# FEEDBACK-AWARE MONTE CARLO TREE SEARCH FOR EFFICIENT INFORMATION SEEKING IN GOAL-ORIENTED CONVERSATIONS

Harshita Chopra & Chirag Shah

University of Washington, Seattle hchopra3@cs.washington.edu chiraqs@uw.edu

#### ABSTRACT

The ability to identify and acquire missing information is a critical component of effective decision making and problem solving. With the rise of conversational artificial intelligence (AI) systems, strategically formulating information-seeking questions becomes crucial and demands efficient methods to guide the search process. We introduce a novel approach to adaptive question-asking through a combination of Large Language Models (LLM) for generating questions that maximize information gain, Monte Carlo Tree Search (MCTS) for building and using a decision tree across multiple samples, and a hierarchical feedback mechanism to learn from past interactions. We present two key innovations: (1) an adaptive MCTS algorithm that balances exploration and exploitation for efficient search over potential questions; and (2) a clustering-based feedback algorithm that leverages prior experience to guide future interactions. Each incoming sample is assigned to a cluster based on its semantic similarity with previously observed samples. Our UCT (Upper Confidence bound for Trees) formulation selects optimal questions by combining expected rewards, a function of information gain, with a clusterspecific bonus that decays with depth, to emphasize the importance of early-stage questions that have proven effective for narrowing the solution space in similar samples. Experiments in medical diagnosis and troubleshooting domains demonstrate that our method leads to an average of 12% improvement in success rates and a 10x reduction in the average number of LLM calls made per conversation for the search process, in comparison to the state of the art.<sup>1</sup>

### **1** INTRODUCTION

The field of conversational AI has witnessed significant advancements in recent years, enabling systems to engage in complex, goal-directed dialogues. A fundamental requirement in these systems is the ability to identify and seek out missing information which is required to find a solution. In domains that involve complex reasoning and problem-solving, the system must interact over multiple rounds of question-answering and gather information that is not immediately available. This requires strategically formulating questions in a sequence that progressively reduces uncertainty. As Large Language Models (LLMs) continue to push the boundaries of natural language understanding and generation, they provide the foundation for developing approaches to enhance this informationseeking capability of conversational systems.

LLMs have emerged as powerful tools for planning using trees Zhao et al. (2024); Putta et al. (2024); Hu et al. (2024). Notably, the *Tree of Thoughts* (ToT) Yao et al. (2023) approach has demonstrated the potential of leveraging hierarchical structures to improve problem-solving and reasoning in language models. However, expanding the full tree can be computationally inefficient and expensive, especially in domains with large search spaces. One of the fundamental problems lies in designing systems capable of dynamically adapting their interaction strategy to converge on a specific solution

<sup>&</sup>lt;sup>1</sup>Code: https://github.com/harshita-chopra/misq-hf



Figure 1: Overview of MISQ-HF. The system combines Monte Carlo Tree Search with hierarchical feedback for efficient information seeking. Green nodes (•) and red nodes (•) represent  $v^A$  and  $v^N$  respectively, corresponding to possibilities after affirmative and negative answers to a question. The purple arrows indicate the information-seeking question loop (steps 4-8). Towards the end of decision-making, targeted questions are generated, denoted by the loop of dashed arrows (steps 9-6-7-8). Finally, when the user response  $a_t$  (step 7) confirms the target detection, the conversation is terminated (step 10) and the feedback mechanism updates the reward bonus along the successful questioning trajectory in the decision tree, enabling the system to adapt its question selection strategy for similar future cases.

efficiently. To address these challenges, Monte Carlo Tree Search (MCTS) has emerged as a promising technique for balancing exploration and exploitation in decision-making processes. MCTS has been successfully applied in various domains, from game-playing AI Silver et al. (2016; 2018) to robotics Dam et al. (2022), and is now being leveraged in natural language processing tasks.

Building upon this foundation, our work introduces a novel framework that combines MCTS with a hierarchical feedback mechanism to achieve adaptive and efficient information-seeking. Inspired by prior work on uncertainty-aware planning, our approach presents three key contributions:

- MCTS on Decision Tree for Question Generation: Traditional tree-based approaches can struggle with the vast search spaces. By employing MCTS, our system efficiently constructs and explores promising branches of the decision tree without exhaustive expansion. We propose a variant of MCTS that optimizes question selection by considering both immediate and long-term information gain. We use the standard Upper Confidence Bound for Trees (UCT) and propose a modified variant that assigns higher weights to questions that have effectively narrowed the solution space in past interactions.
- 2. Cluster-Based Feedback Mechanism: We leverage similarity-based clustering using text embeddings of the initial problem description provided by the user. The system learns from historical successes by propagating cluster-specific bonus rewards through the decision tree after a successful interaction. This bonus reward decays with depth, emphasizing the importance of early-stage questions in narrowing the solution space by targeting high-level information that eliminates large subsets of possibilities. These questions are broadly applicable across multiple cases, making them valuable for generalizing strategies, whereas later-stage questions tend to focus on finer details unique to each interaction.
- 3. Efficiency through reduced LLM calls: Our method minimizes computational overhead by strategically limiting LLM interactions during tree construction and simulation. We maintain a single decision tree for a given dataset which expands as more data points stream in. We also employ lightweight rollout policies limited by predefined depth during the simulation step of MCTS. The exploration-exploitation balance avoids exhaustive tree expansion and significantly reduces the number of LLM calls without compromising performance.

These innovations collectively enable our system to dynamically adapt its questioning strategy based on historical patterns, achieving greater efficiency and relevance. For instance, the integration of feedback mechanisms ensures that successful strategies for specific clusters are reinforced over time, while depth-aware MCTS ensures computational resources are focused on the most promising paths. An overview of our framework, **MISQ-HF** (Monte Carlo Tree Search for Information Seeking Questions with Hierarchical Feedback), is illustrated in Figure 1.

We evaluate our approach across diverse conversational tasks. Results demonstrate that our system outperforms existing baselines in both task success and efficiency in scenarios requiring complex reasoning and hierarchical decision-making. We also highlight the individual contributions of depth-aware MCTS and cluster-based feedback in enhancing the system's performance.

