

Reasoning Like a Doctor: Improving Medical Dialogue Systems via Diagnostic Reasoning Process Alignment

Kaishuai Xu¹, Yi Cheng^{1*}, Wenjun Hou^{1,2*}, Qiaoyu Tan³, Wenjie Li¹

¹The Hong Kong Polytechnic University, HKSAR, China

²Southern University of Science and Technology, Shenzhen, China

³New York University Shanghai, Shanghai, China

{kaishuaii.xu, alyssa.cheng}@connect.polyu.hk, houwenjun060@gmail.com
qiaoyu.tan@nyu.edu, cswjli@comp.polyu.edu.hk

Abstract

Medical dialogue systems have attracted significant attention for their potential to act as medical assistants. Enabling these medical systems to emulate clinicians' diagnostic reasoning process has been the long-standing research focus. Previous studies rudimentarily realized the simulation of clinicians' diagnostic process by fine-tuning language models on high-quality dialogue datasets. Nonetheless, they overly focus on the outcomes of the clinician's reasoning process while ignoring their internal thought processes and alignment with clinician preferences. Our work aims to build a medical dialogue system that aligns with clinicians' diagnostic reasoning processes. We propose a novel framework, EMULATION, designed to generate an appropriate response that relies on abductive and deductive diagnostic reasoning analyses and aligns with clinician preferences through thought process modeling. Experimental results on two datasets confirm the efficacy of EMULATION. Crucially, our framework furnishes clear explanations for the generated responses, enhancing its transparency in medical consultations.¹

1 Introduction

Medical dialogue systems as a fundamental tool in facilitating effective healthcare consultations have garnered sustained attention in recent years (Zeng et al., 2020; Liu et al., 2022b). Especially with the rise of large language models (LLMs), these systems have shown promising potential to act as virtual medical assistants that can aid clinicians in accurate and efficient diagnosis (Bao et al., 2023; Chen et al., 2023; OpenAI, 2023). In this research area, how to enable these medical systems to emulate clinicians' diagnostic reasoning process has been the long-standing research focus ever since

*Equal Contribution.

¹Our codes are available at <https://github.com/kaishxu/Emulation>.

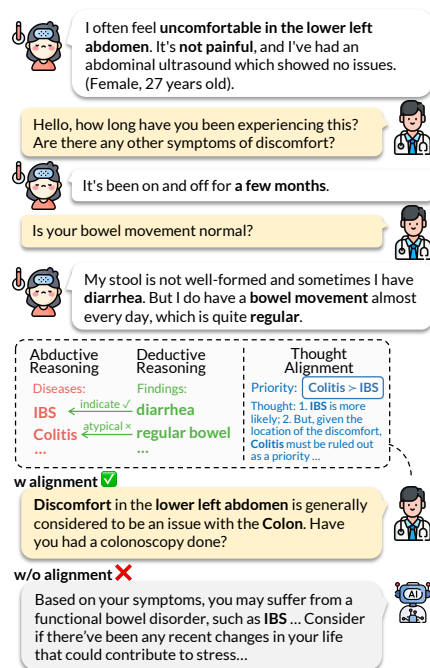


Figure 1: An example of a medical dialogue. **IBS** is the abbreviation for irritable bowel syndrome.

its inception (Holyoak and Morrison, 2005). To achieve this, previous research has paid tremendous efforts in constructing high-quality datasets of medical consultation dialogues, based on which they fine-tuned language models, rudimentarily realizing the simulation of clinicians' diagnostic process that can inquire about symptoms and make a diagnosis (Bao et al., 2023; Chen et al., 2023; Liu et al., 2022b; Xu et al., 2023).

Nonetheless, existing research remains focused on the "output" of the clinician's reasoning process (e.g., the symptoms and diseases mentioned in clinicians' utterances), still neglecting the clinician's internal thought process and decision-making mechanisms. As a result, available systems overly rely on the co-occurrence patterns in the training data, prone to inquire about the most frequent symptoms and diagnose the most common disease. In fact,

a real clinician’s thought process goes far beyond this. In clinical medicine, there is a concept known as “Clinician Preference” (Silverman et al., 2016; Holyoak and Morrison, 2005; Murray et al., 2007). It refers to the inclinations that healthcare professionals exhibit during diagnostic reasoning. It has been widely acknowledged that clinician preferences are affected by many factors (Silverman et al., 2016), much more than the symptom and disease frequency. For example, in Figure 1, the clinician prefers to discuss examinations in the next response to rule out colitis rather than consider factors that cause irritable bowel syndrome. This is because the location of the discomfort typically indicates a more severe issue in the colon. For the existing systems that are constructed purely through dialogue data fine-tuning, it is difficult to capture such subtle reasoning preferences (Xiong et al., 2023).

In this paper, we aim to develop a medical dialogue system that can align with the internal diagnostic reasoning process of clinicians. To this end, we must first model a diagnostic analysis process. For the multi-turn medical dialogue, clinicians generally adopt an iterative, abductive, and deductive analysis that discovers an explanation of the patient’s condition and evaluates the effectiveness of the explanation (Holyoak and Morrison, 2005). This analysis provides a robust and comprehensive foundation for accurate diagnoses. Then, we need to align response generation with the clinician preference based on the analysis and dialogue context. The thought process of how clinicians reason and generate responses is a vital resource for learning their preferences. Extracting the thought process and modeling it together with responses can help learn the diagnostic reasoning process.

Based on the above motivation, we propose a novel medical dialogue system framework, EMULATION, which emulates clinicians’ diagnostic reasoning processes to generate an appropriate response that relies on ample diagnostic analysis and aligns with clinician preferences in consultation. First, an abductive reasoning module investigates potential diseases that can explain a patient’s condition. Then, a deductive reasoning module comprehensively analyzes the relation between clinical findings and potential diseases. Finally, the thought alignment module adjusts the potential disease priority that may be discussed next and generates thought processes that align with the clinician preference based on the above analysis. To learn the clinician preference, we build a new diagnostic

thought process dataset with the help of an LLM.

Our key contributions are outlined as follows: (1) We propose a novel medical dialogue system framework, EMULATION, that emulates clinicians’ diagnostic reasoning processes and aligns with clinician preferences. This is the first work that explores clinician preferences and internal thought processes. (2) We build a diagnostic thought process dataset that is employed to align response generation with clinician preferences. (3) Experimental results demonstrate the effectiveness and explainability of EMULATION.

2 Preliminary

2.1 Problem Formulation

In our work, we conceptualize a medical dialogue as a sequence $U = \{(U_k^P, U_k^D)\}_{k=1}^T$, where each pair (U_k^P, U_k^D) comprises an utterance from a patient followed by an utterance from a doctor. Each doctor’s utterance is annotated with a list of diseases $E_t = \{e_i\}$ that could be relevant to the patient’s condition mentioned in the dialogue. Given a dialogue history $U_t = \{U_1^P, U_1^D, \dots, U_t^P\}$ up to the t -th patient utterance, our system’s objective is to generate a contextually appropriate and medically informed t -th doctor’s utterance U_t^D .

2.2 Diagnostic Reasoning

Diagnostic reasoning involves a detailed examination of how doctors think and make decisions in medicine (Holyoak and Morrison, 2005). It serves as a foundational element for advanced cognitive activities, such as formulating diagnostic conclusions and grasping the underlying pathology of diseases. The framework of diagnostic reasoning has been a focal point of medical cognition research (Patel and Ramoni, 1997). A widely recognized perspective holds that diagnosis represents an iterative, abductive, and deductive process of formulating and evaluating potential explanations for a patient’s abnormal condition (Elstein et al., 1978; Holyoak and Morrison, 2005). We summarize this perspective into two reasoning processes: **Abductive Reasoning** aims to create a plausible diagnosis to explain observed clinical findings; **Deductive Reasoning** further tests the available diagnoses by determining whether the findings support, refute, or are unrelated to the diagnoses. In our method, the first process efficiently explores a disease knowledge base to identify several possible diseases that explain the patient’s condition in the current turn of

the conversation. The second process comprehensively inspects the relationship between the clinical findings and possible diseases. These two processes are conducted iteratively across the conversation. Besides, the clinician preference for medical decision analysis plays an essential role in the diagnostic reasoning process (Silverman et al., 2016). This preference highlights behaviors and actions in medical conversations unique to each clinician, which LLMs may lack. It is the nuanced preference difference that distinguishes a superior clinician from other clinicians (Yu et al., 2024). We model a **Thought Alignment** process to align response generation with the general clinician preference in consultation.

2.3 Disease Annotation

Given GPT-4’s demonstrated effectiveness in several medical licensing examinations (OpenAI, 2023), we employ it to annotate potential diseases relevant to the patient’s condition automatically. We implement two ways to generate lists of diseases considered **before** and **after** examining the current doctor’s response. For the first list, we construct a prompt with the dialogue history U_t , utilizing the model’s diagnostic skills to identify potential diseases. For the second list, we create a prompt that incorporates both the history U_t and the ground truth response U_t^D , enabling the model to deduce the diseases that doctors might discuss. Then, we link the inferred diseases with those in an external medical knowledge base and obtain two lists of potential diseases, E_t^{pri} and E_t^{post} . We merge two lists as one E_t for each doctor’s response. The details are described in the Appendix A.2.

