerous hallucinations or lacking comprehensive coverage of critical aspects. Existing approaches attempt to address these challenges through domain-specific fine-tuning, but this proves resource-intensive and difficult to scale across models. To improve the comprehensiveness and factuality of medical responses, we propose a novel approach utilizing structured medical reasoning. Our method guides LLMs through an seven-step cognitive process inspired by clinical diagnosis, enabling more accurate and complete answers without additional training. Experiments on the MedLFQA benchmark demonstrate that our approach achieves the highest Factuality Score of 85.8, surpassing fine-tuned models. Notably, this improvement transfers to smaller models, highlighting the method’s efficiency and scalability. Our code and datasets are available.

1 Introduction

Language lies at the heart of healthcare delivery, serving as the primary medium through which medical knowledge is communicated and healthcare decisions are made. Recent advances in LLMs have shown promising potential in medical applications, particularly in tasks requiring deep clinical knowledge and reasoning (Singhal and et al., 2022; Tu and et al., 2023; Labrak et al., 2024; Jeong et al., 2024b; Liu et al., 2023). Among these applications, medical QA stands out as a critical task for evaluating how effectively LLMs encode and apply clinical knowledge. This task is particularly challenging, as it requires an understanding of medical context, the recall of relevant knowledge, and the

ability to reason with expert-level information.

While existing LLMs have demonstrated impressive performance on medical benchmarks, such as MedMCQA (Pal et al., 2022) and USMLE (Examination, 2023), these evaluations primarily focus on knowledge retrieval rather than comprehensive reasoning. This limitation has led to the development of more challenging benchmarks like MedLFQA (Jeong et al., 2024a), which demands detailed, factually accurate, and clinically relevant long-form responses.

However, when tasked with generating long-form medical answers, LLMs frequently struggle with two critical challenges. First, they often produce responses containing dangerous hallucinations or factual errors. As shown in Figure 6, consider a patient’s inquiry regarding the concurrent use of Zyrtec and Hydroxyzine: while an accurate response should detail potential adverse interactions between these antihistamines and emphasize the importance of professional medical guidance, LLMs sometimes generate misleading or incorrect information about drug compatibility. Such inaccuracies could lead to serious clinical consequences if relied upon for medical decision-making.

As shown in the Figure 4, the second challenge lies in generating comprehensively informed responses. When asked about a health condition like boils and carbuncles, an adequate response must cover multiple aspects: how they are treated, potential causes, and common symptoms. Current models sometimes provide incomplete responses, addressing only certain parts of the question while leaving out critical information. This lack of comprehensive detail can result in patients or healthcare providers missing important aspects of diagnosis and treatment.

Existing approaches have primarily relied on two strategies to address these challenges: supervised fine-tuning with domain-specific datasets and chain-of-thought (CoT) prompting (Kalajdziewski,

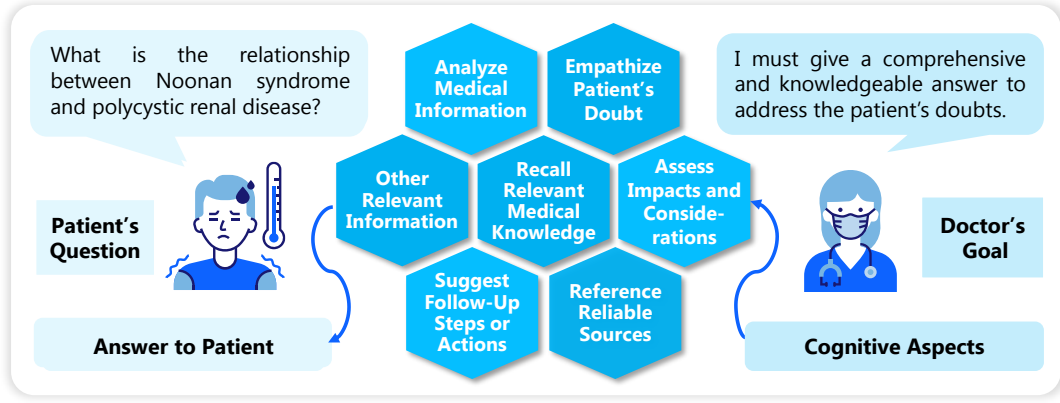


Figure 1: A flowchart showing the doctor’s cognitive process to answer a patient’s question, involving medical analysis, relevant information, and follow-up steps.

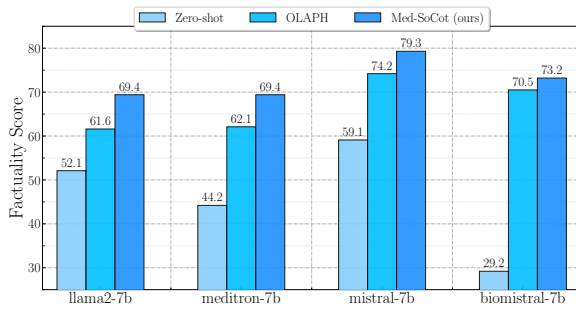


Figure 2: Factuality Scores for different models (LLaMA2-7B, Meditron-7B, Mistral-7B, BioMistral-7B) across three methods: Zero-shot, OLAPH, and Med-SoCoT (Ours).

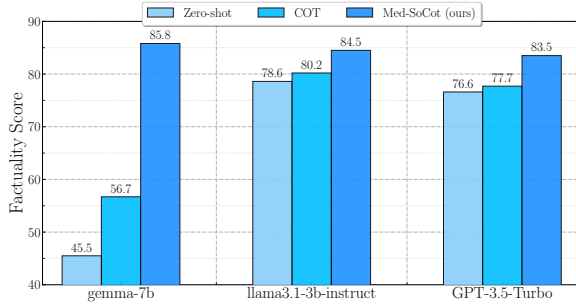


Figure 3: Factuality Scores for different models (Gemma-7B, LLaMA3.1-3B-INSTRUCT, GPT-3.5-Turbo) across three methods: Zero-shot, CoT and Med-SoCoT (Ours).

2024). While fine-tuning can improve performance, it requires substantial computational resources and risks catastrophic forgetting of pre-trained knowledge. CoT prompting has shown promise in enhancing reasoning capabilities, but its effectiveness in generating comprehensive medical responses remains limited, particularly for models with restricted context lengths.

Drawing inspiration from cognitive science research on how humans process complex information (Thagard, 1996; Breedlove et al., 2010; Gazzaniga et al., 2018; Krathwohl, 2002; Duncan, 2013), we propose a novel medical structured

output approach. As shown in the Figure 1, this method guides LLMs through a systematic process of medical reasoning, breaking down the response generation into distinct cognitive steps: understanding the question, recalling relevant medical knowledge, analyzing implications, and synthesizing comprehensive answers. By explicitly structuring the output process, we enable LLMs to better organize and validate their medical knowledge, leading to more accurate and complete responses (Zhao et al., 2024).

Our research demonstrates that **Medical Structured Output CoT (Med-SoCoT)** effectively mitigates hallucinations and "echo" phenomena while improving the factuality in a training-free manner, enabling even smaller models to generate high-quality medical responses. For models with limited context length, overly long prompts can hinder the complete output of structured content. To address this, we adopt a step-by-step reasoning approach, where the model generates one piece of structured content at a time, and the outputs are concatenated for summarization. Notably, this approach notably achieves performance improvements on par with fine-tuned models, demonstrating that prompt engineering alone can substantially enhance factuality without requiring additional training.