### 2 RELATED WORK

The integration of LLMs with tree search algorithms has proven effective for enhancing multi-step reasoning and decision-making. Guez et al. (2018) introduced MCTSnets, combining tree search with neural networks for tasks like Sokoban. While it focuses on discrete planning, our approach extends MCTS to hierarchical conversations, optimizing question selection with depth-aware strategies and semantic embeddings. Yu et al. (2023) applied MCTS for dialogue planning with LLM-based simulations, but their method differs by using open-loop search without feedback mechanisms and focusing solely on persuasion tasks. Hui & Tu (2024) proposed the Reflection on Search Trees (RoT) framework for generating state-specific guidelines to improve search performance. Unlike RoT, which prioritizes reasoning efficiency, our approach maintains a single decision tree while incorporating cluster-specific rewards to guide question selection. Zhu et al. (2021) framed retrieval and answering as a partially observed Markov decision process to adaptively gather evidence but does not address learning from feedback in conversational settings. In the clinical domain, Li et al. (2024) introduced MEDIQ, an adaptive framework leveraging LLMs to identify missing information and ask follow-up questions. While MEDIQ focuses on reliability in high-stakes contexts, our method generalizes by prioritizing early-stage questions to narrow the solution space effectively.

Our work draws inspiration from the Uncertainty of Thoughts (UoT) algorithm by Hu et al. (2024), which enhances LLMs' ability to seek information through effective questioning. UoT uses uncertainty-based rewards driven by information gain and a reward propagation scheme to optimize question selection, improving performance across various domains and baselines like Chain of Thoughts Wei et al. (2022) and Tree of Thoughts Yao et al. (2023). Extending these principles, we integrate a depth-aware MCTS algorithm for question selection in hierarchical conversations and a cluster-based feedback system that leverages historical successes. The depth-aware bonus reward prioritizes early-stage questions, balancing efficiency and adaptability.

## 3 Methodology

### 3.1 PROBLEM FORMULATION

The task is modeled as a sequential interaction between a Questioner (an LLM) and an Answerer (a human; here, simulated by an LLM). Let  $\Omega$  be the space of possible target items, with an unknown target  $\omega \in \Omega$ . The interaction proceeds over turns  $t = 1, \ldots, T$ , where T is the maximum allowed turns. At each turn, the system asks a question  $q_t$  and receives an answer  $a_t$ , which is either binary (yes/no) or open-ended. The history  $h_t = \{q_1, a_1, \ldots, q_{t-1}, a_{t-1}\}$  represents all previous interactions. At each step i, the current possibility set  $\Omega_i$  is updated based on prior interactions. The updated set  $\Omega_i$  contains all elements consistent with the history  $h_t$ . Questions in the first  $\delta * T$  turns, where  $\delta \in (0, 1)$ , are selected strategically by traversing the decision tree of potential questions, and the remaining turns are reserved for making informed decisions about what the outcome is, based on  $h_i$  and  $\Omega_i$ . Here  $\delta$  denotes the proportion of turns reserved for asking information-seeking questions. The process ends when the Questioner identifies  $\omega$  or reaches the maximum number of turns T.

Upon reaching terminal states in the decision tree (where  $|\Omega_t| \leq 2$ ), the Questioner LLM transitions from asking information-seeking questions to making targeted questions about specific outcomes. In the **Closed Set** scenario, where the target is one of the items within a predefined set  $\Omega$ , the tree construction is well-defined. However, this space is often unknown in real-world situations, leading to an **Open Set** scenario, where the models operate without prior knowledge of the outcomes. To address this, we follow the approach used in Hu et al. (2024), where we directly prompt the Question Generator to first define the initial possibility space  $\Omega$  based on the problem description and then update it progressively based on the history of interactions  $h_t$ . In practice, it is important to declare the set of all possible items  $\Omega$  in the initial prompt for the Questioner LLM to avoid longer dialogues and failures due to wild guesses in the Closed Set scenario. Hence, we ensure that the Questioner LLM is explicitly informed about  $\Omega$  once at the beginning of each interaction. This helps constrain the items guessed in the targeting questions. To justify this modification, we report results with and without the  $\Omega$ -aware prompt.

#### 3.2 DECISION TREE OF QUESTIONS

The question generation process employs a hierarchical approach that combines LLMs with Monte Carlo Tree Search (MCTS) to efficiently construct a tree of potential questions represented by nodes. In the information-seeking phase, at each turn t, the system either traverses or expands the tree to find the most optimal questions to ask, with the goal of maximizing information gain. This tree is cached and re-used across all data-points belonging to a given dataset.

The LLM-based question generator takes two inputs: the current possibility set  $\Omega_t$  and the ancestral context  $\mathcal{A}_{\Omega t}$  which consists of the sequence of question-answer pairs corresponding to all the ancestor nodes of the current node that led to the formation of  $\Omega_t$ . This context prevents the generation of redundant or previously asked questions, thereby providing a clearer path of deduction. Let v represent a node in the decision tree at turn t. With slight abuse of notation, we define  $\Omega_v$  as the possibility set corresponding to node v. For each non-terminal node v, where  $|\Omega_v| > 2$ , the LLM generates m candidate questions which form its immediate child nodes:

$$Q_v = \{q_v^{(1)}, q_v^{(2)}, ..., q_v^{(m)}\} = LLM_{gen}(\Omega_v, \mathcal{A}_{\Omega v})$$
(1)

For each generated question,  $LLM_{gen}$  performs a binary partitioning of the possibility set  $\Omega_v$  into two disjoint subsets  $\Omega_v^A$  and  $\Omega_v^N$ , corresponding to affirmative and negative responses respectively, such that  $\Omega_v = \Omega_v^A \cup \Omega_v^N$ . The LLM is prompted to generate questions that maximize information gain by creating balanced partitions where:  $|\Omega_v^A| - |\Omega_v^N| \to 0$ .

Each question spawns two child nodes  $v^A$  and  $v^N$  corresponding to affirmative and negative responses, storing  $\Omega_v^A$  and  $\Omega_v^N$  respectively. For example, if the possibility set is  $\Omega_v = \{$ flu, pneumonia, enteritis, asthma $\}$ , the question "Do you have difficulty breathing?" might partition the set as:  $\Omega_v^A = \{$ pneumonia, asthma $\}, \quad \Omega_v^N = \{$ flu, enteritis $\}.$ 

#### 3.3 INFORMATION GAIN

To quantify the effectiveness of each question in reducing uncertainty, entropy-based metrics are employed Hu et al. (2024). Following the formulation of reward structure in UoT, let  $p_v^A = |\Omega_v^A|/|\Omega_v|$  and  $p_v^N = 1 - p_v^A$ . The expected information gain at v is:

$$IG_v(X) = -p_v^A \log p_v^A - p_v^N \log p_v^N$$
<sup>(2)</sup>

The reward function, given by  $R_{IG}(v)$ , achieves its maximum value when the subsets  $\Omega_v^A$  and  $\Omega_v^N$  have equal probabilities, signifying the greatest reduction in uncertainty. It attains its minimum value when one of the subsets has a probability of zero, indicating no reduction in uncertainty. To normalize and sharpen rewards, a scaling parameter  $\lambda > 0$  is employed. Formally,

$$R_{IG}(v) = \frac{-p_v^A \log p_v^A - p_v^N \log p_v^N}{1 + \lambda^{-1} |p_v^A - p_v^N|}.$$
(3)