3 Method

In this section, we introduce a diagnostic reasoning framework, which analyzes patient conditions through abductive reasoning and deductive reasoning (Holyoak and Morrison, 2005), and aligns responses with the clinician preferences. As shown in Figure 2, our framework includes three modules. The *Abductive Reasoner* (§3.1) first generates potential diseases based on the new clinical findings of each turn. Then, the *Deductive Reasoner* (§3.2) comprehensively analyzes the relation between potential diseases and new clinical findings. Finally, the *Thought Alignment* (§3.3) module adjusts the disease priority that might be discussed in subsequent dialogues and performs thought process

modeling to generate the next response.

Given the medical dialogue history U_t at the t -th conversation turn, we first use a pre-trained LLM to summarize new clinical findings from the recent utterances U_{t-1}^D and U_t^P . The clinical findings are presented using the SOAP note (Cameron and Turtle-Song, 2002), a documentation technique utilized by healthcare professionals for recording notes (e.g., chief complaint) in a patient’s medical record. For example, we provide the LLM with a prompt to “*summarize the clinical findings from the recent conversation between doctor and patient, adhering to the structure of the SOAP note*”. Each of the summarized clinical findings is in the phrase format and merged into a set $S_t = \{s_j\}_{j=1}^m$, where m represents the number of clinical findings at the t -th conversation turn.

3.1 Abductive Reasoner

The Abductive Reasoner aims to find the most possible diseases that explain the patient’s abnormal condition. Previous studies rely solely on the diagnosis capability of black-box language models or focus on limited diseases in one dataset. We design a two-step pipeline to achieve a comprehensive and explainable abduction. The first step employs a disease retriever to narrow down the scope of the disease. Then, the second step leverages a pre-trained LLM to generate a potential disease list with detailed explanations based on the clinical findings and external medical knowledge.

Disease Retrieval. We concatenate all findings as the query and retrieve disease documents that are relevant to the findings from an external medical knowledge base. A dense retriever with a BERT encoder is applied to collect relevant documents. The sequences of the query and disease document, each prefixed with a “[CLS]” token, are fed independently into the encoder. The hidden state corresponding to the “[CLS]” token from each sequence is chosen to serve as their respective representations $\mathbf{h} \in \mathbb{R}^d$ and $\mathbf{h}_{doc} \in \mathbb{R}^d$. We compute relevance scores through dot product as $\mathbf{h}^\top \cdot \mathbf{h}_{doc}$. The diseases corresponding to the top- K documents ranked by the scores are considered candidate diseases $\hat{E}_t^0 = \{e_i\}_{i=1}^K$. Our approach utilizes Contrastive Learning (Gao et al., 2021) to train the retrieval model. The loss function applied to the

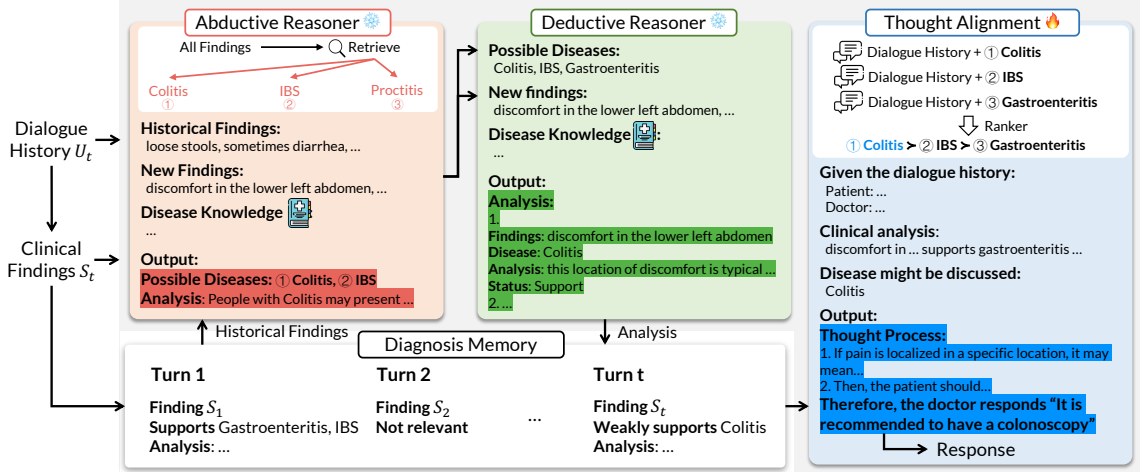


Figure 2: Illustration of our Framework EMULATION. The symbols of a snowflake or a flame represent the module mainly operating in a prompt-based way or having undergone fine-tuning, respectively.

training is specified as follows:

$$\mathcal{L}_{CL} = -\log \frac{\exp(\mathbf{h}^\top \cdot \mathbf{h}_{doc}^+)}{\exp(\mathbf{h}^\top \cdot \mathbf{h}_{doc}^+) + \sum_C \exp(\mathbf{h}^\top \cdot \mathbf{h}_{doc}^-)}, \quad (1)$$

where \mathbf{h}_{doc}^+ and \mathbf{h}_{doc}^- denotes document representations of the relevant diseases in this turn E_t and the irrelevant ones in the knowledge base \mathcal{C} . We adopt in-batch negative sampling for training.

Diagnosis Refinement. After the retrieval, we have assembled a list of potential diseases; however, this approach unavoidably results in the inclusion of diseases irrelevant to the patient’s condition. Furthermore, the disease retrieval, based on vector similarity, can not elucidate the rationale behind the presence of specific diseases within the potential list. In pursuit of an explainable diagnosis, we utilize a pre-trained LLM to identify possible diseases from this list, offering an explanation grounded in clinical findings. The diagnosis knowledge of each disease is provided since it can appropriately reduce the hallucination issue of LLMs. For instance, we prompt the LLM with “*given the findings and knowledge, select possible diseases that can explain the new clinical findings while satisfying past findings*”. To achieve a stable disease identification, we adopt a majority vote inspired by Wang et al. (2023). Specifically, we randomly divide the list into batches of the same size (since the maximum input tokens of LLMs limit the diseases and knowledge that can be input). Such division is repeated several times to construct different batch groups.

We utilize batch groups to generate several refined disease lists and calculate the voting score of each disease as follows:

$$v(e_i) = \sum_{j=1}^B \mathbb{1}(e_i \in \hat{E}_t^j), e_i \in \hat{E}_t^0, \quad (2)$$

where B stands for the number of batch groups, and \hat{E}_t^j represents the j -th refined list. The final refined list includes diseases with a voting score exceeding $B/2$, denoted as $\hat{E}_t' = \{e_i\}_{i=1}^{K'}$.

3.2 Deductive Reasoner

While the Abductive Reasoner has offered an analysis of the relationship between potential diseases and new clinical findings, this analysis predominantly aims to confirm a disease rather than exclude one or sift through information irrelevant to the current diagnosis. It is crucial for clinicians to thoroughly examine the information presented in the dialogue history. However, previous studies often overlook these non-affirmative analyses in their diagnostic processes, which is not conducive to improving diagnostic accuracy.

We introduce the Deductive Reasoner to evaluate the affirmative and non-affirmative relations between new clinical findings and possible diseases. Given the clinical findings S_t and the refined disease list \hat{E}_t' , this reasoner applies a pre-trained LLM with a prompt as “*analyze if the new clinical findings support, oppose, or are irrelevant to the possible diseases*”. We improve the accuracy of analysis by incorporating additional diagnostic knowledge of each disease. The content of specific tags, i.e., findings, disease, and status, is extracted

as diagnosis memory M_t for subsequent response generation.

3.3 Thought Alignment

The clinician preference in consultation is critical, as it differentiates an experienced clinician from a novice clinical student. Prior research has always concentrated on diagnosing the most likely disease and discussing extra symptoms or treatment based on that diagnosis, which does not accurately reflect the dynamics of an actual healthcare consultation. To address this, we have developed the Thought Alignment module, which is designed to adapt preferences akin to that of an expert clinician. The module first prioritizes diseases to be discussed in the subsequent dialogue. Then, it models the thought process and response of real clinicians to align with clinician preferences.

Disease Priority Alignment We implement a disease ranker to establish the priority of each disease within the refined disease list \hat{E}'_t . The backbone of this ranker is the BERT encoder. For the input, we combine the dialogue history with each potential disease, formatting it as follows: “[CLS] {history} the next response will discuss: {disease}”. The relevance score of the dialogue history to the discussed disease is calculated:

$$r(U_t, e_i) = \text{MLP}(\text{repr}(U_t; e_i)), e_i \in \hat{E}'_t, \quad (3)$$

where $\text{repr}(\cdot)$ extracts the hidden state of the “[CLS]” token as the representation, and $\text{MLP}(\cdot)$ projects the representation to a scalar. The ranking model is also trained using a contrastive loss function. To better adjust the priority within the refined disease list, the selection of negative samples includes the union of this refined list \hat{E}'_t and the pre-annotated disease list E_t^{pri} . The positive diseases are from E_t^{post} . We compute relevance scores between the dialogue history and all diseases in the refined list and then reorder the list to \hat{E}''_t .