Our research focuses on two key questions: First, whether structured output can enhance the factual accuracy of medical question answering; and second, whether smaller models can achieve structured outputs solely through prompt engineering. Our contributions are summarized as follows:

- We integrate chain-of-thought reasoning with structured output to develop novel medical prompts, demonstrating that general-purpose LLMs can produce comprehensive answers

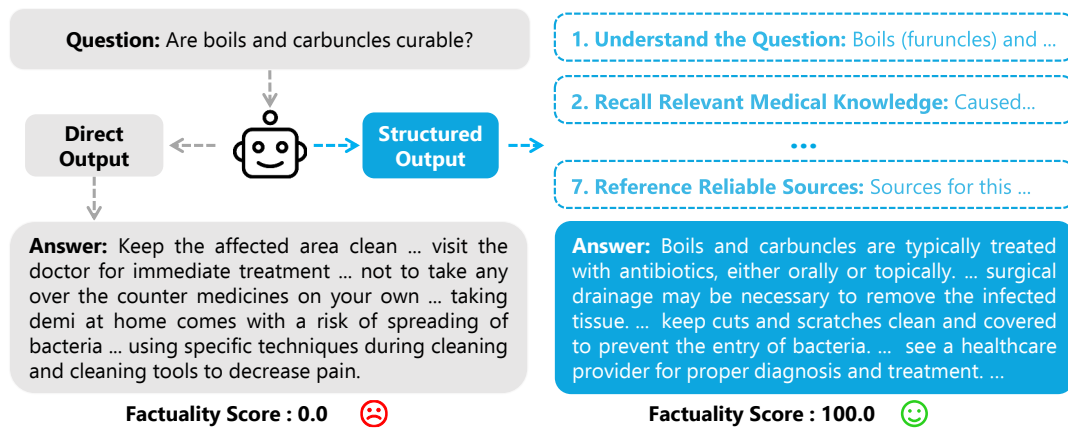


Figure 4: Model: Gemma2-7B. LLMs often generate unreliable answers, due to cognitive limitations. Structured output helps LLMs analyze problems step by step, leading to more complete and accurate answers.

with reduced hallucinations without additional training. Our approach achieves the highest Factualty Score of 85.8, surpassing even fine-tuned models, which achieve a maximum Factualty Score of 74.2.

- For models with shorter context lengths, we employed stepwise reasoning without altering the prompt content, achieving an 12.7% improvement in the comprehensiveness of the model’s output, increasing Factualty Score from 61.6 to 69.4.
- For smaller models (3B), our approach remains highly effective, boosting Factualty Score from 78.6 in Zero-shot settings to 84.5. This improvement highlighting the robustness of our method across model sizes.

as it allows for detailed and comprehensive responses to open-ended questions. This is especially important in biomedical and clinical settings, where accuracy and depth are vital for patient care (Fan et al., 2019; Stelmakh et al., 2023). To address these needs, Minbyul developed the MedLFQA dataset by combining several existing QA datasets (Abacha et al., 2019). MedLFQA provides a valuable resource for evaluating how large language models (LLMs) handle patient queries, offering insights into the quality of their long-form responses. The dataset includes five distinct medical QA datasets and encompasses various fields like question, answer, and essential details, enabling a thorough assessment of the comprehensiveness and accuracy of LLM-generated medical content.

2 Related Work

2.1 Medical question answering

The landscape of medical question answering (QA) is diverse, encompassing various formats that cater to different aspects of clinical reasoning and knowledge application. These formats include multiple-choice QA (Jin et al., 2020; Hendrycks et al., 2021; Jin et al., 2019; Pal et al., 2022; Xiong et al., 2024), fragment-based QA (Krithara et al., 2023), open-domain QA (Lee et al., 2019), factoid QA (Jin et al., 2022), multi-turn dialogue QA (Bao et al., 2023), knowledge-base QA (Lozano et al., 2023), and clinical decision-making QA (Hager et al., 2024). Each format presents unique challenges, ranging from simple knowledge retrieval to complex reasoning and multi-step decision-making, reflecting the multifaceted nature of real-world medical problem-solving.

Among these, Long-Form Question Answering (LFQA) is particularly crucial in the medical field,

2.2 Evaluation of Long-Text Generation

To evaluate the quality of long text generation by LLMs, ROUGE metrics, such as ROUGE-1, ROUGE-2, and ROUGE-L, are commonly used to assess text similarity at different granularities. To provide a more comprehensive evaluation, we use the Words Composition metric from MedLFQA, which is calculated as shown in Equation A.1. This metric averages the F1 scores of ROUGE-1, ROUGE-2, and ROUGE-L, offering a more accurate assessment of the text’s quality in terms of word-level, bigram-level, and sentence structure performance.

Some metrics, such as Hallucination Score and Comprehensiveness Score (Manes et al., 2024a), are used to assess the inclusion of key statements in the clinical domain. Specifically, Hallucination Score is a metric designed to measure how many clinical statements in the language model’s response contradict factual information. This score

is calculated using Equation A.2. Where S represents all the statements that must be included (Must Have, MH) and those that are optional (Nice to Have, NH), such that $|S| = |\text{MH}| + |\text{NH}|$. Additionally, Comprehensiveness Score is a metric used to measure how many key clinical statements are included in the language model’s response. This score is calculated using Equation A.3. To predict the entailment of the response, we used a classification model based on BioBERT (Lee et al., 2019), which was trained on NLI datasets (Bowman et al., 2015; Williams et al., 2018).

The Factuality Score is calculated using Equation A.4. Our work is based on these fine-grained and cost-effective evaluation metrics to understand the factuality and lexical structure of LLMs when generating long text.

2.3 Structured Output

In traditional NLP literature, structured data typically refers to the extraction of entity-relation-entity triples and attributes from plain text, followed by the construction of knowledge graphs or databases with specific data formats or schemas (Jiang et al., 2023b; Li et al., 2023). In previous work, "structuring" was used to reorganize input sentences into a three-layer structure based on the internal semantic relationships between them. This method captures the dependencies and connections among elements in long-text inputs, thereby enhancing LLMs’ understanding of knowledge structure and relationships. Earlier studies have explored how structuring input information can improve LLMs’ cognitive abilities (Liu et al., 2024). In the work on GPT-4o, structured output refers to organizing the content generated by a language model into a specific format or pattern (OpenAI, 2023). This output often includes predefined labels, fields, or sections to make it more convenient for users to interpret and use. Structured output not only helps reduce redundancy in generated content but also ensures that the text meets certain requirements or standards.

In this paper, instead of focusing on entity-level information extraction or restructuring the input content, we primarily leverage the model’s structured output capabilities to guide it in completing structured thinking. This enhances the model’s ability to recognize the structure and relationships between recalled knowledge and the posed questions.

3 Methodology

3.1 Overview

Structured reasoning plays a crucial role in modern medical applications by enhancing diagnostic accuracy and decision-making efficiency through the systematic organization of complex medical knowledge. By emulating human cognitive processes, it ensures a logical and coherent approach to clinical problem-solving. This is particularly vital for large language models (LLMs) in medical decision-making, where, despite vast knowledge, LLMs often produce fragmented outputs without guidance. A structured framework allows these models to analyze data systematically, consider relevant clinical factors, and generate contextually appropriate, reliable conclusions, improving both accuracy and consistency.

To implement this, we propose a framework based on established clinical reasoning theories (Higgs, 2020; Young et al., 2020; Shin, 2019; Pelaccia et al., 2011). It breaks down diagnostic and therapeutic reasoning into systematic steps, ensuring a logical progression from patient assessment to treatment planning. Supported by evidence-based strategies and structured models, this approach yields precise, actionable outputs that enhance clinical decision-making.