The expected reward of asking a question,  $R_e(v)$ , is obtained by recursively adding the immediate reward  $R_{IG}(v)$  (information gain at the current node) and the expected rewards of its child nodes. The total information gained until node v of an interaction trajectory is denoted by the accumulated reward,  $R_a(v)$ , which is obtained by starting at the root and propagating down to the node v.

$$R_e(v) := \begin{cases} R_a(v), & v \text{ is a leaf,} \\ p_v^A R_e(v^A) + p_v^N R_e(v^N) & \text{otherwise.} \end{cases}$$
(4)

$$R_a(v) := R_{IG}(v) + \begin{cases} 0 & v \text{ is root,} \\ R_a(parent(v)) & \text{otherwise.} \end{cases}$$
(5)

where  $v^A$  and  $v^N$  are the child nodes of v corresponding to affirmative and negative responses respectively. For each non-terminal response-specific node  $v^{\alpha}$ ,  $\alpha \in \{A, N\}$ , the expected reward  $R_e(v^{\alpha})$  is the average expected reward of its child nodes, denoted as *children*( $v^{\alpha}$ ). Formally,

$$R_e(v^{\alpha}) = \frac{1}{|children(v^{\alpha})|} \sum_{v' \in children(v^{\alpha})} R_e(v')$$
(6)

#### 3.4 MONTE CARLO TREE SEARCH

The decision tree of questions is constructed iteratively when MCTS is executed over K iterations at each decision-making step across samples. Each iteration includes four phases:

**1. Selection:** Considering the current node as the root, a child node is selected based on the widely used UCT formulation that balances exploration and exploitation:

$$UCT(v) = \frac{R_{\text{total}}(v)}{N_v} + C\sqrt{\frac{\ln N_p}{N_v}}$$
(7)

where  $R_a(v)$  is the accumulated reward,  $N_v$  is the visit count of node v,  $N_p$  is its parent node's visit count, and C is the exploration constant.

**2. Expansion:** The expansion step aims to create child nodes by generating potential questions that split the possibility set, thereby narrowing down the search space of the tree. When the selected node is non-terminal ( $|\Omega_v| > 2$ ) and does not have child nodes, candidate questions are formulated according to Equation 1. If child nodes already exist, we skip this step and proceed to simulation.

**3.** Simulation: A rollout policy is used to estimate the expected reward of the selected node by simulating a random interaction up to a predefined depth  $d_s$  or a terminal state, whichever occurs first. In the rollout phase, we traverse  $d_s$  levels down from the selected node, choosing one of its child nodes at random, following a single level of expansion if child nodes did not exist. This step provides an estimate of the long-term utility of selecting a particular question.

**4. Backpropagation:** After reaching a leaf node u at the end of the simulation, the expected reward  $R_e(u)$  is backpropagated up the tree to update the ancestor nodes along the path to the selected (simulated) node. During backpropagation, the cumulative reward,  $R_{\text{total}}(u)$  of every ancestor node of u is updated, and their visit counts are incremented by 1. Formally, for each node v in the path from u to the root:

$$R_{\text{total}}(v) \leftarrow R_{\text{total}}(v) + R_e(u) \tag{8}$$

This backpropagation process ensures that nodes closer to successful outcomes, i.e., nodes leading to higher rewards, are more likely to be selected in subsequent iterations.

Finally, after completing K iterations of MCTS, the system asks the question having the highest expected information gain. This corresponds to finding  $v^*$ , the child node of the current root having the maximum expected reward:

$$v^* = \arg \max_{v' \in \mathbb{C}(\text{root})} R_e(v') \tag{9}$$

#### 3.5 FEEDBACK MECHANISM AND CLUSTER-BASED REWARD ADJUSTMENT

Learning from experience is a critical component of decision-making systems. We introduce a clustering-based approach for dynamic reward adjustments by propagating feedback through the tree after a successful conversation. When the system successfully identifies the target, it triggers a series of updates to the bonus rewards of nodes that led to the target.

#### 3.5.1 CLUSTER CREATION AND ASSIGNMENT

When a new sample arrives, its text embedding is computed and compared with existing cluster medoids to determine its assignment. A similarity threshold  $\tau$  is used to decide whether to assign the embedding to an existing cluster or create a new one. Let e represent the embedding of the current data point and  $\mathbf{m}_k$  denote the medoid of cluster k. The similarity condition is given by:

Similarity(
$$\mathbf{e}, \mathbf{m}_k$$
) =  $\frac{\mathbf{e} \cdot \mathbf{m}_k}{\|\mathbf{e}\|_2 \|\mathbf{m}_k\|_2} \ge \tau$  (10)

where  $\cdot$  represents the dot product, and  $\|\cdot\|_2$  represents the L2 norm. If no existing cluster satisfies this condition, a new cluster is created with e as its initial medoid. If assigned to an existing cluster, e is added to that cluster's set of embeddings  $C_k$ , and the medoid  $\mathbf{m}_k$  is recomputed as:

$$\mathbf{m}_{k} = \operatorname*{argmax}_{\mathbf{x} \in \mathcal{C}_{k}} \sum_{\mathbf{y} \in \mathcal{C}_{k}} \frac{\mathbf{x} \cdot \mathbf{y}}{\|\mathbf{x}\|_{2} \|\mathbf{y}\|_{2}}$$
(11)

#### 3.5.2 BONUS REWARDS FOR FEEDBACK PROPAGATION

For each node v, we maintain a dictionary  $B_k(v)$  which maps each cluster k to the bonus reward (initialized as zero) earned on reaching the target successfully. Once a target is correctly identified,  $B_k(v)$  is updated for each node along the path from the current node back to the root.  $B_k(v)$  is a proportion of the node's cumulative reward and depends on the absolute depth of the node, given by  $d_v$  controlled by an exponential decay as we move upward in the tree. The bonus reward for node v corresponding to cluster k is updated after every success:

$$B_k(v) \leftarrow B_k(v) + \beta \cdot R_{\text{total}}(v) \cdot \gamma^{d_v} \tag{12}$$

where  $\beta$  is a task-specific scaling factor (denoting proportion of total reward), and  $\gamma \in (0, 1)$  is a decay factor that controls the influence of the bonus, keeping it higher for early-stage questions which have proven effective to narrow down the large possibility set, and lower for questions near the terminal nodes which are often more specific to each case.

