

---

# Reasoning with Language Model is Planning with World Model

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Large language models (LLMs) have shown remarkable reasoning capabilities,  
2 particularly with chain-of-thought (CoT) prompting. However, LLMs can still  
3 struggle with problems that are easy for humans, such as generating action plans  
4 for executing tasks in a given environment, or performing complex math or logical  
5 reasoning. The deficiency stems from the key fact that LLMs lack an internal *world*  
6 *model* to predict the world *state* (e.g., environment status, intermediate variable  
7 values) and simulate long-term outcomes of actions. This prevents LLMs from  
8 performing deliberate planning akin to human brains, which involves exploring  
9 alternative reasoning paths, anticipating future states and rewards, and iteratively  
10 refining existing reasoning steps. To overcome the limitations, we propose a new  
11 LLM reasoning framework, **Reasoning via Planning (RAP)**. RAP repurposes the  
12 LLM as both a world model and a reasoning agent, and incorporates a principled  
13 planning algorithm (based on Monte Carlo Tree Search) for strategic exploration  
14 in the vast reasoning space. During reasoning, the LLM (as agent) incrementally  
15 builds a reasoning tree under the guidance of the LLM (as world model) and re-  
16 wards, and efficiently obtains a high-reward reasoning path with a proper balance  
17 between exploration *vs.* exploitation. We apply RAP to a variety of challenging rea-  
18 soning problems including plan generation, math reasoning, and logical inference.  
19 Empirical results on these tasks demonstrate the superiority of RAP over various  
20 strong baselines, including CoT and least-to-most prompting with self-consistency.  
21 RAP on LLaMA-33B surpasses CoT on GPT-4 with 33% relative improvement in  
22 a plan generation setting.

## 23 1 Introduction

24 Large language models (LLMs) have exhibited emergent reasoning abilities in a wide range of  
25 tasks [5, 10, 44, 2]. Recent approaches further boost their ability by prompting LLMs to generate  
26 intermediate reasoning steps (e.g., chain-of-thought, CoT [59]) or answer a series of subquestions  
27 (e.g., least-to-most prompting [66]). However, LLMs still face difficulties with tasks that humans find  
28 easy. For example, in creating action plans to move blocks to a target state, GPT-3 [5] achieves a  
29 success rate of only 1%, compared to 78% for humans [57]; these models also struggle when solving  
30 complex tasks that require multiple steps of math, logical, or commonsense reasoning [65, 22, 41, 6].

31 Humans possess an internal **world model**, a mental representation of the environment [28, 27, 15],  
32 which enables humans to simulate actions and their effects on the world’s state for deliberate **planning**  
33 during complex tasks of motor control, imagery, inference, and decision making [54, 55, 4, 49, 17, 33].  
34 For example, to make an action plan towards a goal, planning with the world model involves exploring  
35 various alternative courses of actions, assessing the likely outcomes by rolling out possible future  
36 scenarios, and iteratively refining the plan based on the assessment [25, 14, 52, 19, 48, 21]. This is  
37 in stark contrast to the current LLM reasoning, which instinctively generates a reasoning trace in

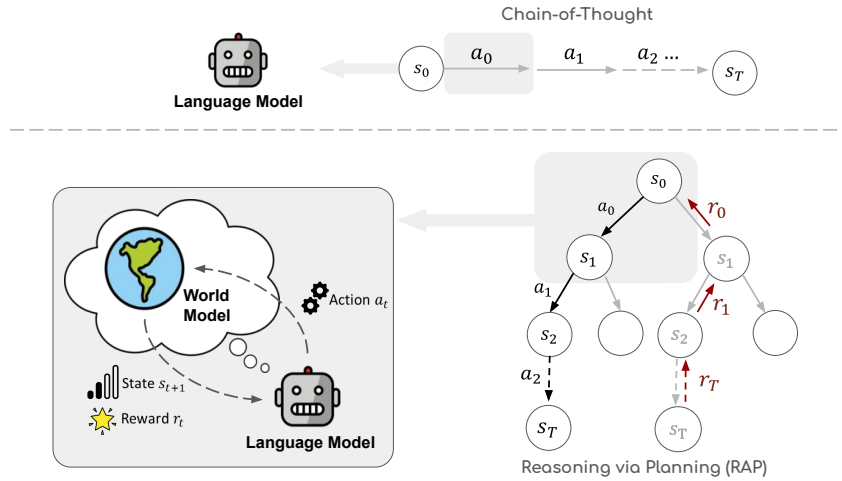


Figure 1: An overview of Reasoning via Planning (RAP). Compared with previous LLM reasoning methods like Chain-of-Thought [59], we explicitly model the world state from a world model (repurposed from the language model), enabling us to leverage advanced planning algorithms to solve the reasoning problems.