Thought Process Alignment Thought processes are a reflection of clinician preferences in consultation. Motivated by studies on distilling multi-step reasoning capabilities (Fu et al., 2023; Chae et al., 2023), we design a thought distillation method to empower a model that can generate thought processes consistent with clinician preferences. Specifically, we first leverage a pre-trained LLM to deduce a plausible thought process for each doctor’s response. The prompt to the LLM is organized

Description	MedDG	KaMed
# Thought processes	122K	285K
Avg. # steps	4.17	4.13
Avg. # tokens per step	40.98	41.21
Avg. # total tokens	236.39	239.53

Table 1: Statistics of MEDIATOR

as “complete the thought process based on the dialogue context”. We augment the prompt with 3 human-annotated thought processes in the Chain-of-Thought (CoT) format. The thought extraction can be defined as:

$$Y_t \sim \text{LLM}(Y_t|U_t, U_t^D, \text{prompt}). \quad (4)$$

Then, we gather the thought process for each dialogue turn and develop a thought process alignment model using the autoregressive language modeling approach as our training objective. The loss function can be defined as:

$$\mathcal{L}_G = -\sum_j \log p(y_{jt}|U_t, M_{\leq t}, E_t^{post}, Y_{t,<j}), \quad (5)$$

where $M_{\leq t}$ denotes the diagnosis memory until the t -th turn and E_t^{post} denotes the diseases discussed in the next response. During the inference, we utilize the top- K'' diseases in the list \hat{E}''_t instead of E_t^{post} and their related analyses to generate the thought process. The final response is extracted from “Therefore, the doctor responds {response}”.

4 Thought Process Corpus

In this section, we present the medical dialogue thought process corpus, MEDIATOR, where each dialogue turn is annotated with a chain-of-thought reasoning path. We utilize the powerful reasoning ability of GPT-4 and its impressive professional skills in the medical domain for automatic annotation. The annotated reasoning path is expected to reflect the thought process of clinicians who analyze the patient’s condition and determine what to discuss. Two datasets, MedDG (Liu et al., 2022b) and KaMed (Li et al., 2021), are annotated following the few-shot prompt in §3.3. As shown in Table 1, the former is annotated with 122K thought processes and the latter with 285K. Each thought process includes a reasoning path consisting of approximately four steps. An example of the generated thought process is displayed in Table 2. The doctor investigates the new clinical findings (i.e., discomfort under certain conditions) and plans to discuss allergy issues after multi-step reasoning.

Patient: What’s going on with my throat feeling itchy and making me want to gag? (Female, 22 years old)
Doctor: How long have you been experiencing this?
Patient: About half an hour. And my temperature is just under 37 degrees.
Doctor: Have you had pharyngitis before?
Patient: No, but I have chronic gastritis.
Doctor: So, this hasn’t happened before?
Patient: Only when I have a cold, I would feel itchy and cough. Also, I tend to feel uncomfortable when the seasons change. Other than that, no.

Thought Process:

1. The patient mentions symptoms of itchy throat and gagging during colds or when the seasons change. This indicates respiratory issues.
2. The doctor considers the possibility of an allergic reaction, as ... Seasonal changes frequently introduce variations in allergens and irritants in the environment, which might explain the patient’s symptoms.
3. The patient’s symptoms are not frequent, but appear to be triggered by specific conditions or events. This further supports the allergies or environmental causes.
4. The patient mentions having chronic gastritis, but the relation to the current symptoms is unclear ... but usually does not cause an itchy throat. So, confirming allergy issues is a priority.

Therefore, the doctor responds, "Have you done any cleaning today? Or been exposed to dust mites?"...

Table 2: An example of the generated thought process.

Datasets	Knowledge	Consistency	Rationality
MedDG	90%	89%	82%
KaMed	87%	91%	79%

Table 3: Human evaluation results of automatically generated thought processes. Values in the table represent average valid percentages.

Human assessments are carried out to evaluate the quality of the generated thought processes. We randomly select 100 samples from each dataset and ask three medical students who have undergone clinical internships to assess them. The evaluation employs three metrics: (1) **Knowledge**: whether the knowledge used in the thought process is accurate; (2) **Consistency**: whether the thought process is consistent with the dialogue history; (3) **Rationality**: whether the thought process starts with premises and uses logical progression to derive responses. Table 3 presents the evaluation results, demonstrating that the thought processes in MEDATOR adequately fulfill the above three criteria.

5 Experiments

5.1 Datasets

Our experiments utilize two medical dialogue datasets, MedDG (Liu et al., 2022b) and KaMed

(Li et al., 2021). Dialogues in these datasets exhibit a clear multi-turn context, with each dialogue averaging about 10 turns. The MedDG dataset comprises 17,860 dialogues, focusing on 12 gastroenterology-related diseases. The dataset is divided into 14,862/1,999/999 for training, validation, and testing. The KaMed dataset includes more than 63,000 dialogues, spanning a wide range of diseases across approximately 100 hospital departments. We clean privacy-sensitive content following DFMed (Xu et al., 2023) and divide the dataset into 29,159/1,532/1,539 for training, validation, and testing.

5.2 Baseline methods

We compare our method with two categories of baselines: LLMs equipped with Chinese medical conversation abilities and language models that are fine-tuned on target datasets.

Medical LLMs. (1) **HuatuoGPT-II** (Chen et al., 2023) is a medical LLM performing state-of-the-art in several Chinese medical tasks. (2) **DISC-MedLLM** (Bao et al., 2023) is a medical LLM with strong multi-turn consultation capabilities. (3) **GPT-4** (OpenAI, 2023) is one of the most advanced pre-trained LLMs designed by OpenAI.

Fine-tuned Models. (1) **VRBot** (Li et al., 2021) is a medical dialogue generation (MDG) model with entity tracking and predicting. (2) **GPT-2** (Radford et al., 2019) is a transformer decoder-based language model. (3) **BART** (Lewis et al., 2020) is a transformer-based encoder-decoder model. (4) **Qwen-7B** (Bai et al., 2023) is a strong base language model focusing on Chinese and English. (5) **DFMed** (Xu et al., 2023) enhances MDG with entity and dialogue act flow learning.

5.3 Implementation Details

All the prompt-based operations in EMULATION are implemented with gpt-3.5-turbo, i.e., clinical findings extractions, abductive diagnosis refinement, and deductive diagnosis analysis. We apply the MedBERT² pre-trained in the medical domain as the backbone of the disease retriever and the disease alignment model. The corpus of disease documents, utilized for disease retrieval and to enhance both abductive and deductive reasoning, is derived from a specialist-certified online medical knowledge base *xiaohe*³. We retrieve the top 50 ($K=50$)

²<https://github.com/trueto/medbert>

³<https://www.xiaohe.cn/medical>

Methods		MedDG					KaMed				
		B-1	B-4	R-1	R-2	E-F	B-1	B-4	R-1	R-2	E-F
Zero-shot LLMs	HuatuoGPT-II	42.45	24.78	15.85	4.24	9.45	40.89	22.9	18.06	4.65	11.46
	DISC-MedLLM	40.72	22.6	20.13	6.6	10.15	38.05	20.26	20.48	5.93	13.54
	GPT-4	42.19	23.32	13.99	3.47	13.15	41.88	23.34	13.94	3.1	13.86
Fine-tuned Models	VRBot	29.69	16.34	24.69	11.23	12.78	30.04	16.36	18.71	7.28	12.08
	GPT-2	35.27	19.16	28.74	13.61	16.14	33.76	17.82	26.80	10.56	17.26
	BART	34.94	19.06	29.03	14.40	16.66	33.62	17.64	27.91	11.43	19.20
	Qwen-7B	35.11	19.03	30.19	15.01	18.05	34.00	17.66	28.34	12.18	19.88
	DFMed	42.83	22.90	29.72	14.31	22.92	40.50	20.92	28.33	11.73	22.31
	EMULATION	42.35	22.76	30.91 [†]	15.17 [†]	24.03 [†]	39.87	19.79	28.54 [†]	12.33 [†]	24.27 [†]

Table 4: Automatic evaluation results on two datasets. † denotes statistically significant differences ($p < 0.05$).

Methods	Proactivity	Accuracy	Helpfulness	LQ	Average
DISC-Med	4.52	4.41	4.71	4.98	4.66
DFMed	4.73	4.48	4.79	4.86	4.72
EMULATION	4.69	4.65	4.82	4.99	4.79

Table 5: Automatic evaluation results based on GPT-4.

relevant diseases for further refinement. The refined disease list is sent to the disease alignment model to obtain the top 5 ($K''=5$) diseases that may be discussed in subsequent responses. Then, we employ the pre-trained language model Qwen-7B-Chat⁴ to train the thought process alignment model, which has seven billion parameters and proficient Chinese understanding and generation ability. We train the model using the LoRA approach with $r=64$ and $\alpha=16$. All experiments are carried out on a system equipped with four RTX 3090 GPUs. Other details are presented in the Appendix A.3.

5.4 Automatic Evaluation

We assess the generated responses using three automated metrics: BLEU-1/2/4 (**B-1/4**) (Papineni et al., 2002), evaluating n -gram precision; ROUGE-1/2 (**R-1/2**) (Lin, 2004), assessing n -gram recall; and Entity-F1⁵ (**E-F**) (Liu et al., 2022b), which gauges the accuracy of medical entities, such as diseases and medications.