3.2 Medical Structured Reasoning Process

Structured reasoning is a cognitive approach that organizes complex medical problems into systematic decision-making steps. Inspired by clinical reasoning in medicine (Higgs, 2020; Young et al., 2020; Shin, 2019; Pelaccia et al., 2011), we propose a structured framework comprising seven key cognitive processes. This framework provides large language models (LLMs) with a clear reasoning pathway that mimics human medical cognition, thereby improving accuracy, interpretability, and reducing reasoning errors.

1. Understand the Question

The first step in medical reasoning is accurately understanding the patient’s problem. This involves an in-depth analysis of the chief complaint, medical history, and contextual factors. Research indicates that a well-defined problem statement is the foundation of a successful diagnosis (Ng et al., 2024). For LLMs, precise input interpretation is critical in minimizing errors and setting the context for subsequent logical analysis.

2. Recall Relevant Medical Knowledge

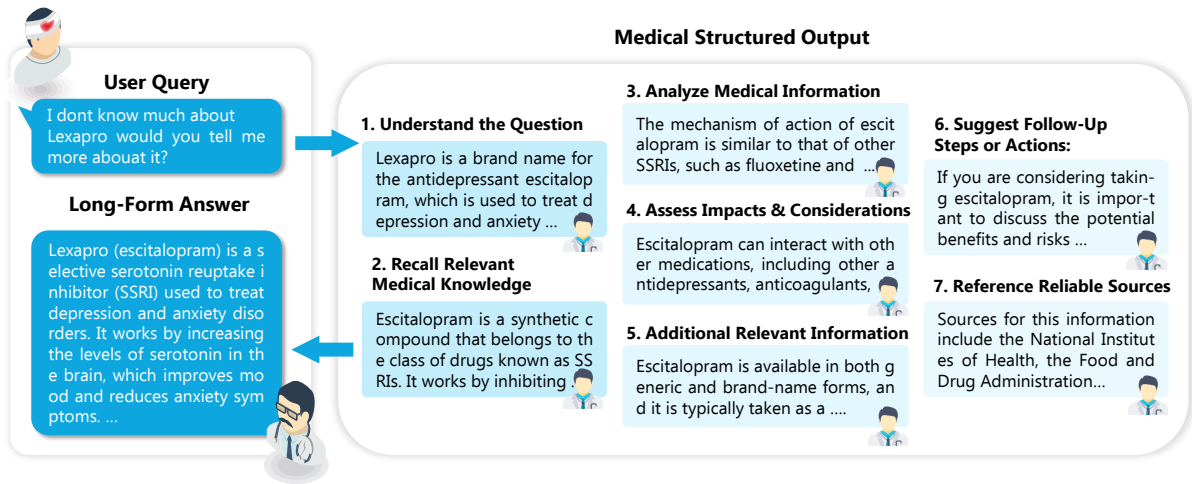


Figure 5: Model: Gemma-7B, Dataset: LiveQA. An overview of structured medical reasoning process shows a step-by-step framework to improve the comprehensiveness and factuality of medical QA.

Once the problem is clearly defined, clinicians recall relevant disease patterns, risk factors, and physiological mechanisms. The disease script theory suggests that experienced physicians rely on stored patterns for rapid decision-making (Thampy et al., 2019). This step is essential for LLMs, as an effective knowledge retrieval mechanism can reduce hallucinations and ensure medically accurate responses. Moreover, robust knowledge recall directly supports the next step—logical analysis of the medical information.

3. Analyze Medical Information

Using the retrieved knowledge, clinicians systematically compare and interpret the collected patient data to generate differential diagnoses. The dual-process theory highlights that analytical reasoning is particularly crucial in complex or unfamiliar cases (Higgs, 2020). For LLMs, this step requires strong logical reasoning capabilities to avoid misinterpretations caused by data noise or biases. Furthermore, accurate knowledge recall (Step 2) enhances this analytical process by providing the necessary context for reasoning.

4. Assess Impacts and Considerations

After forming diagnostic hypotheses, clinicians evaluate potential treatment risks, ethical considerations, and resource constraints. Multilevel decision frameworks, such as the Four Quadrant Model, help balance medical, ethical, and logistical factors (Corrao and Argano, 2022). This step bridges logical analysis (Step 3) and actionable decision-making (Step 5). For LLMs, incorporating structured assessment criteria reduces bias and ensures recommendations align with clinical best practices.

5. Provide Additional Relevant Information

During diagnosis and treatment, clinicians must

ensure that all necessary medical information is conveyed clearly to patients and colleagues. Studies show that effective patient education fosters trust and improves treatment adherence (Richards and et al., 2020). This step is crucial for LLMs, as providing clear, accessible explanations enhances transparency and interpretability. Additionally, structured information delivery reinforces the justification for proposed actions (Step 6).

6. Suggest Follow-Up Steps or Actions

A well-structured medical decision should include a personalized follow-up plan, specifying additional tests, treatment goals, and monitoring schedules. Research indicates that structured follow-up plans help prevent complications and improve patient adherence (Bae et al., 2023). This step naturally follows information provision (Step 5), ensuring continuity of care. For LLMs, integrating clinical pathways into their reasoning process enables them to generate practical and clinically relevant recommendations.

7. Reference Reliable Sources

Finally, clinical decisions must be grounded in peer-reviewed literature and evidence-based guidelines (e.g., Cochrane Database, PubMed). Medical education research emphasizes that reliance on high-quality information sources is critical for maintaining diagnostic and treatment consistency (Young et al., 2020; Pelaccia et al., 2011). This final step strengthens the credibility of previous recommendations and mitigates misinformation risks in LLM-generated medical reasoning.

8. Conclusion

This structured framework in medical decision-making ensures a cohesive and logical process by following a sequential reasoning path, where each

step builds upon the previous one. It begins with problem identification, integrates relevant medical knowledge, and systematically analyzes information to enable evidence-based decisions. By assessing potential impacts and ethical considerations, it refines reasoning to ensure contextually appropriate choices. Additionally, it provides supplementary information and follow-up recommendations, supporting both immediate and long-term care strategies.

Unlike unconstrained generative methods, structured outputs follow predefined frameworks, enhancing consistency, accuracy, and interpretability while reducing errors from fragmented recall or misinterpretation. This approach improves clinical reasoning by guiding models through a step-by-step diagnostic process, minimizing critical oversights. It also strengthens explainability and verifiability, ensuring AI recommendations are transparent, traceable, and aligned with medical guidelines, reinforcing LLM reliability.

3.3 Structured Output Generation

Instructional Guidance To ensure the generated content is logical and coherent, we use prompts or guidewords to direct the model step by step in producing content in a specific format. As shown in Figure 8, by clearly specifying the required sections or types of information, the model sequentially fills in the corresponding information, ensuring that the generated content adheres to the expected structure. Additionally, by imposing constraints on the output, the model avoids repetition and redundancy, making the content more concise and intuitive.

Template Generation By predefining structured templates, the model can populate each field of the template step by step. To enhance the comprehensiveness and accuracy of medical question answering, we designed seven key stages of prompts to guide the model in generating structured content.

Stepwise Generation Given the complexity of certain tasks that may exceed the model's context processing capabilities. For example, LLaMA2-7B-HF supports a context length of only 4k tokens, we divide them into multiple stages, with the model outputting structured content one step at a time. Each generation focuses on a specific stage, and through multiple iterations, the model progressively constructs the complete structured content. After generating the content for all stages, the model combine the information to generate a comprehensive long-form response.