To incorporate these bonus rewards into decision-making, we modify the UCT formula used in the selection step to handle similar cases while maintaining exploration capabilities. We introduce a cluster-specific bonus term to bias selection toward nodes associated with successful outcomes for similar data points. The modified UCT formula becomes:

$$UCT_{\text{feedback}}(v,k) = \frac{R_{\text{total}}(v)}{N_v} + C\sqrt{\frac{\ln N_p}{N_v}} + B_k(v)$$
(13)

This modification prioritizes nodes from clusters associated with successful outcomes when processing samples similar to their medoids. Algorithm 1 shows the pseudocode of the proposed framework.

#### 4 **EXPERIMENTS**

#### 4.1 DATASETS

We demonstrate the effectiveness of our approach by focusing on three diverse domains: Medical Diagnosis, Troubleshooting, and 20 Questions. We use the following datasets pre-processed by Hu et al. (2024). In Medical Diagnosis, a patient initially reports a brief description of their symptoms, based on which the doctor asks questions to diagnose the disease. The maximum number of turns T was limited to 6 in the experiments. Two datasets were used. The **DX** dataset by Xu et al. (2019) contains 104 doctor-patient dialogues and five diseases in its test set. The MedDG dataset, which originally included over 17,000 conversations across 15 disease types, was refined by removing inconsistent samples. We used 454 high-quality samples for evaluation. Open-ended responses are allowed in MedDG to test the system's generalization capabilities in less constrained scenarios. Both datasets limit interactions to 6 turns. In the Troubleshooting domain, customer support technicians interact with users to identify faults in systems such as cars or electronic devices. We use the FloDial dataset by Raghu et al. (2021), containing 153 dialogues across 153 unique fault types. Maximum number of turns T was limited to 20. In the 20 Questions domain, the task involves identifying a target item by asking up to 20 yes-or-no questions. The Common dataset includes 111 items spanning categories such as animals, places, food, and objects, and the Things dataset Hebart et al. (2019), was filtered to include 200 items.

#### 4.2 **BASELINES**

The proposed system is compared with two primary baselines: Direct Prompting (**DP**) and the Uncertainty of Thoughts (**UoT**) framework Hu et al. (2024). DP directly queries the LLM to generate the next question without any structured planning, serving as a minimal-effort benchmark to

#### Algorithm 1 MISQ-HF **Require:** Dataset S, question ratio $\delta$ , Embedding model, cluster-embeddings hashmap C, similarity threshold $\tau$ , maximum turns T, MCTS iterations K, LLM 1: Initialize $B_k(v) = 0, \forall k \in \{1, 2, ..., |\mathcal{C}|\}$ for each node v 2: for sample $s \in S$ do 3: $\mathbf{e}_s \leftarrow \text{Embedding}[description(s)]$ 4: $C_k \leftarrow \text{AssignCluster}(\mathbf{e}_s, C, \tau)$ ⊳ Use Eq. 10-11 5: Initialize $\Omega_v = \Omega$ at current root node v 6: Initialize conversation history $h = \emptyset$ and t = 07: while t < T and target(s) undetected do if $t < \delta * T$ and $|\Omega_v \geq 2|$ then 8: $v^* \leftarrow \mathrm{MCTS}(v, K, \mathcal{C}_k)$ 9: ⊳ Use Eq. 3-8, 13 10: $q_t \leftarrow question(v^*)$ ⊳ Use Eq. 9 11: $a_t \leftarrow \text{UserResponse}(q_t, h)$ $v \leftarrow child(v^*|a_t)$ 12: 13: else 14: $q_t \leftarrow \text{TargetingPrompt}(\Omega_v)$ 15: $a_t \leftarrow \text{UserResponse}(q_t, h)$ 16: end if 17: $h \leftarrow h \cup \{q_t, a_t\}$ 18: $t \leftarrow t + 1$ 19: end while 20: if *target*(s) *detected* then Successful conversation 21: ▷ Feedback Propagation $v' \leftarrow v$ 22: while $parent(v') \neq \emptyset$ do 23: Update $B_k(v')$ ⊳ Use Eq. 12 24: $v' \leftarrow parent(v')$ 25: end while 26: end if 27: end for

highlight improvements achieved by strategic decision-making. UoT uses tree-based planning with exhaustive expansion of all question paths to maximize information gain. While effective in reducing uncertainty, its full tree exploration leads to significant computational overhead, especially in large possibility spaces. This baseline assesses the efficiency improvement of the proposed MISQ-HF approach, which selectively expands promising branches. Comparisons with Chain-of-Thought (CoT) Wei et al. (2022) and Tree-of-Thoughts (ToT) Yao et al. (2023) are omitted, as prior work Hu et al. (2024) shows UoT consistently outperforms these methods.

#### 4.3 EVALUATION

We use three key metrics to evaluate the question-asking system's effectiveness and efficiency. First, Success Rate (SR) measures the percentage of cases where the system correctly identifies the target item within the maximum allowed turns. Second, tracks the average number of turns needed to reach the correct answer in successful interactions. Finally, we introduce a novel metric to track Question Generation Calls (QGC), defined as the number of LLM prompting calls required for question generation. Lower values indicate better computational efficiency. Together, these metrics provide a comprehensive view of the system's performance, balancing success rate, conversation efficiency, and computational overhead.

The frequency of QGC for exhaustive tree search and MCTS differs significantly. An LLM call creates m questions, each splitting  $\Omega$  into affirmative and negative response sets, thereby giving rise to 2m nodes. For UoT, the first simulation requires  $1 + 2m + (2m)^2 + \dots + (2m)^{d_s} = \frac{(2m)^{d_s+1}-1}{2m-1}$  calls, and  $(2m)^{d_s}$  calls for subsequent turns if paths are expanded till  $d_s - 1$ , where  $d_s$  is the depth-limit for simulation. For MCTS, the first expansion requires 1 call, and simulation for selecting a question at each turn requires  $K \cdot d_{avg}$ , where K is the number of iterations and  $0 \le d_{avg} \le d_s$  because traversing requires 0 LLM calls when the nodes are already expanded. Exhaustive search scales exponentially with  $d_s$ , while MCTS scales linearly, highlighting its efficiency for deeper searches in the simulation phase.