Table 4 displays the response generation results of all baseline methods. Our framework outperforms the baselines in most metrics, especially the medical entity accuracy. It demonstrates the effectiveness of our framework in diagnostic reasoning and response generation. Specifically, EMULATION performs better than available medical LLMs trained on large-scale medical dialogues and knowledge bases in R-1/2 and E-F. It demonstrates that modeling the diagnostic reasoning process can help generate more accurate and targeted responses. Be-

⁴<https://github.com/QwenLM/Qwen>

⁵<https://github.com/lwgkzl/MedDG>

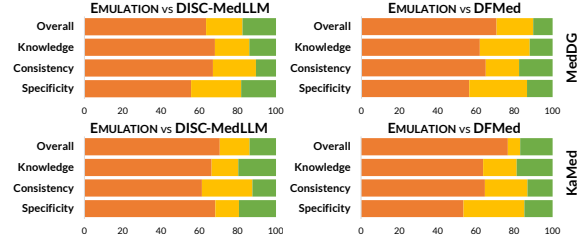


Figure 3: Human evaluation results of baseline methods. Red is for **Win**, yellow for **Tie**, and green for **Lose**.

sides, our framework achieves better n -gram recall and entity accuracy than the state-of-the-art model DFMed. The B-1/4 scores are a bit lower, which may be because the fine-tuned parameters are too limited to approach the linguistic pattern of the gold utterances. However, benefiting from the construction of diagnostic reasoning processes, EMULATION can generate responses with more consistent content and higher entity accuracy.

We also adopt another multi-dimensional automatic evaluation using GPT-4, which follows the method in DISC-Med (Bao et al., 2023). This evaluation focuses on **Proactivity**, **Accuracy**, **Helpfulness**, and **Linguistic Quality (LQ)**. Our EMULATION performs better than other baselines in the above aspects as shown in Table 5.

5.5 Human Evaluation

We conduct human evaluations on comparison with two baselines: DISC-MedLLM and DFMed. We randomly select 100 samples and ask three medical students to evaluate them according to **Knowledge**, **Consistency**, and **Specificity**. Besides, the overall quality is compared. EMULATION outperforms two baselines in all aspects, as shown in Figure 3. Details and case studies are in the Appendix A.5 and A.6. Cross-inference cases are displayed.

Methods	MedDG			KaMed		
	B-4	R-2	E-F	B-4	R-2	E-F
EMULATION	22.76	15.17	24.03	19.79	12.33	24.27
w/o Abd. Reasoning	20.30	13.77	17.47	18.02	11.31	19.75
w/o Ded. Reasoning	22.34	15.07	23.77	19.63	12.05	24.11
w/o Dis. Alignment	21.93	14.69	21.24	18.84	11.87	22.43
w/o Thot. Alignment	22.31	14.95	23.82	19.72	12.12	24.08

Table 6: Ablation study on two datasets

5.6 Analysis of Diagnostic Reasoning Process

To further evaluate the effectiveness of our framework, we analyze several variations of our EMULATION as detailed below: (1) **w/o Abd. Reasoning**, which omits abductive reasoning along with any related deductive reasoning, relying solely on dialogue history to form thought processes. (2) **w/o Ded. Reasoning**, which eliminates the process of deductively analyzing clinical findings and potential diseases. (3) **w/o Dis. Alignment**, which forgoes learning the disease priority, opting to select the most relevant diseases identified through abductive reasoning. (4) **w/o Thot. Alignment**, which directly generates responses based on the results of abductive and deductive reasoning.

Table 6 presents the comprehensive results of our ablation study. There is a noticeable decline in effectiveness across various metrics for the ablation models, underscoring the indispensable contribution of each module within our framework. Particularly, the variant *w/o Abd. Reasoning* experiences a significant drop in response quality due to the lack of initial diagnosis and subsequent analysis. This decline is attributed to the typical practice of clinicians to center medical conversations around certain diseases; thus, omitting precise abductive and deductive reasoning processes leads to a lack of focus in the dialogue. Besides, the *w/o Dis. Alignment* variant shows a marked decrease in performance, reinforcing the importance of disease priority alignment with clinician practices for achieving consistent responses. The outcomes from *w/o Ded. Reasoning* and *w/o Thot. Alignment* further affirm that detailed analyses and thought generation enhance the quality of response generation.

5.7 Analysis of Disease Alignment

We evaluate the diagnosis accuracy to determine the impact of aligning disease priorities. The metric Intersection over Union (IoU) between predicted diseases and plausible diseases E_t^{post} is employed. As shown in Figure 4, aligned diagnoses substan-

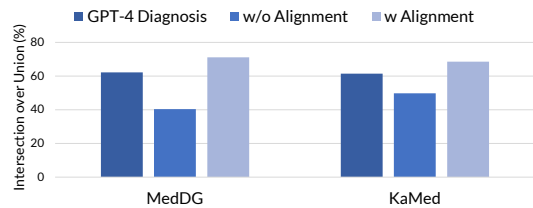


Figure 4: The top-5 diagnosis results with or without disease priority alignment.

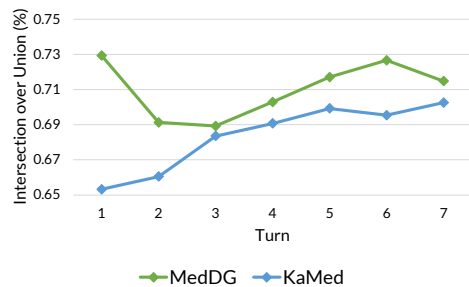


Figure 5: The top-5 diagnosis results of EMULATION at different conversation turns.

tially outperform unaligned diagnoses that are derived directly from abductive reasoning. Remarkably, aligned diagnoses are marginally more accurate than the diseases diagnosed by GPT-4, i.e., E_t^{pri} . We also display the diagnosis results in different conversation turns as shown in Figure 5. Overall, the diagnosis accuracy continuously improves as the number of conversation turns increases. The diagnostic accuracy on the MedDG dataset was initially high but decreased in the second turn. It is because the patients in this dataset provide relatively sufficient information in the first turn, allowing EMULATION to make a good rough diagnosis. However, with the introduction of additional information, the difficulty of diagnosis increases, and the difference between EMULATION’s diagnosis and that of real doctors becomes greater.

6 Related Work

Medical Dialogue Systems Medical Dialogue Systems (MDS) aim to provide healthcare services to patients. A significant and early area of research focuses on automated diagnosis in the form of task-oriented dialogue systems, which prioritize the quick identification of underlying symptoms and the provision of a final diagnosis and do not offer further consultations (Liao et al., 2020; Lin et al., 2019; Chen et al., 2022; Liu et al., 2022a). The study by Wei et al. (2018) presented a dataset with symptom annotations and established an MDS

using reinforcement learning. Xu et al. (2019) incorporated a knowledge graph into MDS to control the sequence of symptom inquiries. Tchang et al. (2022) enhanced system dependability by applying an exploration-confirmation approach and giving precedence to severe diseases.

The emergence of large-scale medical dialogue datasets such as MedDialog (Zeng et al., 2020), MedDG (Liu et al., 2022b), and KaMed (Li et al., 2021), along with pre-trained language models (Lewis et al., 2020; Radford et al., 2019), has sparked increased interest in medical dialogue generation (Liu et al., 2022b, 2021; Lin et al., 2021; Zhao et al., 2022). The research by Liu et al. (2022b) tackled medical dialogue generation by emphasizing entity prediction and entity-centric response creation. Li et al. (2021) presented a semi-supervised variation reasoning system supplemented by a patient state tracker and a physician action network. Xu et al. (2023) introduced a dual flow (i.e., dialogue act and entity flows) modeling approach to enhance dialogue understanding and guide response generation using acts and entities. Dou et al. (2023) applies LLMs to medical dialogue generation in a plug-and-play way.

Medical Large Language Models Given the astonishing performance of GPT-4 in several medical examinations, an increasing number of researchers are directing their attention toward developing medical LLMs. ChatDoctor (Li et al., 2023) is equipped with an external Wikipedia knowledge base and trained on real medical conversations. DoctorGLM (Xiong et al., 2023) is developed using a medical dialogue and question answering dataset supplemented by ChatGPT-translated documents. HuatuoGPT-2 (Chen et al., 2023) and DISC-MedLLM (Bao et al., 2023) try to construct a unified domain adaption framework that uses ChatGPT to convert available documents into pre-training and fine-tuning instructions.

Available medical dialogue systems and medical LLMs try to learn from the outcome of the diagnostic reasoning (i.e., high-quality medical dialogue datasets) but ignore the internal thought process of real clinicians and alignment with clinician preferences. Our work seeks to construct a medical dialogue system that aligns with the internal diagnostic reasoning process of real clinicians.

7 Conclusion

This paper proposes a novel medical dialogue system framework, EMULATION, that emulates clinicians' diagnostic reasoning processes to generate appropriate responses grounded in abductive and deductive diagnostic reasoning analysis and alignment with clinician preferences. Besides, a new diagnostic thought process corpus is presented and utilized to model the clinician preference. Experimental results demonstrate the effectiveness of EMULATION on two datasets. One promising area for future work is applying the EMULATION framework in telemedicine consultations. With the increasing demand for telemedicine, it is essential to explore how this framework can enhance virtual patient-doctor interactions. Future studies could investigate the framework's effectiveness in improving patient satisfaction and overall consultation quality in remote settings.