3.4 Optimization Strategies

To further enhance the comprehensiveness and credibility of medical question answering, we propose the following optimization strategies:

One-shot Example After the chain of thought, provide an example closely related to the question to help validate the reasoning process. This example should directly guide towards the final answer, ensuring the accuracy and operational effectiveness of the reasoning chain.

Conciseness of the Problem Statement Ensure the problem is presented immediately after the chain of thought, avoiding unnecessary spaces, line breaks, or redundant information to improve input clarity and processing efficiency. Refining the problem statement helps reduce potential misinterpretations by the model and improves response time.

Instruction Reinforcement For critical instructions, repeat them multiple times in the prompt to reinforce their execution. For example, both in the chain of thought and the one-shot example, the "Structured Outputs" tag should be used, and it should be appended again at the end to clearly guide the model's output.

Specialized Markers Use specific markers to extract important content, such as "ANSWER END," to facilitate subsequent processing and analysis.

Content and Word Count Limitation We perform quality checks on each step of the step-by-step reasoning, including content and format checks to remove information that is irrelevant to the question in the model output, so that the intermediate reasoning steps do not affect the quality of the answer in the final step. Ensure that the chain of thought is free from redundancy, with each phase of the reasoning chain kept to a maximum of 200 words to avoid repetitive outputs. The final answer should be limited to 512 tokens, ensuring it is concise and precise.

Content Output Strategy In the first stage, the reasoning chain may output longer structured information (up to 4096 tokens). However, in the second stage, the answer extraction should be limited to 512 tokens to ensure efficiency and focus.

Through the above methods and strategies, we ensure that the structured output for medical question answering is comprehensive, accurate, and meets the high standards required for clinical and research applications.

4 Experiments

We evaluated the effectiveness of Medical Structured Outputs CoT (Med-SoCoT) in medical question answering based on general-purpose LLMs, comparing it with the strong baseline method OLAPH (Jeong et al., 2024a), which was fine-tuned, whereas our approach was training-free. Our experimental setup closely follows previous work to ensure a fair comparison. Additionally, we validated the effectiveness of our method on other open-source models.

4.1 Experimental Setup

We used the MedLFQA benchmark (Jeong et al., 2024a) for evaluation and adhered to existing work by using factuality as the primary evaluation metric. To facilitate reproducibility, we used five datasets in the benchmark: LiveQA, MedicationQA, HealthSearchQA, K-QA Golden, and K-QA Silver.

Comparison Method We compared our approach with the state-of-the-art (SoTA) OLAPH method, which utilizes automated evaluation to generate synthetic preference sets and employs preference optimization through iterative learning.

Model Configuration We applied the structured output method to several popular LLMs, including LLaMA2-7B-HF (Touvron et al., 2023), Mistral-7B (Jiang et al., 2023a), BioMistral-7B (Labrak et al., 2024), Gemma-7B (Team and et al., 2024), GPT-3.5-Turbo. By default, we set the decoding temperature to 0.0 to minimize the randomness in the output of the LLMs. We applied the step-by-step reasoning approach to the LLaMA2-7B-HF, Mistral-7B, BioMistral-7B, and Meditron-7B models, and the direct output approach to the Gemma-7B, LLaMA3.1-3B-INSTRUCT, and GPT-3.5-Turbo models, and reported the Factuality Score.

4.2 Overall Performance

This chapter examines whether general and medical models enhance factual scores via structured output and whether models with limited context length benefit from step-by-step reasoning. Additionally, it analyzes the impact of structured output on text quality, including Words Composition indicators. Below is a detailed analysis:

Table 1 shows that both general models (LLaMA2-7B, Mistral-7B) and medical fine-tuned models (BioMistral-7B, Meditron-7B) significantly improve Factuality Score using the Med-SoCoT

method. This demonstrates the effectiveness of structured output. Notably, Mistral-7B achieves a factual score of 79.3, surpassing other methods.

Table 1 shows that for models with limited context length, step-by-step reasoning enhances both structured output and factuality scores. For example, after applying Med-SoCoT, Mistral-7B achieved high factuality scores, particularly on the HealthSearchQA datasets, with scores of 85.8.

In terms of Words Composition, the Med-SoCoT method outperformed the OLAPH method, showing significant improvements across multiple models. For example, Mistral-7B and BioMistral-7B achieved Words Composition scores of 11.3 and 17.3, respectively, surpassing OLAPH.

4.3 Additional Experiments

According to Table 2, in the additional experiments, we compared the effects of standard Chain-of-Thought (CoT) prompts and structured output prompts, while also evaluating the structured output capabilities of both the 3B small model and general large models. The experimental results show that structured output significantly improves the factuality score while maintaining text quality, for both the open-source Gemma-7B model and the closed-source GPT-3.5-Turbo model. For instance, when Gemma-7B adopted the Med-SoCoT method, its factuality score increased from 56.7 to 85.8, demonstrating a remarkable improvement. Similarly, GPT-3.5-Turbo’s factuality score improved from 77.7 to 83.5 with structured output, further confirming the effectiveness of this approach.

Additionally, for the latest 3B small model, LLaMA3.1-3B-INSTRUCT, its strong instruction-following ability and contextual understanding allow it to generate structured content in a single step without requiring iterative reasoning. Experimental results show that after adopting the Med-SoCoT method, LLaMA3.1-3B-INSTRUCT’s factuality score increased from 80.2 to 84.5, approaching the performance of GPT-3.5-Turbo. This indicates that the small-scale model has strong potential in structured output tasks.

4.4 Ablation Study

Our method employs the Med-SoCoT prompt, and we conducted ablation studies on the MedLFQA dataset using the Gemma-7B model to assess the impact of key steps on performance. The studies focused on the contribution of steps like "Understanding the Question," "Analyzing Medical Infor-

Method	LiveQA		MedicationQA		HealthSearchQA		K-QA Golden		K-QA Silver		Average	
	Words	Fact.	Words	Fact.	Words	Fact.	Words	Fact.	Words	Fact.	Words	Fact.
LLaMA2-7B w/ Zero-shot	7.4	58.1	4.4	48.8	11.0	62.4	6.9	50.4	6.1	40.7	7.2	52.1
LLaMA2-7B w/ OLAPH	7.7	71.2	5.1	57.1	12.0	68.2	8.4	64.3	7.5	47.4	8.1	61.6
LLaMA2-7B w/ Med-SoCoT (Ours)	7.8	76.9	7.1	65.0	13.2	75.1	10.5	72.5	12.2	57.3	10.2	69.4
Mistral-7B w/ Zero-shot	8.5	59.6	5.4	56.6	15.8	78.7	9.8	57.9	8.4	42.8	9.6	59.1
Mistral-7B w/ OLAPH	6.6	66.8	6.3	67.5	13.7	83.8	10.9	75.2	18.2	77.5	11.1	74.2
Mistral-7B w/ Med-SoCoT (Ours)	9.6	80.6	8.2	70.7	14.5	85.8	12.1	80.8	12.2	78.8	11.3	79.3
BioMistral-7B [♣] w/ Zero-shot	4.7	27.4	2.1	12.9	7.0	41.1	7.5	37.0	5.4	27.5	5.3	29.2
BioMistral-7B [♣] w/ OLAPH	13.5	69.1	12.5	70.9	18.4	76.9	17.3	75.6	17.2	59.8	15.8	70.5
BioMistral-7B [♣] w/ Med-SoCoT (Ours)	14.9	71.1	15.5	73.6	20.5	78.7	17.5	79.4	17.8	63.0	17.3	73.2
Meditron-7B [♣] w/ Zero-shot	6.5	49.5	3.7	44.0	7.4	45.7	6.0	44.6	5.5	37.3	5.8	44.2
Meditron-7B [♣] w/ OLAPH	8.0	66.7	5.9	62.7	8.7	50.4	10.4	71.3	11.0	59.6	8.8	62.1
Meditron-7B [♣] w/ Med-SoCoT (Ours)	10.1	69.9	6.4	63.6	9.3	70.3	12.5	72.2	13.4	71.2	10.3	69.4

Table 1: Performance of models (LLaMA2-7B, Mistral-7B, BioMistral-7B, Meditron-7B) on MedLFQA datasets, evaluated on Words Composition (Words) and Factuality Score (Fact.). Models are tested in Zero-shot (no prompt), fine-tuned with OLAPH, and with our Med-SoCoT method. Models with [♣] are medically fine-tuned. The best result for each model is bolded, and average values are in blue.