Model	Method	$\Omega$ -aware	DX		MedDG			FloDial			
Model			SR↑	MSC↓	QGC↓	SR↑	MSC↓	$QGC\downarrow$	SR↑	MSC↓	QGC↓
	UoT	×	72.11	1.54	0.36	79.51	2.09	4.95	34.64	6.84	43.76
	MISQ	×	75.00	2.17	0.05	86.56	3.39	0.40	35.29	9.09	3.99
Llama 3 3	MISQ-HF	×	80.76	1.94	0.21	86.78	3.29	0.78	39.86	9.09	4.07
70B Instruct	DP	$\checkmark$	88.46	3.15	-	84.14	3.93	-	21.56	13.72	
	UoT	$\checkmark$	79.80	1.65	0.77	89.86	2.16	4.84	60.78	8.47	44.61
	MISQ	$\checkmark$	92.30	1.28	0.48	92.29	3.44	3.59	62.74	9.73	5.16
	MISQ-HF	$\checkmark$	98.07	1.84	0.04	93.39	3.35	0.54	67.97	9.81	3.97
	DP	$\checkmark$	50.00	3.50	-	76.43	3.91	-	16.99	14.23	_
Mixtral	UoT	$\checkmark$	76.92	1.43	0.45	83.70	2.19	5.70	39.21	7.01	45.11
8*7B Instruct	MISQ	$\checkmark$	63.46	2.63	0.08	76.55	3.33	0.17	47.71	10.45	1.66
	MISQ-HF	$\checkmark$	76.92	2.40	0.06	84.58	3.08	0.33	49.01	9.62	1.46
GPT-40	DP	$\checkmark$	73.07	3.48	-	81.27	3.98	-	43.79	14.86	_
	UoT	$\checkmark$	82.69	1.18	0.17	88.79	2.03	1.81	59.47	8.14	41.86
	MISQ	$\checkmark$	87.50	1.97	0.03	89.20	3.46	0.60	74.50	10.15	4.10
	MISQ-HF	$\checkmark$	99.03	2.19	0.05	90.30	3.42	0.41	72.54	10.36	2.94

Table 1: Results on Medical Diagnosis and Troubleshooting Domain in a Closed Set scenario.  $\Omega$ *aware* denotes whether the Questioner Model was informed about the possibility space  $\Omega$  once at the beginning of the conversation.

#### 4.4 EXPERIMENTAL SETUP

**Models.** We employed three different LLMs of varying sizes as the Questioner. Llama 3.3 70B Instruct Grattafiori et al. (2024) and Mixtral 8\*7B Instruct Jiang et al. (2024) were accessed via the AWS Bedrock API (2023). GPT-40 was accessed via API provided by OpenAI et al. (2024). The user (Answerer) was simulated by Llama 3.3 70B Instruct in all tasks. We prompted the model with ground truth details (e.g., the patient's disease or fault description) and maintained a separate conversation history. The temperature was set to 0. The ratio of turns was defined by  $\delta = 0.6$ .

**MCTS implementation.** We set the number of iterations K = 10 and exploration constant C = 0.2. The maximum simulation depth  $d_s$  was set to 3 to balance computational efficiency with search effectiveness. For each  $\Omega_v$ , the question generator creates m = 3 potential questions to maintain diversity. For the reward calculation in  $R_{IG}(v)$ , the scaling parameter  $\lambda$  was set to 0.4.

**Feedback mechanism.** We used a decay factor  $\gamma = 0.9$  for the bonus rewards. The cluster similarity threshold  $\tau$  was set to 0.9 in terms of cosine similarity, and the bonus scaling factor  $\beta$  was set to 0.2 for all tasks. Problem descriptions were embedded using DistilBERT Sanh et al. (2019) for the troubleshooting domain, and Clinical-BERT Wang et al. (2023) for medical diagnosis. To optimize computational efficiency, we cache the decision tree, allowing reuse across all samples within each dataset. The system tracks Mean QGC through a thread-safe counter.

### 5 RESULTS AND DISCUSSION

Table 1 shows our results in the Closed Set scenario. Our approach requires fewer QGC while maintaining higher SR compared to baseline methods, indicating both improved effectiveness and computational efficiency. Notably, the  $\Omega$ -aware Questioner consistently performs better as compared to when unaware. To demonstrate the importance of the feedback mechanism, we also report results using MISQ, our framework without the hierarchical feedback component, to rationalize the design choice and support an ablation study.

In the Medical Diagnosis domain, MISQ-HF achieves superior performance across both datasets, with an average reduction of 6.77 times in Mean QGC across all three LLMs. Using Llama 3.3 70B as a Questioner, MISQ-HF required only 0.04 Mean QGC compared to UoT's 0.77 on DX, showing a 19.25 times reduction. Similarly, a reduction of 8.97 times is observed on MedDG. This efficiency stems from selective tree expansion rather than exhaustive exploration. In the medical domain, the cost of misdiagnosis far outweighs the overhead of a few additional clarifying questions. While MSC is marginally higher for MISQ-HF compared to UoT, this trade-off can be overseen by the higher success rate of diagnosis, which directly impacts patient outcomes.

Method	Ω-		Common			Thing		
memou	aware	SR↑	MSC↓	QGC↓	SR↑	MSC↓	QGC↓	
Llama 3.3	3 70B Ins	truct						
UoT	×	39.63	8.27	4.08	19.00	9.78	4.48	
MISQ	×	41.44	8.43	5.05	23.5	9.57	1.57	
DP	$\checkmark$	45.94	13.70	-	32.50	13.27	-	
UoT	$\checkmark$	61.26	9.94	7.92	35.50	11.43	3.40	
MISQ	$\checkmark$	74.77	9.90	4.74	59.50	10.68	3.31	
Mixtral 8	*7B Instr	uct						
DP	$\checkmark$	8.10	14.33	-	7.50	13.46	-	
UoT	$\checkmark$	28.82	11.56	4.34	12.50	13.52	5.91	
MISQ	$\checkmark$	37.83	11.38	2.39	20.00	11.50	0.06	
GPT-40								
DP	$\checkmark$	63.06	14.72	-	40.50	14.16	-	
UoT	$\checkmark$	74.77	8.59	5.88	47.00	9.13	2.75	
MISQ	$\checkmark$	85.58	8.51	4.86	55.50	9.54	2.19	

Table 2: Results on 20 Questions Datasets in the Closed Set scenario

In the Troubleshooting domain using the FloDial dataset, MISQ-HF continues to excel with higher SR, surpassing UoT by 20.68% improvement on average across the three LLMs. Our framework also achieves notable computational efficiency as compared to UoT, with an average of 18.63 times reduction in Mean QGC. While troubleshooting conversations typically require more turns due to fault complexity, MISQ-HF maintains a similar MSC (a difference of less than 3 turns on average) while delivering improved performance. With GPT-40, MISQ-HF achieved an SR of 72.54% with 27 clusters, the largest consisting of 31 samples.

Table 2 shows performance on the 20 Questions domain, evaluated using the Common and Things datasets. Due to a lack of problem description or initial hints about the target, we do not evaluate MISQ-HF here. For both datasets, MISQ consistently outperforms UoT and DP in all three metrics. These results demonstrate the scalability of the proposed approach to larger possibility spaces.