Limitations

While our framework outperforms various baseline approaches in medical dialogue generation, there is still room for progress. The corpus for the diagnostic thought process is constructed by inferring from doctors' responses. Although it has been assessed that the thought processes in the corpus demonstrate a logical sequence, individual steps within these processes might not mirror those of a real clinician precisely. To enhance consistency in generated thought processes, additional human-annotated thought processes should be collected to conduct further alignment.

Ethics Statement

Our designed system is intended to improve medical consultations for patient care. All datasets were anonymized upon their publication in dataset papers. Nonetheless, due to the training of our model on a limited number of samples for certain conditions, there's a possibility that the responses might contain inaccurate information regarding diagnosis, treatment, and safety measures. We advise treating our system as a supplementary resource and seeking professional medical advice when necessary. Additionally, user interactions with the system could potentially expose sensitive data (e.g., user-reported gender), and the online LLM API services should be substituted by a local open-source model if our system is deployed. Therefore, we urge users to meticulously assess the ethical

considerations of the generated responses. Moreover, the scientific tools utilized in our research, such as NLTK, ROUGE, Transformers, and various GitHub repositories, are openly accessible for academic purposes. The application of these tools in this study adheres to their designated purposes.

Acknowledgment

This work was supported by the Research Grants Council of Hong Kong (15207920, 15213323) and the National Natural Science Foundation of China (62076212).

References

- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. [Qwen technical report](#). *CoRR*, abs/2309.16609.
- 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](#). *CoRR*, abs/2308.14346.
- Susan Cameron and Imani Turtle-Song. 2002. Learning to write case notes using the soap format. *Journal of Counseling & Development*, 80(3):286–292.
- Hyungjoo Chae, Yongho Song, Kai Tzu-iunn Ong, Taeyoon Kwon, Minjin Kim, Youngjae Yu, Dongha Lee, Dongyeop Kang, and Jinyoung Yeo. 2023. [Dialogue chain-of-thought distillation for commonsense-aware conversational agents](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 5606–5632. Association for Computational Linguistics.
- Junying Chen, Dongfang Li, Qingcai Chen, Wenxiu Zhou, and Xin Liu. 2022. [Diaformer: Automatic diagnosis via symptoms sequence generation](#). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 4432–4440. AAAI Press.
- Junying Chen, Xidong Wang, Anningzhe Gao, Feng Jiang, Shunian Chen, Hongbo Zhang, Dingjie Song, Wenya Xie, Chuyi Kong, Jianquan Li, Xiang Wan, Haizhou Li, and Benyou Wang. 2023. [Huatuogpt-ii, one-stage training for medical adaption of llms](#). *CoRR*, abs/2311.09774.
- Chengfeng Dou, Zhi Jin, Wenpin Jiao, Haiyan Zhao, Yongqiang Zhao, and Zhengwei Tao. 2023. [PlugMed: Improving specificity in patient-centered medical dialogue generation using in-context learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5050–5066, Singapore. Association for Computational Linguistics.
- Arthur S Elstein, Lee S Shulman, and Sarah A Sprafka. 1978. *Medical problem solving: An analysis of clinical reasoning*. Harvard University Press.
- Yao Fu, Hao Peng, Litu Ou, Ashish Sabharwal, and Tushar Khot. 2023. [Specializing smaller language models towards multi-step reasoning](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 10421–10430. PMLR.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Keith J Holyoak and Robert G Morrison. 2005. *The Cambridge handbook of thinking and reasoning*. Cambridge University Press.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*, pages 7871–7880. Association for Computational Linguistics.
- Dongdong Li, Zhaochun Ren, Pengjie Ren, Zhumin Chen, Miao Fan, Jun Ma, and Maarten de Rijke. 2021. [Semi-supervised variational reasoning for medical dialogue generation](#). In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 544–554. ACM.
- Yunxiang Li, Zihan Li, Kai Zhang, Ruilong Dan, Steve Jiang, and You Zhang. 2023. Chatdoctor: A medical chat model fine-tuned on a large language model meta-ai (llama) using medical domain knowledge. *Cureus*, 15(6).
- Kangenbei Liao, Qianlong Liu, Zhongyu Wei, Baolin Peng, Qin Chen, Weijian Sun, and Xuanjing Huang.

2020. [Task-oriented dialogue system for automatic disease diagnosis via hierarchical reinforcement learning](#). *ArXiv*, abs/2004.14254.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Shuai Lin, Pan Zhou, Xiaodan Liang, Jianheng Tang, Ruihui Zhao, Ziliang Chen, and Liang Lin. 2021. [Graph-evolving meta-learning for low-resource medical dialogue generation](#). In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 13362–13370. AAAI Press.
- Xinzhu Lin, Xiahui He, Qin Chen, Huaixiao Tou, Zhongyu Wei, and Ting Chen. 2019. [Enhancing dialogue symptom diagnosis with global attention and symptom graph](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5033–5042, Hong Kong, China. Association for Computational Linguistics.
- Wenge Liu, Yi Cheng, Hao Wang, Jianheng Tang, Yafei Liu, Ruihui Zhao, Wenjie Li, Yefeng Zheng, and Xiaodan Liang. 2022a. ["my nose is running." "are you also coughing?": Building A medical diagnosis agent with interpretable inquiry logics](#). In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 4266–4272. ijcai.org.
- Wenge Liu, Jianheng Tang, Yi Cheng, Wenjie Li, Yefeng Zheng, and Xiaodan Liang. 2022b. [Meddg: An entity-centric medical consultation dataset for entity-aware medical dialogue generation](#). In *Natural Language Processing and Chinese Computing - 11th CCF International Conference, NLPCC 2022, Guilin, China, September 24-25, 2022, Proceedings, Part I*, volume 13551 of *Lecture Notes in Computer Science*, pages 447–459. Springer.
- Wenge Liu, Jianheng Tang, Xiaodan Liang, and Qingling Cai. 2021. [Heterogeneous graph reasoning for knowledge-grounded medical dialogue system](#). *Neurocomputing*, 442:260–268.
- Elizabeth Murray, Lance Pollack, Martha White, and Bernard Lo. 2007. Clinical decision-making: physicians’ preferences and experiences. *BMC Family Practice*, 8(1):1–10.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Vimla L Patel and Marco F Ramoni. 1997. Cognitive models of directional inference in expert medical reasoning.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Jonathan Silverman, Suzanne Kurtz, and Juliet Draper. 2016. *Skills for communicating with patients*. crc press.
- Arsène Fansi Tchango, Rishab Goel, Julien Martel, Zhi Wen, Gaétan Marceau-Caron, and Joumana Ghosn. 2022. [Towards trustworthy automatic diagnosis systems by emulating doctors’ reasoning with deep reinforcement learning](#). In *NeurIPS*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Zhongyu Wei, Qianlong Liu, Baolin Peng, Huaixiao Tou, Ting Chen, Xuanjing Huang, Kam-fai Wong, and Xiangying Dai. 2018. [Task-oriented dialogue system for automatic diagnosis](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 201–207, Melbourne, Australia. Association for Computational Linguistics.
- Honglin Xiong, Sheng Wang, Yitao Zhu, Zihao Zhao, Yuxiao Liu, Linlin Huang, Qian Wang, and Ding-gang Shen. 2023. [Doctorglm: Fine-tuning your chinese doctor is not a herculean task](#). *CoRR*, abs/2304.01097.
- Kaishuai Xu, Wenjun Hou, Yi Cheng, Jian Wang, and Wenjie Li. 2023. [Medical dialogue generation via dual flow modeling](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6771–6784, Toronto, Canada. Association for Computational Linguistics.
- Lin Xu, Qixian Zhou, Ke Gong, Xiaodan Liang, Jianheng Tang, and Liang Lin. 2019. [End-to-end knowledge-routed relational dialogue system for automatic diagnosis](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 7346–7353. AAAI Press.

Kun-Hsing Yu, Elizabeth Healey, Tze-Yun Leong, Isaac S Kohane, and Arjun K Manrai. 2024. Medical artificial intelligence and human values. *New England Journal of Medicine*, 390(20):1895–1904.

Guangtao Zeng, Wenmian Yang, Zeqian Ju, Yue Yang, Sicheng Wang, Ruisi Zhang, Meng Zhou, Jiaqi Zeng, Xiangyu Dong, Ruoyu Zhang, Hongchao Fang, Penghui Zhu, Shu Chen, and Pengtao Xie. 2020. *MedDialog: Large-scale medical dialogue datasets*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9241–9250, Online. Association for Computational Linguistics.

Yu Zhao, Yunxin Li, Yuxiang Wu, Baotian Hu, Qingcai Chen, Xiaolong Wang, Yuxin Ding, and Min Zhang. 2022. *Medical dialogue response generation with pivotal information recalling*. In *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022*, pages 4763–4771. ACM.

A Appendix

A.1 Automatic Evaluation Details

We adopt the calculation approach used in the original dataset paper MedDG (Liu et al., 2022b) and in the most recent paper DFMed (Xu et al., 2023). The “nltk” package with version 4.5.1 is used to calculate BLEU scores. The “rouge” version is 1.0.1. There is a score gap between our study and some baseline studies since the metric package used for calculating BLEU is different. For example, the officially released code for VRBot doesn’t specify the metric package version such as “nlgeval” and “rouge”, while the results will be affected by the package version.