Method	LiveQA		MedicationQA		HealthSearchQA		K-QA Golden		K-QA Silver		Average	
	Words	Fact.	Words	Fact.	Words	Fact.	Words	Fact.	Words	Fact.	Words	Fact.
Gemma-7B w/ Zero-shot	8.3	55.7	6.0	41.1	12.4	58.0	12.4	40.9	6.1	31.9	9.1	45.5
Gemma-7B w/ CoT	12.7	60.4	6.8	55.7	12.9	69.9	8.6	55.9	8.2	41.3	9.8	56.7
Gemma-7B w/ Med-SoCoT (Ours)	12.7	86.6	10.2	83.0	23.3	92.2	17.9	83.8	21.1	83.4	17.0	85.8
LLaMA3.1-3B-INSTRUCT w/ Zero-shot	10.6	82.1	9.6	76.6	21.2	82.8	15.8	82.7	18.8	69.0	15.2	78.6
LLaMA3.1-3B-INSTRUCT w/ CoT	11.0	82.5	10.7	77.6	20.1	83.7	15.3	82.6	18.9	74.7	15.2	80.2
LLaMA3.1-3B-INSTRUCT w/ Med-SoCoT (Ours)	11.5	86.0	10.9	81.6	21.8	87.6	15.9	84.5	19.2	82.7	15.8	84.5
GPT-3.5-Turbo [♠] w/ Zero-shot	28.8	77.7	38.2	79.2	29.7	74.0	35.6	76.3	36.2	75.7	33.7	76.6
GPT-3.5-Turbo [♠] w/ CoT	29.3	78.2	40.6	80.1	30.9	76.5	37.8	76.6	38.6	77.3	35.4	77.7
GPT-3.5-Turbo [♠] w/ Med-SoCoT (Ours)	31.2	87.2	43.2	80.8	37.6	78.1	45.2	87.6	45.3	83.7	40.5	83.5

Table 2: Performance of models (Gemma-7B, LLaMA3.1-3B-INSTRUCT, GPT-3.5-Turbo) on MedLFQA datasets, evaluated on Words Composition (Words) and Factuality Score (Fact.). Models are tested in Zero-shot (no prompt), CoT (Chain-of-Thought), and with our Med-SoCoT method. Models with [♠] are closed source. The best result for each model is bolded, and average values are in blue.

mation," and "Suggesting Follow-Ups" to the Factuality Score. The results are presented in Tables 5, 6, and 7.

Table 5 demonstrates that removing core steps significantly degrades performance, highlighting their role in improving reasoning and coherence. Table 6 confirms that retaining core steps leads to better results, while Table 7 shows that altering the order of core steps harms performance, emphasizing the need for a logical sequence.

Additionally, Table 8 reveals that removing features like One-shot Examples, Instruction Reinforcement, and Specialized Markers leads to a significant drop in the Factuality Score, highlighting their importance in task understanding and clarity.

In summary, our findings validate that the Med-SoCoT prompt, with its carefully structured steps and technical components, is crucial for optimizing performance and reasoning capabilities. Further experimental data and case studies are available in the Appendix.

5 Conclusion

This paper introduces a method centered on structured outputs to improve the factuality and comprehensiveness of LLMs in medical QA tasks. By leveraging structured prompts and supporting both one-shot reasoning and stepwise reasoning, the method enhances long-form text generation without additional training. Structured outputs significantly improve factual accuracy, enabling smaller models to perform comparably to fine-tuned large models, while stepwise reasoning helps mitigate performance limitations in resource-constrained settings. The structured prompts, covering seven key aspects of medical reasoning, consistently outperform baselines in factuality accuracy and comprehensiveness. This approach demonstrates the potential of structured outputs in optimizing LLM performance across specialized domains, offering a cost-effective alternative to model training with broader applicability in areas like legal and technical documentation.

Limitations

Our Med-SoCoT prompt is an initial step toward optimizing structured outputs in medical reasoning tasks. However, its generalizability to non-medical domains and efficiency in resource-constrained settings remain open questions. Future work could explore more advanced prompt designs and adaptive strategies to enhance performance and applicability across diverse tasks and models.

Ethical Considerations

Ethical considerations are central to our research. In this study, we ensure adherence to ethical principles by exclusively using publicly available datasets and employing models that are open-source or widely accepted within the research community. We emphasize transparency in all stages of our work and prioritize the responsible application of technology, particularly in the sensitive domain of medical reasoning, to ensure that our contributions promote fairness, reliability, and societal benefit.

References

Asma Ben Abacha, Eugene Agichtein, Yuval Pinter, and Dina Demner-Fushman. 2017. Overview of the medical question answering task at trec 2017 liveqa. In *Proceedings of the Text REtrieval Conference (TREC)*, Gaithersburg, MD, USA. NIST.

Asma Ben Abacha, Yassine Mrabet, Mark Sharp, Travis R Goodwin, Sonya E Shooshan, and Dina Demner-Fushman. 2019. [Bridging the gap between consumers' medication questions and trusted answers](#). *Studies in health technology and informatics*, 264:25–29.

Juyeon Bae, JuHee Lee, and et al. 2023. [Development of the clinical reasoning competency scale for nurses](#). *BMC Nursing*, 22(1):138.

Zhijie Bao, Wei Chen, Shengze Xiao, Kuang Ren, Jiaao Wu, Cheng Zhong, Jiajie Peng, Xuanjing Huang, and Zhongyu Wei. 2023. [Disc-medllm: Bridging general large language models and real-world medical consultation](#).

Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#).

S. Marc Breedlove, Neil V. Watson, and Mark R. Rosenzweig. 2010. *Biological Psychology: An Introduction to Behavioral, Cognitive, and Clinical Neuroscience*. Sinauer Associates, Incorporated Publishers.

Salvatore Corrao and Christiano Argano. 2022. [Rethinking clinical decision-making to improve clinical reasoning](#). *Frontiers in Medicine*, 9.

John Duncan. 2013. [The structure of cognition: Attentional episodes in mind and brain](#). *Neuron*, 80(1):35–50.

U.S. Medical Licensing Examination. 2023. [U.s. medical licensing examination \(usmle\) examination description and guidelines](#).

Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [Eli5: Long form question answering](#).

Michael S. Gazzaniga, Richard B. Ivry, and George R. Mangun. 2018. *Cognitive Neuroscience: The Biology of the Mind*, 5th edition. W.W. Norton.

Paul Hager, Friederike Jungmann, Robbie Holland, Kunal Bhagat, Inga Hubrecht, Manuel Knauer, Jakob Vielhauer, Marcus Makowski, Rickmer Braren, Georgios Kaissis, and Daniel Rueckert. 2024. [Evaluation and mitigation of the limitations of large language models in clinical decision-making](#). *Nature Medicine*, 30(9):2613–2622.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#).