Results in the Open Set scenario are provided in Table 3 of Appendix. Our experiments show that the MISQ-HF outperforms other baselines on datasets with different sizes of possibility spaces. Its hierarchical feedback mechanism enables the system to learn from successful questioning trajectories in past conversations, especially in domains like troubleshooting and medical diagnosis, where similar cases often recur. Compared to exhaustive methods like UoT, MISQ-HF achieves ~10x reduction in Mean QGC while delivering ~12% improvement in SR across these domains. Performance without hierarchical feedback (MISQ) further emphasizes the benefits of this approach. Example conversations and prompts are provided in the Appendix.

**Limitations.** The system currently does not incorporate a mechanism to learn from mistakes in failure cases, which could potentially refine its decision-making process. Penalizing suboptimal or redundant questions also requires careful design of reward function and remains an open challenge. Incorporating confidence metrics to better quantify uncertainty in risk-sensitive domains is another potential improvement. Addressing these limitations is a promising direction for future work.

### 6 CONCLUSION

In this paper, we introduced Monte Carlo Tree Search for Information Seeking Questions with Hierarchical Feedback (MISQ-HF), a novel framework for adaptive question-asking with a goal to optimize conversational efficiency and accuracy. By leveraging selective tree expansion and clusterbased feedback, our approach reduces computational overhead while maintaining high success rates across diverse domains such as medical diagnosis, troubleshooting, and general information-seeking tasks. We demonstrate the potential of combining structured and efficient planning with historical learning to dynamically adapt question-asking strategies. Our contributions aim to refine the decision-making process, enabling conversational AI systems to handle increasingly complex and open-ended scenarios with better planning.

#### ACKNOWLEDGMENTS

This work is supported by the National Science Foundation (NSF) award IIS-2336769. The authors thank Hafsah Usman, student of the University of Washington, Seattle, for her valuable assistance in running a part of the experiments.

#### REFERENCES

- AWS Bedrock API. Build Generative AI Applications with Foundation Models Amazon Bedrock AWS, 2023. URL https://aws.amazon.com/bedrock/.
- Tuan Dam, Georgia Chalvatzaki, Jan Peters, and Joni Pajarinen. Monte-Carlo Robot Path Planning. *IEEE Robotics and Automation Letters*, 7(4):11213–11220, October 2022. ISSN 2377-3766. doi: 10.1109/LRA.2022.3199674. URL https://ieeexplore.ieee.org/document/9860036.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, and Austen Gregerson et al. The Llama 3 Herd of Models, November 2024. URL http://arxiv.org/abs/2407.21783. arXiv:2407.21783.
- Arthur Guez, Theophane Weber, Ioannis Antonoglou, Karen Simonyan, Oriol Vinyals, Daan Wierstra, Remi Munos, and David Silver. Learning to search with MCTSnets. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 1822–1831. PMLR, July 2018. URL https://proceedings.mlr.press/v80/guez18a.html.
- Martin N. Hebart, Adam H. Dickter, Alexis Kidder, Wan Y. Kwok, Anna Corriveau, Caitlin Van Wicklin, and Chris I. Baker. THINGS: A database of 1,854 object concepts and more than 26,000 naturalistic object images. *PLOS ONE*, 14(10):e0223792, October 2019. ISSN 1932-6203. doi: 10.1371/journal.pone.0223792. URL https://journals.plos.org/plosone/ article?id=10.1371/journal.pone.0223792.
- Zhiyuan Hu, Chumin Liu, Xidong Feng, Yilun Zhao, See-Kiong Ng, Anh Tuan Luu, Junxian He, Pang Wei Koh, and Bryan Hooi. Uncertainty of Thoughts: Uncertainty-Aware Planning Enhances Information Seeking in LLMs. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, November 2024. URL https://openreview.net/forum? id=CVpuVe1N22&noteId=aTI8PGp047.
- Wenyang Hui and Kewei Tu. RoT: Enhancing Large Language Models with Reflection on Search Trees, July 2024. URL http://arxiv.org/abs/2404.05449. arXiv:2404.05449 [cs].
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, Lélio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mixtral of Experts, January 2024. URL http://arxiv.org/abs/2401.04088. arXiv:2401.04088 [cs].
- Shuyue Stella Li, Vidhisha Balachandran, Shangbin Feng, Jonathan S. Ilgen, Emma Pierson, Pang Wei Koh, and Yulia Tsvetkov. MediQ: Question-Asking LLMs and a Benchmark for Reliable Interactive Clinical Reasoning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, November 2024.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and Irwan Bello et al. GPT-4 Technical Report, March 2024. URL http://arxiv.org/abs/2303.08774. arXiv:2303.08774 [cs].

- Pranav Putta, Edmund Mills, Naman Garg, Sumeet Motwani, Chelsea Finn, Divyansh Garg, and Rafael Rafailov. Agent Q: Advanced Reasoning and Learning for Autonomous AI Agents, August 2024. URL http://arxiv.org/abs/2408.07199. arXiv:2408.07199.
- Dinesh Raghu, Shantanu Agarwal, Sachindra Joshi, and Mausam. End-to-End Learning of Flowchart Grounded Task-Oriented Dialogs. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 4348–4366, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main. 357. URL https://aclanthology.org/2021.emnlp-main.357/.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *ArXiv*, abs/1910.01108, 2019.
- David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, January 2016. ISSN 1476-4687. doi: 10.1038/nature16961. URL https://www.nature.com/articles/ nature16961.
- David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144, December 2018. doi: 10. 1126/science.aar6404. URL https://www.science.org/doi/10.1126/science. aar6404.
- Guangyu Wang, Xiaohong Liu, Zhen Ying, Guoxing Yang, Zhiwei Chen, Zhiwen Liu, Min Zhang, Hongmei Yan, Yuxing Lu, Yuanxu Gao, Kanmin Xue, Xiaoying Li, and Ying Chen. Optimized glycemic control of type 2 diabetes with reinforcement learning: a proof-of-concept trial. *Nature Medicine*, 29(10):2633–2642, October 2023. ISSN 1546-170X. doi: 10.1038/s41591-023-02552-9. URL https://www.nature.com/articles/s41591-023-02552-9.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, pp. 24824–24837, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 978-1-71387-108-8.
- Lin Xu, Qixian Zhou, Ke Gong, Xiaodan Liang, Jianheng Tang, and Liang Lin. End-to-End Knowledge-Routed Relational Dialogue System for Automatic Diagnosis. *Proceedings of the* AAAI Conference on Artificial Intelligence, 33(01):7346–7353, July 2019. ISSN 2374-3468. doi: 10.1609/aaai.v33i01.33017346. URL https://ojs.aaai.org/index.php/AAAI/ article/view/4722.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), Advances in Neural Information Processing Systems, volume 36, pp. 11809–11822, 2023. URL https://proceedings.neurips.cc/paper\_files/paper/2023/ file/271db9922b8d1f4dd7aaef84ed5ac703-Paper-Conference.pdf.
- Xiao Yu, Maximillian Chen, and Zhou Yu. Prompt-Based Monte-Carlo Tree Search for Goaloriented Dialogue Policy Planning. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 7101–7125, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.439. URL https://aclanthology.org/2023. emnlp-main.439/.