A.2 Disease Annotation

We annotate the potential diseases for each dialogue turn with the help of GPT-4 using two different prompts. One is “Generate which disease the patient may suffer from based on the medical conversation and explain why. The diseases should be ranked by their possibility according to the conversation.” This prompt aims to leverage the diagnostic capability of GPT-4 to infer potential diseases. The other is “Generate which disease the doctor is considering or intending to rule out based on the doctor’s response and explain why. The diseases should be mainly ranked by their relationship to the response.” This prompt aims to deduce diseases from the doctor’s response. The difference between these two prompts is whether the current doctor’s response is given. Examples

of two disease annotation prompts are shown in Figure 6 and Figure 7.

After inferring diseases using two prompts respectively, we need to link each inferred disease to the disease in an external knowledge base. We divide this linking into two steps: (1) Coarse matching and (2) GPT-4 assisted matching. For the coarse matching, we first build a dense retriever to calculate the relevance score between disease name and disease documents. The disease documents are from the knowledge base. This retriever can bridge the gap between different aliases of one disease. Then, we retrieve the top 10 relevant diseases as a preliminary list. For the GPT-4 assisted matching, we organize a prompt that requires GPT-4 to select from the preliminary disease list. The prompt is “Select diseases from the candidate disease list that describe the same one as our target.” An example of a disease match prompt is shown in Figure 8.

A.3 Implementation Details

Baseline Methods. For the medical LLMs, we use the HuatuoGPT2-13B⁶ and DISC-MedLLM⁷ to generate responses in zero-shot way as they have already been fine-tuned on medical dialogue datasets and other large-scale medical knowledge datasets. These two models are all based on language models with 13B parameters. We utilize the default decoding parameters to generate responses.

For fine-tuned baseline methods, we employ generation results from the original papers.

Fine-tuned modules in our framework. For training disease retriever in §3.1, we use the MedBERT as the encoder and employ contrastive learning. The batch size is 8 with 6 gradient accumulation steps. The learning rate is set at $3e-5$. We train 6 epochs for the MedDG dataset and 10 epochs for the KaMed dataset. We select the checkpoint with the highest disease recall rate in validation datasets.

For training the disease priority alignment model, we also use the MedBERT as the encoder and employ contrastive learning. The batch size is 2 with 8 gradient accumulation steps. The learning rate is set at $1e-5$. We train 3 epochs for two datasets and select the checkpoint with the highest IoU in validation datasets.

For training the thought process generation model, we use the Qwen-7B-Chat model and adopt LoRa fine-tuning. The LoRa settings are: r equals

⁶<https://github.com/FreedomIntelligence/HuatuoGPT-II>

⁷<https://github.com/FudanDISC/DISC-MedLLM>

Prompt for disease annotation without response

Given the medical conversation history as follows:

Patient: Is there any medicine that can relieve severe stomach movement and noise? (Female, 15 years old).

Doctor: Hello, how long has this been happening?

Patient: About two or three weeks.

Doctor: This symptom is medically known as borborygmi, generally caused by the sounds of gastrointestinal motility, which is normal and not too worrisome. However, stomach rumbling can also occur when one is hungry. If you frequently experience this, it might also be due to indigestion, for which you can take digestive aids like domperidone, digestive tablets, or engage in some aerobic exercises that help digestion. Additionally, this phenomenon could be related to eating too quickly and swallowing air or high stress levels. Pay attention to your diet and try to relax.

Patient: And I've also had prolonged diarrhea, with mucus in my stool. It feels too uncomfortable with the intestinal movement.

Generate which disease the patient may suffer from based on the medical conversation and explain why. The generated diseases are enclosed in angle brackets, i.e., <.>, <.>, etc., and each bracket pair contains one disease. The diseases should be ranked by their possibility according to the conversation.

Figure 6: Prompt for disease annotation without the doctor's response.

Prompt for disease annotation with response

Given the medical conversation history as follows:

Patient: Is there any medicine that can relieve severe stomach movement and noise? (Female, 15 years old).

Doctor: Hello, how long has this been happening?

Patient: About two or three weeks.

Doctor: This symptom is medically known as borborygmi, generally caused by the sounds of gastrointestinal motility, which is normal and not too worrisome. However, stomach rumbling can also occur when one is hungry. If you frequently experience this, it might also be due to indigestion, for which you can take digestive aids like domperidone, digestive tablets, or engage in some aerobic exercises that help digestion. Additionally, this phenomenon could be related to eating too quickly and swallowing air or high stress levels. Pay attention to your diet and try to relax.

Patient: And I've also had prolonged diarrhea, with mucus in my stool. It feels too uncomfortable with the intestinal movement.

The doctor's ground truth response:

This condition might be related to chronic enteritis. You could take some berberine tablets and Intestine Inflammation Relief.

Generate which disease the doctor is considering or intending to rule out based on the doctor's response and explain why. The generated diseases are enclosed in angle brackets, i.e., <.>, <.>, etc., and each bracket pair contains one disease. The diseases should be mainly ranked by their relationship to the response.

Figure 7: Prompt for disease annotation with the doctor's response.

Prompt for disease match

Select the same diseases as the raw disease text from the candidate diseases.

Examples:

Given the raw disease text:

infectious disease

Candidate diseases:

<Infectious diseases> <Hereditary diseases> <Infections> <Sexually transmitted diseases> <AIDS with pulmonary tuberculosis> <Hereditary immune nephritis> <Mycobacterium kansasii infection> <Echo and Coxsackie virus in children Infection> <Pre-stage of HIV infection> <Pregnancy complicated by gonorrhoea>

The selected diseases:

<Infectious Disease>

Given the raw disease text:

Hepatitis B (chronic or past infection)

Candidate diseases:

<Acute hepatitis B> <Children's hepatitis B> <Hepatitis B antigenemia> <Hepatitis B> <Neonatal hepatitis B> <Hepatitis B viral arthritis> <Chronic hepatitis B virus> <Hepatitis B carrier> <Hepatitis B virus related nephritis> <Chronic hepatitis C>

The selected diseases:

<Hepatitis B> <Chronic Hepatitis B>

...

Then, for the raw disease text:

Infectious Viral Hepatitis

Candidate diseases:

<Viral hepatitis> <Viral hepatitis in the elderly> <Viral hepatitis B> <Hepatitis virus-related rheumatic diseases> <Autoimmune hepatitis> <Viral hepatitis B in children> <Infectious diseases> <Hepatitis B viral joints inflammation> <Hepatitis B carrier> <Hepatitis>

Please think about the following steps. The output format is:

Step (1) If the raw disease text is describing a disease? If not, please enter 'Not a disease phrase'.

Step (2) If the above is "Yes", then select the same disease as the raw disease text from the candidate diseases and explain why. The selected diseases are enclosed in angle brackets, i.e., <.>, <.>, etc., and each bracket pair contains one disease. If no disease can be selected from the candidate diseases, please output 'Not in the candidate diseases'.

Figure 8: Prompt for disease match.

Datasets	Top-10	Top-25	Top-50	Top-100
MedDG	60.90%	83.99%	92.88%	96.85%
KaMed	64.91%	83.10%	90.64%	95.04%

Table 7: Retrieval results.

64, α equals 16, dropout rate equals 0.05. The learning rate is set at $3e-4$, the batch size for each GPU is 2 with 8 gradient accumulation steps. We train the LoRa parameters for 3 epochs and select the last checkpoint to generate responses.

Prompt modules in our framework. For the abductive reasoner, the prompt is shown in Figure 9. For the deductive reasoner, the prompt is shown in Figure 10.

A.4 Disease Retrieval Results

Table 7 displays the recall rate of different top- K diseases (i.e., $K=10, 25, 50, 100$).

A.5 Human Evaluations on Comparison with Baseline Methods

We compare our framework EMULATION with baseline methods from three aspects. (1) **Knowledge Accuracy (Knowledge)**, which assesses whether the response correctly applies the disease knowledge. (2) **Consistency**, which assesses whether the response is consistent with the dialogue context and doctor’s ground truth response. (3) **Specificity**, which assesses whether the response provides specific diagnoses, prescriptions, treatment plans, or examination requirements rather than general suggestions. We ask our annotators to compare the responses generated by our method with the responses generated by baseline methods.

A.6 Case Studies

Figure 11 and Figure 12 display cases on the MedDG and KaMed datasets, respectively. Besides, we also infer samples in MedDG with the model trained on KaMed and infer samples in KaMed with the model trained on MedDG in Figure 13 and Figure 14.

Prompt for abductive reasoning

Given the medical conversation history as follows:
 Patient: Is there any medicine that can relieve severe stomach movement and noise? (Female, 15 years old).
 Doctor: Hello, how long has this been happening?
 Patient: About two or three weeks.
 Doctor: This symptom is medically known as borborygmi, generally caused by the sounds of gastrointestinal motility, which is normal and not too worrisome. However, stomach rumbling can also occur when one is hungry. If you frequently experience this, it might also be due to indigestion, for which you can take digestive aids like domperidone, digestive tablets, or engage in some aerobic exercises that help digestion. Additionally, this phenomenon could be related to eating too quickly and swallowing air or high stress levels. Pay attention to your diet and try to relax.
 Patient: And I've also had prolonged diarrhea, with mucus in my stool. It feels too uncomfortable with the intestinal movement.