Joy Higgs. 2020. *Developing Clinical Reasoning Capabilities*, pages 1–18. Springer Singapore, Singapore.

Minbyul Jeong, Hyeon Hwang, Chanwoong Yoon, Tae-whoo Lee, and Jaewoo Kang. 2024a. [Olaph: Improving factuality in biomedical long-form question answering](#).

Minbyul Jeong, Jiwoong Sohn, Mujeen Sung, and Jaewoo Kang. 2024b. [Improving medical reasoning through retrieval and self-reflection with retrieval-augmented large language models](#).

Albert Q. Jiang, Alexandre Sablayrolles, and et al. 2023a. [Mistral 7b](#).

Jinhao Jiang, Kun Zhou, Zican Dong, Keming Ye, Wayne Xin Zhao, and Ji-Rong Wen. 2023b. [Structgpt: A general framework for large language model to reason over structured data](#).

Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2020. [What disease does this patient have? a large-scale open domain question answering dataset from medical exams](#).

Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William W. Cohen, and Xinghua Lu. 2019. [Pubmedqa: A dataset for biomedical research question answering](#).

Qiao Jin, Zheng Yuan, Guangzhi Xiong, Qianlan Yu, Huaiyuan Ying, Chuanqi Tan, Mosha Chen, Songfang Huang, Xiaozhong Liu, and Sheng Yu. 2022.

720	Biomedical question answering: A survey of approaches and challenges. <i>ACM Computing Surveys</i> , 55(2):1–36.	775
721		776
722		777
723	Damjan Kalajdzievski. 2024. Scaling laws for forgetting when fine-tuning large language models.	778
724		
725	David R. Krathwohl. 2002. A revision of bloom’s taxonomy: An overview. <i>Theory into Practice</i> , 41(4):212–218.	779
726		780
727		781
728	Anastasia Krithara, Anastasios Nentidis, Konstantinos Bougiatiotis, and Georgios Paliouras. 2023. Bioasqqa: A manually curated corpus for biomedical question answering. <i>Scientific Data</i> , 10(1):170.	782
729		783
730		
731		
732	Yanis Labrak, Adrien Bazoge, Emmanuel Morin, Pierre-Antoine Gourraud, Mickael Rouvier, and Richard Dufour. 2024. Biomistral: A collection of open-source pretrained large language models for medical domains.	784
733		785
734		786
735		787
736		
737	Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. Biobert: a pre-trained biomedical language representation model for biomedical text mining. <i>Bioinformatics</i> , 36(4):1234–1240.	788
738		789
739		790
740		791
741		
742	Haoyang Li, Jing Zhang, Cuiping Li, and Hong Chen. 2023. Resdsq: Decoupling schema linking and skeleton parsing for text-to-sql.	792
743		793
744		794
745	Kai Liu, Zhihang Fu, Chao Chen, Wei Zhang, Rongxin Jiang, Fan Zhou, Yaowu Chen, Yue Wu, and Jieping Ye. 2024. Enhancing llm’s cognition via structurization.	795
746		796
747		
748		
749	Zhengliang Liu, Tianyang Zhong, Yiwei Li, Yutong Zhang, Yi Pan, Zihao Zhao, Peixin Dong, Chao Cao, Yuxiao Liu, Peng Shu, et al. 2023. Evaluating large language models for radiology natural language processing. <i>arXiv preprint arXiv:2307.13693</i> .	797
750		798
751		
752		
753		
754	Alejandro Lozano, Scott L Fleming, Chia-Chun Chiang, and Nigam Shah. 2023. Clinfo.ai: An open-source retrieval-augmented large language model system for answering medical questions using scientific literature.	799
755		800
756		801
757		
758		
759	Itay Manes, Naama Ronn, David Cohen, Ran Ilan Ber, Zehavi Horowitz-Kugler, and Gabriel Stanovsky. 2024a. K-qa: A real-world medical qa benchmark.	802
760		803
761		
762	Itay Manes, Naama Ronn, David Cohen, Ran Ilan Ber, Zehavi Horowitz-Kugler, and Gabriel Stanovsky. 2024b. K-qa: A real-world medical q&a benchmark. <i>arXiv preprint arXiv:2401.14493</i> .	804
763		805
764		
765		
766	Isaac K S Ng, Wilson G W Goh, Desmond B Teo, Kar Mun Chong, Li Feng Tan, and Chia Meng Teoh. 2024. Clinical reasoning in real-world practice: a primer for medical trainees and practitioners. <i>Post-graduate Medical Journal</i> , page qgae079.	806
767		807
768		808
769		809
770		
771	OpenAI. 2023. Introducing structured outputs in the api. https://openai.com/index/introducing-structured-outputs-in-the-api/ . Accessed: 2024-10-22.	810
772		811
773		
774		
	Ankit Pal, Logesh Kumar Umapathi, and Malaikanan Sankarasubbu. 2022. Medmcqa : A large-scale multi-subject multi-choice dataset for medical domain question answering.	812
		813
	Thierry Pelaccia, Jacques Tardif, Emmanuel Tribby, and Bernard Charlin. 2011. An analysis of clinical reasoning through a recent and comprehensive approach: the dual-process theory. <i>Medical Education Online</i> , 16.	814
		815
	Jeremy B. Richards and et al. 2020. Teaching clinical reasoning and critical thinking: From cognitive theory to practical application. <i>CHEST</i> , 158(4):1617–1628.	816
	Hyoung Seok Shin. 2019. Reasoning processes in clinical reasoning: from the perspective of cognitive psychology. <i>Korean Journal of Medical Education</i> , 31(4):299–308.	817
		818
		819
	Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. 2023. Large language models encode clinical knowledge. <i>Nature</i> , 673(7748):173–179.	820
		821
	Karan Singhal and et al. 2022. Large language models encode clinical knowledge.	822
		823
	Ivan Stelmakh, Yi Luan, Bhuwan Dhingra, and Ming-Wei Chang. 2023. Asqa: Factoid questions meet long-form answers.	824
		825
	Gemma Team and et al. 2024. Gemma: Open models based on gemini research and technology.	
	P. Thagard. 1996. <i>Mind: Introduction to Cognitive Science</i> . Bradford Books. MIT Press.	
	Harish Thampy, Emma Willert, and Subha Ramani. 2019. Assessing clinical reasoning: Targeting the higher levels of the pyramid. <i>Journal of General Internal Medicine</i> , 34(8):1631–1636.	
	Hugo Touvron, Louis Martin, and et al. 2023. Llama 2: Open foundation and fine-tuned chat models.	
	Tao Tu and et al. 2023. Towards generalist biomedical ai.	
	Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference.	
	Guangzhi Xiong, Qiao Jin, Zhiyong Lu, and Aidong Zhang. 2024. Benchmarking retrieval-augmented generation for medicine.	
	Meredith E. Young, Alik Thomas, and et al. 2020. Mapping clinical reasoning literature across the health professions: a scoping review. <i>BMC Medical Education</i> , 20(1):107.	
	Wayne Xin Zhao, Kun Zhou, and et al. 2024. A survey of large language models.	

A Evaluation Index Calculation Method

In this section, we describe the formulas used to compute various evaluation metrics.