- Zirui Zhao, Wee Sun Lee, and David Hsu. Large language models as commonsense knowledge for large-scale task planning. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, pp. 31967–31987, Red Hook, NY, USA, May 2024. Curran Associates Inc.
- Yunchang Zhu, Liang Pang, Yanyan Lan, Huawei Shen, and Xueqi Cheng. Adaptive Information Seeking for Open-Domain Question Answering. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 3615–3626, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021. emnlp-main.293. URL https://aclanthology.org/2021.emnlp-main.293/.

## Appendix

### A OPEN SET RESULTS

The results in the Open Set scenario are demonstrated on two datasets where the range of possible outcomes is often unpredictable and varies in size. Problem descriptions were used to initialize  $\Omega$  to a set of 5 possibilities, and it was progressively updated based on conversation history. We used Flo-Dial for troubleshooting, which has a larger possibility space, and DX for medical diagnosis, which has a smaller possibility space. Table 3 shows consistently lower Mean QGC while maintaining similar SR in both domains.

Table 3: Results in the Open Set scenario. Possibility set  $\Omega$  is unknown to the Questioner LLM.

Method		FloDial			DX	
	SR↑	MSC↓	QGC↓	SR↑	MSC↓	QGC↓
DP	16.99	14.80	-	29.80	3.22	-
UoT	28.10	7.56	11.42	35.57	2.35	10.22
MISQ	28.10	8.02	5.33	36.53	2.73	6.52
MISQ-HF	28.75	6.95	5.10	37.50	2.35	6.32

### **B** CONVERSATION PROMPTS FOR 20 QUESTIONS

#### **B.1 QUESTIONER PROLOGUE**

Let us play the game of 20 questions. I am impersonating the thing, X. {inform\_set} You will ask me up to 20 questions which start with 'Is X' and can only be answered by yes or no, and I will answer each one truthfully based on being X. Let us begin. Ask me the first question.

{inform_set}	"X is possibly one of the following: $\Omega$ "
	This is given only once at the beginning of the conversation.

### B.2 ANSWERER/USER SIMULATOR PROLOGUE

Let us play the game of 20 questions. You are the answerer and I am questioner. X is '{target\_item}'. I will ask you up to 20 questions and you should answer each one truthfully based on being X, by saying Yes or No.

Note that you must never reveal X, until I guess it correctly.

If I guess X correctly in my question, directly respond "You guessed it. X is '{target\_item}'." instead of saying yes. Let us begin. Here is my first question.

{target_item}	the ground truth target item based on which user responds
---------------	---

#### **B.3** TARGETING PROMPT

Note that you should guess and ask what X exactly is from now on. X is possibly one of the following: {item\_set\_terminal}, or other from the initial set of possibilities. Use the past conversation for context to avoid redundant questions. The question must start with 'Is X ...'

{item\_set\_terminal}  $\Omega_v$  (the set of items remaining at leaf node v)

**Target Question**: Is X '{item}'?

B.4 QUESTION GENERATION PROMPT

Here are all the X: {item\_set}

Design a question about X that can only be answered by Yes or No. Then classify the possible X above based on this question. If the answer is 'YES', put this X into 'YES: ...', otherwise to 'NO: ...'. Finally calculate how many X in YES and NO. Notably, this question should fulfill that the count of YES and NO are almost the same with a permissible discrepancy of no more than one!

{ancestral\_context}

Based on this information, create most relevant  $\{m\}$  questions to classify the above X correctly. Your response should strictly follow the template:

Question 1: Is X ...? YES: comma-separated, list of things, ... Count of YES: ... NO: comma-separated, list of things, ... Count of NO: ...

{item_set}	$\Omega_v$ (the set of items remaining at node $v$ )
{ancestral_context}	"For context, following questions were already asked to build the above set of possibilities: Question? No; Question? Yes; Question? Yes" if ancestral question nodes exist, else blank string.
{m}	number of potential questions to generate (actions)

### C CONVERSATION PROMPTS FOR MEDICAL DIAGNOSIS

#### C.1 QUESTIONER PROLOGUE

You are a doctor, and your patient self-reports that: {problem\_description}. {inform\_set} You should ask your patient questions in English about symptoms which can only be answered by Yes or No, in order to find what disease this patient suffers from. Use the ongoing conversation for context to avoid redundant questions. Let us begin. Ask me the first question.

### C.2 ANSWERER/USER PROLOGUE

You are the patient suffering from '{target\_item}', and I am the doctor. I will ask you up to 6 questions, and you should answer each one truthfully based on your disease, by saying Yes or No. Note that you must never reveal the disease until I tell it correctly. If I tell the disease correctly in my question, directly respond: "You are right. I am experiencing '{target\_item}'." Let us begin. Here is my first question.

#### C.3 QUESTION GENERATION PROMPT

You are a doctor. Here are all the possible diseases that the patient may suffer from: {item\_set}

Design a question to ask your patient regarding symptoms of their illness that can only be answered by Yes or No. Then classify the possible diseases above based on this question. If the answer is 'YES', put this disease into 'YES: ...', otherwise to 'NO: ...'. Finally calculate how many diseases are in YES and NO. Notably, this question should fulfill that the count of YES and NO are almost the same with a permissible discrepancy of no more than one!

{ancestral\_context}
Based on this information, create most relevant {n} questions to ask (and classify the above diseases). Your response should strictly follow the template:
Question 1: ...?
YES: comma-separated, list of disease names, ...
Count of YES: ...

NO: comma-separated, list of disease names, ... Count of NO: ...

#### C.4 TARGETING PROMPT

Note that you should point out and ask what disease the patient suffers from now. The patient is probably suffering from one of the diseases below: {item\_set\_terminal}, or other from the initial set of possibilities.

Refer to the past conversation regarding the patient's symptoms. Never repeat previously asked questions. Respond with 1 new question only, following the format: "Are you experiencing [disease name]?" Ensure that [disease name] was not asked before.

**Target Question**: Are you experiencing '{item}'?

#### C.5 CLASSIFY PROMPT

Here are all diseases that the patient may suffer from: {item\_set}

{problem\_description} For each disease under this report, if the patient is possibly affected, put this disease into 'YES: ...', otherwise to 'NO: ...'. Your answer should strictly follow the template:

YES: comma-separated, list of disease names, ...

NO: comma-separated, list of disease names, ...