New clinical findings are:
 prolonged diarrhea, mucus in my stool, uncomfortable with the intestinal movement

Candidate disease list:
 Chronic Enteritis, Gastroenteritis, Colitis, Indigestion, Functional Gastrointestinal Disease, ...

Knowledge:
 ...

Select all possible diseases that align with new clinical findings and the medical conversation.

Instructions:
 1. You should carefully consider all patient information (e.g., symptoms, medical history, medicine, age, gender, etc.) and explain why.
 2. You should select from the candidate disease list.
 3. You can select multiple diseases if necessary.
 4. If no disease can align with new clinical findings and the medical conversation, please input "Null".
 5. Your output format is: "... therefore, possible diseases are: ...". Possible diseases are enclosed in angle brackets and in Chinese, i.e., <.>, <.>, etc., and each bracket pair contains one disease.

Figure 9: Prompt for abductive reasoning.

Prompt for deductive reasoning

Given the medical conversation history as follows:
 Patient: Is there any medicine that can relieve severe stomach movement and noise? (Female, 15 years old).
 Doctor: Hello, how long has this been happening?
 Patient: About two or three weeks.
 Doctor: This symptom is medically known as borborygmi, generally caused by the sounds of gastrointestinal motility, which is normal and not too worrisome. However, stomach rumbling can also occur when one is hungry. If you frequently experience this, it might also be due to indigestion, for which you can take digestive aids like domperidone, digestive tablets, or engage in some aerobic exercises that help digestion. Additionally, this phenomenon could be related to eating too quickly and swallowing air or high stress levels. Pay attention to your diet and try to relax.
 Patient: And I've also had prolonged diarrhea, with mucus in my stool. It feels too uncomfortable with the intestinal movement.

New clinical findings are:
 prolonged diarrhea, mucus in my stool, uncomfortable with the intestinal movement

Possible diseases based on the conversation:
 Chronic Enteritis, Gastroenteritis, Colitis

Knowledge:
 ...

Analyze if the new clinical findings support, oppose, or are irrelevant to the possible diseases. You can analyze findings individually or in groups. If no clinical findings are included, the output should be "Null". The 'Status: Support' means supporting the disease, and 'Status: Oppose' means opposing the disease. If the clinical findings are not relevant to all possible diseases, the output should be 'Status: Null'. Your output should be in the following format:

Analysis:
 1. Findings:
 Disease:
 Reason:
 Status: Support

2. Findings:
 Disease:
 Reason:
 Status: Oppose

3. Findings:
 Reason:
 Status: Null

Figure 10: Prompt for deductive reasoning.

Given medical dialogue as follows:

Patient: A 28-year-old male previously had bowel movements roughly every 3-4 days with thick stools. However, he had a good appetite before, but in the recent half year, he has been experiencing discomfort in his stomach and intestines, with a decreased appetite, and now has bowel movements every two days, which are thin and soft. A year ago, a gastroscopy revealed gastritis erosion, and tests for Helicobacter pylori were negative. (Male, 28 years old).

Doctor: Hello! How long have you been experiencing this situation? Is your diet regular? What kind of food do you like? Do you eat a lot of spicy or cold food?

Patient: This situation has lasted for half a year. I'm busy with work, so I only have an irregular meal at noon every day. Not much spicy or cold food.

Doctor: What other symptoms do you have?

Patient: I often feel discomfort in my stomach and intestines, and the regularity and shape of my bowel movements are completely different, with no other noticeable symptoms. I have previously taken stomach comfort granules and Stomach Ease. Do I need to undergo a colonoscopy or any other examination for my condition? Or what should I pay attention to and treat next?

Clinical findings in the dialogue are:

28 years old, male, thick stools, good appetite, discomfort in stomach and intestines, decreased appetite, thin and soft stools, gastritis erosion, half a year, irregular diet, not much spicy or cold food, different bowel movement regularity and shape

Clinical analysis:

- "Stomach and intestinal discomfort, decreased appetite, change in bowel habits, thin and soft stools." supports "Gastritis, Irritable Bowel Syndrome, Functional Gastrointestinal Disease"
- "Diagnosed with gastritis erosion a year ago." supports "Gastroesophageal Reflux Disease"
- "Helicobacter pylori test results were negative." opposes "Gastritis"
- "Patient has had symptoms for half a year." supports "Gastritis"
- "Patient's diet is irregular, usually only eating lunch due to being busy at work." supports "Irritable Bowel Syndrome"
- "Patient rarely eats spicy or cold food." opposes "Gastroesophageal Reflux Disease"
- "Frequent stomach and intestinal discomfort" supports "Gastritis"
- "Change in bowel movement regularity and shape" supports "Irritable Bowel Syndrome"
- "No other significant symptoms" opposes "Peptic Ulcer"

Possible diseases that might be discussed:

Gastritis, Irritable Bowel Syndrome, Peptic Ulcer, Gastroesophageal Reflux Disease, Functional Gastrointestinal Disease

Thought process:

1. The patient reports frequent stomach and intestinal discomfort and significant changes in bowel movement regularity and shape. These symptoms, along with the patient's history of gastritis erosion, suggest the possibility of chronic gastrointestinal issues.
2. The patient mentioned that they have tried "stomach comfort granules" and "Stomach Ease," which are traditional Chinese medicines commonly used for treating gastrointestinal discomfort. The effectiveness of these treatments could provide clues about the nature of the patient's condition.
3. The patient inquires about the need for a colonoscopy, indicating their concern about the severity of their condition. A colonoscopy could provide valuable information about the patient's gastrointestinal health and help identify any abnormalities. But it is not needed at the moment as there are no symptoms related to colitis.
4. Considering the patient's symptoms, medical history, and response to previous treatments, the doctor contemplates recommending a gastroscopy to further investigate the patient's stomach condition. This could help identify any potential issues causing the patient's symptoms.
5. The doctor may advise the patient to pay attention to their diet, as irregular eating habits could exacerbate gastrointestinal problems. **Therefore, the doctor replies, "You could consider a gastroscopy to check. Pay attention to your diet in the meantime."** This includes suggesting a gastroscopy to further investigate the patient's stomach condition and advising the patient to pay attention to their diet, as irregular eating habits could exacerbate their symptoms.

Ground Truth Response:

There's no need for a colonoscopy for the time being, but it's crucial to pay attention to your diet regularly.

DFMed:

Your situation is likely due to gastrointestinal dysfunction. It's recommended to have a light diet, keep warm, avoid spicy, stimulating, and greasy foods, and eat more vegetables and fruits. You can take Trimebutine and Bifidobacterium triplex live bacteria tablets together to improve symptoms. If that doesn't work, you may need to go to the hospital for a colonoscopy examination.

DISC-MedLLM:

Your symptoms may be caused by gastritis or gastrointestinal dysfunction. However, persistent symptoms may require further medical evaluation. 1. Diet: Try changing your dietary habits, avoiding overly greasy, spicy, or irritating foods, and increasing the intake of dietary fiber such as whole grains, vegetables, and fruits to help improve digestion. 2. Medication: Continue taking stomach medication as directed by your doctor. If symptoms persist or worsen, the medication treatment plan may need to be adjusted...

Figure 11: Generated case on the MedDG dataset.

Given medical dialogue as follows:

Patient: What tests are needed to check for tuberculosis? (Female, 2 years and 9 months old).

Doctor: To check for tuberculosis, you need a chest CT, sputum test, tuberculosis infection T-cell test, and tuberculin skin test, among others.

Patient: How much blood needs to be drawn for the tuberculosis infection T-cell test and tuberculin skin test for children aged 2 to 3? What's the general heart rate range for children of this age?

Doctor: A few milliliters of blood. The likelihood of tuberculosis in toddlers is extremely low.

Patient: Our child was hospitalized for pneumonia, treated for 27 days with no significant improvement, even with meropenem administered.

Clinical findings in the dialogue are:

Female, 2 years and 9 months old, child aged 2 to 3, pneumonia, treated for 27 days with no significant improvement, administered meropenem.

Clinical analysis:

- "The child was hospitalized for pneumonia, treated for 27 days with no significant improvement" supports "Tuberculosis, Pneumonia, Tuberculosis tract infection, Drug-resistant tuberculosis "

- "Meropenem was administered" supports "Tuberculosis, Drug-resistant tuberculosis "

Possible disease might be discussed:

Tuberculosis, Pneumonia, Tuberculosis, Respiratory tract infection, Drug-resistant tuberculosis

Thought process:

1. The patient mentions that their child was hospitalized for pneumonia, treated for 27 days but showed no significant improvement. This raises concerns as it suggests that the current treatment regimen might be ineffective.

2. The patient also mentioned that the child had been treated with meropenem, an antibiotic commonly used for severe bacterial infections. This may indicate that the pneumonia is bacterial, which is common in children.

3. However, despite treatment, the lack of improvement raises concerns about the accuracy of the diagnosis and the effectiveness of the treatment. It's also possible that the bacteria causing the pneumonia have developed resistance to meropenem.

4. Considering these factors, the doctor suggests conducting a tuberculin skin test. This test can help determine whether the child has been infected with the bacteria that cause tuberculosis. While tuberculosis is less common in children, it cannot be ruled out when a child is ill for a long time and unresponsive to treatment.