A.1 Words Composition (Words Composition)

The Words Composition metric is calculated as the average of ROUGE-1, ROUGE-2, and ROUGE-L F1 scores:

$$\text{Words Composition} = \frac{\text{ROUGE-1 F1} + \text{ROUGE-2 F1} + \text{ROUGE-L F1}}{3}$$

A.2 Hallucination Score (Hallucination Score)

The Hallucination Score measures the proportion of elements in set S that are contradicted by the prediction \hat{P} :

$$\text{Hallucination Score}(\hat{P}) = \frac{|\{x \in S \mid \hat{P} \text{ contradicts } x\}|}{|S|}$$

A.3 Comprehension Score (Comprehensiveness Score)

The Comprehension Score evaluates the proportion of elements in the medical hypothesis set MH that are entailed by the prediction \hat{P} :

$$\text{Comprehensiveness Score}(\hat{P}) = \frac{|\{x \in MH \mid \hat{P} \text{ entails } x\}|}{|MH|}$$

A.4 Factual Score (Factuality Score)

The Factual Score combines the Comprehension Score and the Hallucination Score to produce a balanced metric, normalized to a scale of 0 to 100:

$$\text{Factuality Score} = \frac{(\text{Comprehensiveness Score} - \text{Hallucination Score}) + 100}{2}$$

B Experiment Settings

Dataset	Format	# of QA pairs	# of Ambiguous Questions	Avg. Length of Answers	Avg. # of MH statements	Avg. # of NH Statements
LiveQA (Abacha et al., 2017)	(Q, A, MH, NH)	100	4	82.8	2.6	2.3
MedicationQA (Abacha et al., 2019)	(Q, A, MH, NH)	666	24	55.5	2.6	2.3
HealthSearchQA (Singhal et al., 2023)	(Q, A, MH, NH)	3,077	96	118.8	2.6	2.3
K-QA Golden (Manes et al., 2024b)	(Q, A, MH, NH)	201	1	88.5	4.4	3.5
K-QA Silver (Manes et al., 2024b)	(Q, A, MH, NH)	904	106	99.9	2.4	2.0

Table 3: Overview of MedLFQA datasets with formats and key statistics.

C Model Settings

Model	Method	Context Length	Type
LLaMA2-7B-HF	STEPWISE	512	General
Mistral-7B	STEPWISE	512	General
BioMistral-7B	STEPWISE	512	Medical specific
Meditron-7B	STEPWISE	512	Medical specific
Gemma-7B	DIRECT	2048	General
LLaMA3.1-3B-INSTRUCT	DIRECT	4096	General
GPT-3.5-Turbo	DIRECT	4096	General

Table 4: Specific Model Settings Adopted by the Med-SoCoT Method

D Ablation Study

Gemma-7B	Step Removed	Factuality Score	Δ
Baseline (7-step Full)	None	71.6	-
w/o Step 1	Understand the Question	66.5	↓ 5.1 (7.1%)
w/o Step 2	Recall Medical Knowledge	70.4	↓ 1.2 (1.7%)
w/o Step 3	Analyze Medical Info	65.0	↓ 6.6 (9.2%)
w/o Step 4	Assess Impacts	69.8	↓ 1.8 (2.5%)
w/o Step 5	Add Relevant Info	70.0	↓ 1.6 (2.2%)
w/o Step 6	Suggest Follow-Ups	66.0	↓ 5.6 (7.8%)
w/o Step 7	Reference Sources	70.5	↓ 1.1 (1.5%)

Table 5: Ablation Study on the Importance of Each Step in the Med-SoCoT Using the Gemma-7B on the MedLFQA.

Gemma-7B	Steps Retained	Factuality Score	Δ
Baseline (7-step Full)	All Steps	71.6	-
Step 1 + Step 3 + Step 6	Core Steps Combination	60.5	↓ 11.1 (15.5%)
Step 1 + Step 3	Only Two Core Steps	55.0	↓ 16.6 (23.2%)
Step 2 + Step 4 + Step 5	Non-Core Steps Combination	58.7	↓ 12.9 (18.0%)

Table 6: Ablation Study on the Effectiveness of Different Step Combinations in the Med-SoCoT Using the Gemma-7B Model on the MedLFQA.

Gemma-7B	Step Order Adjusted	Factuality Score	Δ
Baseline (7-step Full)	Standard Order	71.6	-
Step 3 ↔ Step 6	Swap Step 3 and Step 6	70.5	↓ 1.1 (1.5%)
Step 1 ↔ Step 4	Swap Step 1 and Step 4	70.8	↓ 0.8 (1.1%)
Step 5 ↔ Step 7	Swap Step 5 and Step 7	71.2	↓ 0.4 (0.6%)

Table 7: Ablation Study on the Impact of Step Order Adjustments in the Med-SoCoT Using the Gemma-7B Model on the MedLFQA.

Gemma-7B Prompt Version	Factuality Score	Δ
Baseline (Med-SoCoT)	69.4	-
w/o One-shot Example	65.2	↓ 6.0%
w/o Instruction Reinforcement	66.8	↓ 3.7%
w/o Specialized Markers	67.5	↓ 2.7%
w/o All Features	60.3	↓ 13.1%

Table 8: Ablation Study Results of Gemma-7B on Med-SoCoT Prompt Variations, Evaluating the Impact of One-shot Examples, Instruction Reinforcement, and Specialized Markers on Factuality Score Performance.

Reasoning without prompt
Question: I take zyrtec for my allergies will hydroxyzine be ok to take with it?
Answer: Can Zyrtec (cetirizine) be combined with hydroxyzine? Absolutely! In general, cetirizine for the histamine issues and hydroxyzine you like probably for the itching component. Can I mix Benadryl and hydroxyzine?
Factuality Score: 27.4 ✖

Reasoning with OLAPH
Question: I take zyrtec for my allergies will hydroxyzine be ok to take with it?
Answer: Hydroxyzine is a diphenhydramine. It is an antihistamine that reduces the effects of natural chemical histamine in the body. Hydroxyzine is used to treat or prevent the symptoms of an allergic reaction. ... Zyrtec is a brand name for the drug cetirizine. It is an antihistamine that reduces the effects of natural chemical histamine in the body. Zyrtec is used to treat or prevent the symptoms of an allergic reaction. Zyrtec is also used to treat the itching and swelling of a skin rash caused by a condition called urticaria.
Factuality Score: 58.4 ✖

Figure 6: Dataset: K-QA Silver (Manes et al., 2024b), Model: Mistral-7B (Jiang et al., 2023a) (*Upper*) LLM directly answers questions, resulting in a large number of hallucinations (*Lower*) After training with the OLAPH framework, although it can answer relevant knowledge, it cannot provide accurate answers.

Reasoning with Structured Outputs (Ours)
<p>Question: I take zyrtec for my allergies will hydroxyzine be ok to take with it?</p> <hr/> <p>Structured Output:</p> <p>1. Understand the Question: Zyrtec is an antihistamine used to treat allergic reactions, while Hydroxyzine is also an anti-histamine but used for anxiety and insomnia. Both medications can cause drowsiness, so it is important to be aware of potential interactions and side effects.</p> <p>2. Recall Relevant Medical Knowledge: Hydroxyzine is an antihistamine that is used to treat allergic reactions, anxiety, insomnia, nausea, vomiting, ... , watery eyes, cough, wheeziness, asthma, hay fever, insect bites, stings, allergens, ... , infections, viruses, bacteria, fungi, parasites, toxins, poisons, venoms, snake bit, spider bit.</p> <p>3. Analyze Medical Information: Zyrtec is a brand name for cetirizine, an antihistamine used for allergies. Zyrtec is also used as a sedative, hypnotic, antidepressant, and anti-anxiety medication. Hydroxyzine and Zyrtec are both antihistamines. They are used in different ways. The main difference between hydroxyzine and cetirizine is that hydroxyzine is more sedating than cetirizine. Cetirizine has less sedation than hydroxyzine. Both hydroxyzine and cetirizine are antihistamines. Antihistamines are drugs that block the effects of histamine. ...</p> <hr/> <p>Answer: Zyrtec is one of the most popular anticholinergic drugs used today. This drug is often used by people who suffer from allergies. However, this drug can also be used with other drugs. One of these drugs is hydrocortisone. When you take this medication, you will experience a lot of side-effects. These side effect include dizziness and drowsy. You should not take hydrocodone with Zyrtec. If you are taking this medicine, then you should consult your doctor first.</p> <hr/> <p>Factuality Score: 100.0 ✓</p>

Figure 7: Dataset: K-QA Silver (Manes et al., 2024b), Model: Mistral-7B (Jiang et al., 2023a) Use structured outputs and answer questions that include all [correct statements](#).