### C.6 OPEN SET PROMPTS

#### **Initial Prompt:**

You are a doctor, and your patient self-reports that: {problem\_description}. Please propose {size} diseases that you think your patient may suffer from. Your response should only be a list like: ["disease1", "disease2", ...]

#### **Renewal Prompt:**

Based on the conversation history, please propose {size} diseases that your patient may suffer from. The list of {size} diseases should contain {existing\_items}. Your response should only be a list like: ["disease1", "disease2", ...]

### D CONVERSATION PROMPTS FOR TROUBLESHOOTING

### D.1 QUESTIONER PROLOGUE

You are a technician, and your client self-reports that: {problem\_description}. {inform\_set} You should ask your client questions with specific situations which can only be answered by Yes or No, in order to find which issue this client is facing. Use the ongoing conversation for context to avoid redundant questions. Let us begin. Ask me the first question.

#### D.2 ANSWERER/USER PROLOGUE

You are the client with a device that has '{target\_item}' and I am the technician. I will ask you up to 20 questions, and you should answer each one truthfully based on the issue of your device, by saying Yes or No. Note that you must never reveal the issue name until I tell it correctly. If I tell your issue correctly in my question, directly respond: "You are right. My device has issues with '{target\_item}'." Let us begin. Here is my first question.

#### D.3 QUESTION GENERATION PROMPT

You are a technician. Here are all the issues that the client may face: {items\_set}

Design a question to ask your client with a specific situation that can only be answered by YES or NO. Then classify the possible issues above based on this question. If the answer is 'YES', put this issue into 'YES: ...', otherwise to 'NO: ...'. Finally calculate how many issues are in YES and NO. Notably, this question should fulfill that the count of YES and NO are almost the same with a permissible discrepancy of no more than one!

{ancestral\_context} Based on this information, create the most relevant {n} questions to classify the above issues correctly. Your response should strictly follow the template:

Question 1: ...? YES: comma-separated, list of issue names, ... Count of YES: ... NO: comma-separated, list of issue names, ... Count of NO: ...

#### D.4 TARGETING PROMPT

Note that you should now point out and ask what issue the client is facing. Refer to the past conversation for context and never ask redundant questions. Respond with 1 new question only, following the format: "Are you experiencing [issue name]?" Ensure that [issue name] was not asked before.

#### **Target Question**: Are you experiencing '{item}'?

#### D.5 CLASSIFY PROMPT

Here are all the issues that the client may face: {item\_set}

{problem\_description} For each issue under this report, if the client is possibly facing it, put this issue into 'YES: ...', otherwise to 'NO: ...'. Your answer should strictly follow the template: YES: comma-separated, list of issue names, ... NO: comma-separated, list of issue names, ...

D.6 OPEN SET PROMPTS

#### **Initial Prompt:**

You are a technician, and your client self-reports that: {problem\_description}. Please propose {size} issues that you think your client may face with. Your response should only be a list like: ["issue1", "issue2", ...]

#### **Renewal Prompt:**

Based on the conversation history, please propose {size} issues that your client may face with. The list of {size} issues should contain {existing\_items}. Your response should only be a list like: ["issue1", "issue2", ...]

### E EXAMPLES

Two sequences of history are maintained respectively for Questioner LLM and User (Answerer LLM). Each stores the questions asked and answers given so far. At the beginning of the conversa-

tion, the Question LLM is given the Questioner Prologue as the first prompt, and the Answerer LLM is given the Answerer Prologue as the first prompt. The rest of the conversation history remains the same, except for the switching between the "system" role and the "user" role for each LLM. Table 5, 4, 6 show example conversations from three domains.

S.No.	System	User		
<b>Problem Description:</b> My car continues to overheat. The hood is so hot that I think I could melt cheese on the hood! Please help.				
1	Is the issue related to the vehicle's braking system?	No.		
2	Is the issue related to a power or electrical problem?	No.		
3	Is the issue related to overheating or temperature con- trol?	Yes.		
4	Is the issue related to a mechanical part?	Yes.		
5	Is there any visible fluid leaking from the vehicle?	Yes.		
6	Are you experiencing a radiator leak?	You are right. It's a <b>Radiator leak</b> .		

Table 4: An example conversation from FloDial dataset using MISQ-HF with GPT-40.

S.No.	System	User

#### **Problem Description:**

Hello doctor, I have been in this situation for several years. Every summer, I often have diarrhea, but it has gotten worse in the past year or two. There is no regular pattern, sometimes it happens every two to three days when it's mild, and two to three times a day when it's severe. ... I went to a major hospital for examination before, and they found no parasite eggs or fungi, just a slight imbalance in the bacterial colonies, but the medication didn't work. I don't know what my condition is, I wonder if any doctor can provide some answers. Thank you. (Male, 28 years old).

1	Do you experience any symptoms related to the respiratory sys- tem, such as coughing or difficulty breath- ing?	No, I don't experience any symptoms related to the respiratory sys- tem, such as coughing or difficulty breathing. My main issue is with my stomach and bowel movements. I have been experiencing diarrhea, stomach aches, and urgent bowel movements, especially during the summer. The stool is often loose, sticky, and not well- formed, and it's difficult to flush. I've been having this issue for several years, and it's been getting worse over the past year or two.
2	Do you experience any abdominal pain or discomfort?	Yes, I do experience abdominal pain and discomfort. It usually starts with my stomach aching more and more, and then I have to go to the bathroom. The pain can be quite severe, especially when the symp- toms are bad. Sometimes, the stool is not very loose, but it's accom- panied by mucus and is not well-formed, which makes it difficult to flush. The abdominal pain and discomfort are quite frequent, and it's been a persistent issue for me.
3	Do you experience any symptoms related to digestion, such as nausea or vomiting?	I don't usually experience nausea or vomiting, but I do have diar- rhea, which can be quite loose and sticky at times. The stool is often not well-formed and can be difficult to flush. The frequency of the diarrhea can vary, sometimes it's every two to three days when it's mild, and sometimes it's two to three times a day when it's se- vere. The symptoms can be quite unpredictable and irregular, which makes it challenging to manage.
4	Are you experiencing symptoms of irritable bowel syndrome?	You are right. I have <b>Irritable bowel syndrome</b> .

Table 5: An example conversation from MedDG dataset using MISQ-HF with GPT-40.

S.No.	System	User
1	Is X a living organism?	Yes.
2	Is X an animal that primarily lives in water?	No.
3	Is X a type of food?	No.
4	Is X a mammal?	No.
5	Is X capable of flight?	No.
6	Is X a reptile?	Yes.
7	Is X capable of changing its skin color for camouflage?	Yes.
8	Is X a 'Chameleon'?	You guessed it. X is Chameleon.

Table 6: An example conversation from Common dataset using MISQ with GPT-40.