Therefore, the doctor replies, "A tuberculin skin test can be done." This indicates that, given the lack of improvement with current treatments, a tuberculin skin test is recommended to explore other possible causes for the child's persistent symptoms. This reflects the doctor's thoroughness in considering all possible diagnoses and the importance of re-evaluating the diagnosis when treatment is ineffective.

Ground Truth Response:

I suggest you get tested to rule out tuberculosis.

DFMed:

That means there has been no improvement, further examination is needed.

DISC-MedLLM:

Your child has been hospitalized for pneumonia for 27 days and has been treated with meropenem, but there hasn't been significant improvement. In this situation, we need to reassess your child's condition to determine the next steps in treatment. Meropenem is a potent antibiotic commonly used to treat severe infections, but if there hasn't been noticeable improvement after its use, we need to consider the following. Check for resistance: It's possible that the bacteria causing the infection have developed resistance to meropenem, resulting in poor treatment outcomes. Further bacterial culture and sensitivity testing are needed to identify the most effective antibiotic. Rule out complications: Pneumonia can lead to other complications such as pleural effusion or empyema, which may require specific management.

Figure 12: Generated case on the KaMed dataset.

Given medical dialogue as follows:

Patient: I keep having diarrhea no matter what medicine I take (Female, 20 years old).

Doctor: Hello, how long has this condition been going on? Do you have hyperthyroidism or diabetes?

Patient: About half a month, almost a month now. It gets a little better with medication. I don't have any related diseases. I'm not sure if it's because of my diet, but it's the same as before I got sick.

Doctor: Is the stool formed? Is there any mucus, pus, or jelly-like substance?

Patient: It's basically not formed, just viscous. It gets a bit better when I start taking medicine, but then it worsens after a while.

Doctor: Can a light diet, less greasy and spicy food, improve it a bit?

Patient: No blood. It seems like it, but I haven't paid much attention to this aspect, because I'm studying away from home, and it's hard to control my diet. I've been taking vitamin C for about two to three months before I got sick.

Clinical findings in the dialogue:

Diarrhea, recurrent, female, 20 years old, about half a month to a month, slight improvement with medication, no hyperthyroidism or diabetes, stool not formed, viscous, no blood, studying away from home, poor dietary control, taking vitamin C for two to three months

Clinical analysis:

- "Persistent diarrhea" supports "Hyperthyroidism"
- "Diarrhea for almost a month, slightly improves with medication." supports "Inflammatory bowel disease"
- "No known related diseases." opposes "Colorectal cancer"
- "The patient's stool is basically not formed and is viscous." supports "Gastroenteritis"
- "The patient's stool is basically not formed and is viscous." opposes "Food intolerance"
- "Condition slightly improves with medication but worsens after a while." supports "Irritable Bowel Syndrome"
- "Not paying much attention to diet" supports "Food intolerance"
- "Taking vitamin C before getting sick" supports "Irritable Bowel Syndrome"

Possible disease might be discussed:

Gastroenteritis, Irritable Bowel Syndrome, Food intolerance, GRED

Thought process:

1. The patient reports no blood in the stool, which is a good sign as it rules out some serious conditions such as gastrointestinal bleeding.
2. The patient mentions that symptoms seem to worsen when not paying attention to her diet. This suggests that her diet could be a contributing factor to her symptoms. Diet plays a key role in gastrointestinal health, and certain foods can exacerbate symptoms like diarrhea.
3. The patient also mentions having taken vitamin C for about two to three months. While vitamin C is generally safe, excessive intake can cause side effects such as diarrhea. However, given the duration the patient has been taking vitamin C and the duration of the symptoms, vitamin C is unlikely to be the primary cause of her symptoms.
4. Considering the patient's symptoms, dietary habits, and vitamin C intake, the doctor concludes that the patient's diet could be contributing to her symptoms.

Therefore, the doctor replies, "Not paying attention to diet, improper eating, can aggravate diarrhea." This response emphasizes the importance of diet in managing gastrointestinal symptoms and suggests the patient pay attention to her diet to help alleviate her symptoms. The doctor's response is based on the patient's symptoms, dietary habits, and vitamin C intake. It reflects a comprehensive understanding of the patient's condition and a practical approach to management.

Ground Truth Response:

Control your diet. You can take some berberine and montmorillonite powder. The gastrointestinal function has not yet recovered.

DFMed:

Your condition is considered to be caused by intestinal dysfunction. It is recommended that you follow a light diet, eat less oily and fatty food, and avoid spicy food to see if it improves. You can take some probiotics to regulate the intestinal flora. Have you had a colonoscopy?

DISC-MedLLM:

You can try now to have a light diet, less oily and greasy food, and avoid spicy foods to see if there's any improvement. Additionally, you mentioned taking vitamin C continuously before getting sick, which might have impacted your symptoms to some extent. It's recommended to keep your diet light and try to avoid spicy foods as much as possible. Moreover, drink plenty of water to maintain adequate hydration, which can help improve intestinal function. If symptoms persist or worsen, it's advised to seek medical attention promptly for further examination to determine the cause. Hope you recover soon!

Figure 13: Generated MedDG case with the model trained on KaMed.

Given the medical dialogue:

Patient: My child has an upper respiratory infection and a fever. What medicine should he take? (Male, 3 years old).

Doctor: Hello, what is the baby's temperature?

Patient: 37.5°C.

Doctor: When feverish, make sure the child drinks plenty of water, stays warm, and if the fever doesn't exceed 38.5°C, use physical cooling methods primarily, keep the air in the room circulating, bathe with lukewarm water, drink more water, and if the child's temperature exceeds 38.5°C, appropriate use of antipyretics or fever suppositories is possible.

Patient: After taking oral honeysuckle granules, is it necessary to take anti-inflammatory medication?

Doctor: For a lukewarm bath: the water temperature should be 32-34°C, wipe for more than 10 minutes, especially wipe the skin folds, such as the neck, armpits, elbows, etc., and drink more warm boiled water. Is there a runny nose? Is the throat red and swollen?

Patient: No runny nose.

Doctor: The throat is red, right?

Patient: Went to a clinic and the doctor prescribed honeysuckle after checking the throat.

Doctor: It's okay to use honeysuckle granules, and you can use throat spray.

Patient: It's said to be viral, can we take some azithromycin?

Doctor: Did you get a blood test?

Patient: No.

Doctor: First, observe with honeysuckle granules for two days. Azithromycin targets mycoplasma infections, and without tests, it's not recommended for the baby.

Patient: Can we take some cephalosporins then?

Clinical findings in the dialogue:

Child with upper respiratory infection, fever, male, 3 years old, took honeysuckle granules orally, no runny nose, red throat, viral, no routine blood test, baby's temperature 37.5°C

Clinical analysis:

- "Upper respiratory infection, fever" supports "Upper respiratory infection, influenza"
- "Upper respiratory infection, fever" opposes "Pneumonia"
- "The child has a fever with a temperature of 37.5°C." supports "Upper respiratory infection, common cold, influenza"
- "Patient has taken honeysuckle granules." supports "Upper respiratory infection, common cold"
- "The child was brought to a clinic, and a doctor checked the throat." supports "Upper respiratory infection, tonsillitis, pharyngitis"
- "The doctor prescribed honeysuckle." supports "Upper respiratory infection, tonsillitis, pharyngitis"
- "The patient's condition was described as viral." supports "Viral infectious diseases, upper respiratory infection, infection, influenza, infectious mononucleosis"

Possible disease might be discussed:

Upper respiratory infection, Infection, Viral Infectious Diseases, Acute Upper respiratory infection, Infectious Disease

Thought process:

1. The patient inquires about giving the child cephalosporin, an antibiotic. The doctor considers the child's symptoms and current treatment plan.
2. The doctor notes the child has already been given honeysuckle granules, a traditional Chinese medicine commonly used for treating symptoms of colds and flu. The doctor also suggests using a throat spray to soothe the child's throat.
3. The doctor considers the necessity of using antibiotics. Antibiotics are typically used to treat bacterial infections, not viral infections. Given the child's symptoms and another doctor's assessment, the child may have a viral infection, leading the doctor to believe antibiotics might not be necessary at this time.
4. The doctor also considers the potential side effects and risks of antibiotics, such as antibiotic resistance and adverse reactions. Antibiotics should only be used when absolutely necessary, such as when a bacterial infection is confirmed.

Therefore, the doctor replies, "First, continue observing with honeysuckle granules and use the throat spray. I do not recommend using cephalosporins." This response suggests continuing the current treatment plan without antibiotics, reflecting the doctor's consideration of the child's symptoms, the effectiveness of the current treatment, and the potential risks and benefits of using antibiotics. The doctor advises continuing to monitor the child's condition, and if symptoms persist or worsen, to seek further medical advice.

Ground Truth Response:

Because there has been no blood test, there is insufficient indication for the use of antibiotics. Therefore, it is not recommended at this time. Try using honeysuckle granules and throat spray for two days to see if there is any improvement.

DFMed:

Yes, cephalosporins are antibiotics and can be taken if a routine blood test indicates a bacterial infection.

DISC-MedLLM:

Yes, cephalosporins are antibiotics and are ineffective against viral infections.

Figure 14: Generated KaMed case with the model trained on MedDG.