Prompt

Medical Structured Output Prompt (part1):

Task: You are a helpful assistant. Step-by-Step Thinking for Structured Medical Question Answering.

General Instructions: - Generate detailed and structured medical responses based on the given medical question. Answers should be grounded in current medical knowledge, covering all key aspects of the question.

- Ensure the answer includes background, etiology, symptoms, diagnosis, treatment, and prevention.
- The answer should be logically organized and provide accurate, comprehensive medical information.
- ## Task Instructions: - Generate a comprehensive response based on the input question. The response should cover everything from background information to diagnosis and treatment recommendations, ensuring a structured and coherent output.
- The answer should address as many aspects of the medical question as possible, considering risk factors, complications, and related medical conditions.
- Consider the relationship between diseases and medications.
- Do not output duplicate content.
- Each thought process should not exceed 200 words.

Output Structure: - The output should follow the structured template below to ensure the completeness and professionalism of the medical response.

- Please ensure that the output contains Long-Form Answer

Figure 8: Medical Structured Output Prompt (part1).

Prompt

Medical Structured Output Prompt (part2):

Chain of Thought:

1. Understand the Question:

- Explain the background or definition of the medical issue. Provide a brief description of basic concepts and possibly affected systems or organs.
- Identify and define key medical terms and concepts.
- Clarify the specific information or details requested.

2. Recall Relevant Medical Knowledge: - Retrieve information related to the disease, medication, or procedure.

- Consider anatomy, physiology, pathology, pharmacology, and current medical guidelines.

3. Analyze Medical Information: - Combine 1. understanding the question and 2. relevant medical knowledge to connect the issue with pertinent medical knowledge using clinical reasoning.

- Consider possible explanations, mechanisms, or interactions.

4. Assess Impacts and Considerations:

- Evaluate any risks, side effects, or contraindications.
- Consider specific patient factors (age, comorbidities, allergies).

5. Provide Additional Relevant Information:

- Include important details that help in understanding.
- Mention any exceptions, alternative options, or preventive measures.

6. Suggest Follow-Up Steps or Actions:

- If necessary, recommend consulting a healthcare professional.
- Advise on monitoring, follow-up, or further evaluation.

7. Reference Reliable Sources:

- Base responses on evidence from authoritative medical texts or guidelines.
- Cite clinical studies, professional organizations, or regulatory agency information.

8. Long-Form Answer:

- Combine the above reasoning to accurately and comprehensively answer the question. Provide a "long-form answer" that contains 400-500 words. The word count must not be less than 400 words.

END

-Please end the output here.

Figure 9: Medical Structured Output Prompt (part2).

Prompt

Medical Structured Output Prompt (one-shot):

Please refer to the following questions, along with examples of chain of thought and long-form answers.
Question: What is the relationship between Noonan syndrome and polycystic renal disease?

Chain of Thought:

1. Understand the Question: Noonan syndrome is a genetic disorder characterized by distinct facial features, short stature, heart defects, and developmental delays. It affects multiple systems in the body, including the cardiovascular, musculoskeletal, and endocrine systems. Polycystic renal disease, particularly autosomal dominant polycystic kidney disease (ADPKD), is a genetic condition leading to the formation of numerous cysts in the kidneys, resulting in kidney enlargement and impaired function. The question seeks to explore the potential relationship between these two conditions, particularly any shared genetic or pathological mechanisms.

2. Recall Relevant Medical Knowledge: Noonan syndrome is primarily caused by mutations in genes involved in the RAS-MAPK signaling pathway, particularly the PTPN11 gene. It affects approximately 1 in 1,000 to 1 in 2,500 births. On the other hand, polycystic renal disease is commonly caused by mutations in the PKD1 or PKD2 genes. ADPKD has a prevalence of about 1 in 400 to 1 in 1,000. Understanding the genetic basis and clinical manifestations of both conditions is crucial for identifying potential links between them.

3. Analyze Medical Information: The relationship between Noonan syndrome and polycystic renal disease may stem from shared genetic pathways or phenotypic associations. Some studies suggest that patients with Noonan syndrome exhibit renal anomalies, including renal agenesis or structural abnormalities, although true polycystic kidney disease is less commonly reported. This indicates a potential overlap in genetic vulnerabilities that could lead to renal pathologies in Noonan syndrome patients. The mechanisms may involve disruptions in cellular signaling pathways that are pivotal for kidney development and function.

4. Assess Impacts and Considerations: Patients with Noonan syndrome may have additional comorbidities that can influence renal health, such as hypertension or congenital heart defects, which could complicate the presentation of renal disease. Conversely, individuals with polycystic kidney disease are at risk of hypertension and kidney failure, potentially impacting their overall health and necessitating careful monitoring. Genetic counseling may be beneficial for families with a history of either condition to better understand the risks and implications of genetic inheritance.

5. Provide Additional Relevant Information: While there is limited direct evidence linking Noonan syndrome and ADPKD, awareness of renal complications in Noonan syndrome patients is important for clinicians. Furthermore, certain genetic syndromes may predispose individuals to multiple anomalies, making regular screenings for renal function essential in affected individuals.

6. Suggest Follow-Up Steps or Actions: For individuals diagnosed with Noonan syndrome, it is advisable to perform regular renal assessments, including ultrasound examinations to check for any renal structural anomalies. Genetic counseling can provide insights into the risks of polycystic kidney disease and the implications for family planning. Patients should be educated on signs of kidney dysfunction, such as changes in urination patterns, hypertension, or abdominal pain.

7. Reference Reliable Sources: Sources for this information include clinical guidelines from the National Kidney Foundation, the American Academy of Pediatrics, and recent genetic studies published in peer-reviewed journals regarding the genetics of Noonan syndrome and polycystic kidney disease.

8. Long-Form Answer: Noonan's syndrome is an eponymic designation that has been used during the last 8 years to describe a variable constellation of somatic and visceral congenital anomalies, which includes groups of patients previously referred to as male Turner's, female pseudo-Turner's and Bonnevie-Ullrich syndromes. It is now recognized that both sexes may show the stigmas of this condition and, unlike Turner's syndrome, there is no karyotype abnormality although there is often a familial pattern. The most commonly observed anomalies include webbing of the neck, hypertelorism, a shield-shaped chest and short stature. Congenital heart disease, principally pulmonary stenosis, and sexual infantilism often with cryptorchidism in the male subject are additional associated anomalies in this syndrome. Renal anomalies have been described rarely and usually consist of rotational errors, duplications and hydronephrosis. We report the first case of an infant who displayed many of the stigmas of Noonan's syndrome and also showed early evidence of frank renal failure secondary to renal dysplasia with cystic disease.

END

Figure 10: Text Generation Prompt of Med-SoCoT (one-shot).