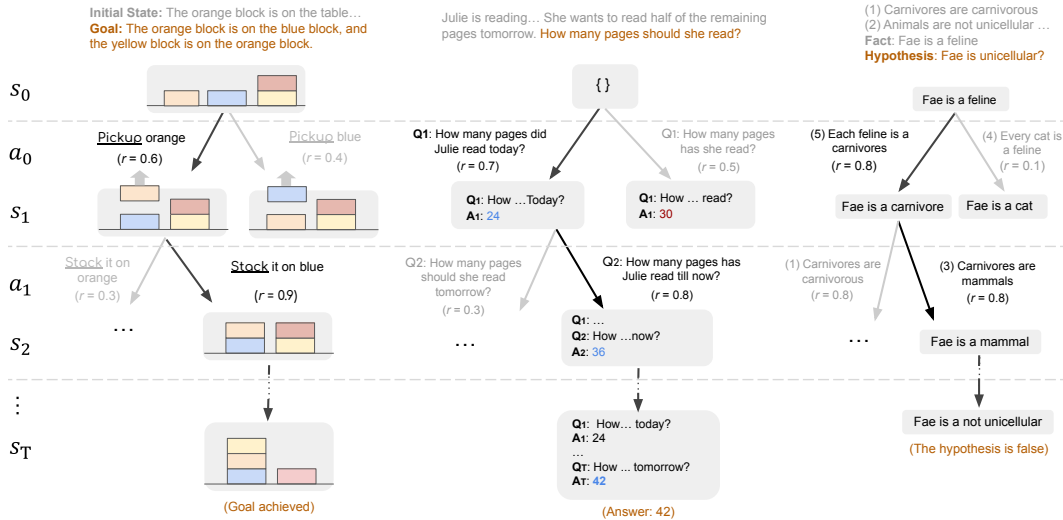


Figure 2: Examples of RAP for plan generation (left), math reasoning (middle), and logical reasoning (right).

38 an autoregressive manner. In particular, we identify several key limitations of the current reasoning  
 39 with LLMs, including (1) the lack of an internal world model to simulate the *state* of the world (e.g.,  
 40 the configuration of blocks, the values of intermediate variables), which is the foundation of human  
 41 planning; (2) the absence of a *reward* mechanism to assess and guide the reasoning towards the  
 42 desired state; and due to both of these limitations, (3) the incapability of balancing *exploration vs.*  
 43 *exploitation* to efficiently explore the vast reasoning space.

44 To address these limitations, this paper proposes a new framework, **Reasoning via Planning (RAP)**,  
 45 that enables LLMs to reason in a manner close to humans’ conscious planning. RAP augments  
 46 the LLM with a world model, and reasons with principled planning (specifically *Monte Carlo Tree*  
 47 *Search, MCTS*) to produce high-reward reasoning traces after efficient exploration (Figure 1). Notably,  
 48 we acquire the world model by repurposing the LLM itself with appropriate prompts. During the  
 49 reasoning, the LLM strategically builds a reasoning tree by iteratively considering the most promising  
 50 reasoning steps (*actions*) and using the world model (the same, repurposed LLM) to look ahead for  
 51 future outcomes. The estimated future rewards are then backpropagated to update the LLM’s beliefs  
 52 about the current reasoning steps, guiding it to refine the reasoning by exploring better alternatives.  
 53 Our MCTS-based planning effectively maintains a proper balance between exploration (of unvisited  
 54 reasoning traces) and exploitation (of the best reasoning steps identified so far).

55 We show RAP is a general framework applicable to a diverse range of challenging problems and  
56 achieves substantial improvements over recent popular LLM reasoning methods. In Blocksworld  
57 [57] for 2/4/6-step plan generation, RAP achieves an average success rate of 64% while CoT fails  
58 almost completely. Moreover, LLaMA-33B with RAP surpasses GPT-4 with CoT by 33% relative  
59 improvement. In math reasoning (GSM8K [11]) and logical inference (PrOntoQA [47]), RAP also  
60 consistently improves over CoT, least-to-most prompting, and their self-consistency variants.

## 61 2 Related Work

62 **Reasoning with LLMs.** In the realm of LLMs [22, 41, 6], reasoning typically entails decomposing  
63 complex questions into sequential intermediate steps (a.k.a. chains) before producing the final  
64 answer, exemplified by chain-of-thought (CoT) prompting and its variants [43, 59, 32]. The basic  
65 CoT approaches, which generate chains all at once, can induce additional errors as the step count  
66 increases. One line of improvement methods involves sampling multiple chains and choosing the  
67 best answer via majority voting, such as self-consistency (SC) [58]. Another line of work focuses on  
68 decomposition, aiming to tackle the problem by solving multiple simple subproblems. For instance,  
69 least-to-most prompting [66] reduces the question into subquestions and answers them sequentially.  
70 More relevantly, similar to our reward formulation, some recent works have explored self-evaluation  
71 approaches, which leverage LLMs themselves to provide feedback for intermediate steps and then  
72 continue the reasoning [60, 51, 45]. For example, Paul et al. [45] fine-tune a critic model to provide  
73 structured feedback iteratively in each step, and Madaan et al. [38] directly reuse the same LLM to  
74 generate multi-aspect feedback and refine the previously generated output. Besides, aligned with  
75 our state formulation, Li et al. [34] incorporates latent “situations” into LLMs, referring to the state  
76 of entities from the context. Nevertheless, none of the above methods formally introduce the world  
77 model and instantiates the reward and state into a unified framework.

78 **Search-guided Reasoning with LLMs.** Most of CoT approaches discussed above are based on  
79 a linear reasoning structure. Self-consistency built onto CoT decodes multiple chains parallelly,  
80 but it remains hard to explore the reasoning space sufficiently. Recent efforts have been made to  
81 investigate non-linear reasoning structures by sampling more reasoning steps efficiently guided by  
82 some search algorithms [30, 67, 63, 64]. For example, Jung et al. [30] generate a tree of explanations  
83 to enforce logical consistency, and Xie et al. [63] adopt beam search to decode a better CoT reasoning  
84 chain. More recently, CoRe [67] proposes to fine-tune both the reasoning step generator and verifier  
85 for solving math word problems, also using MCTS for reasoning decoding. Concurrently to our  
86 work, Yao et al. [64] apply heuristic-based approach, like depth-/breadth-first search, to search for  
87 better reasoning paths. Compared with these search-guided methods, RAP is a more principled  
88 framework that combines world model and reward with MCTS planning. The RAP formulation of  
89 LLM reasoning with state, action, and reward also presents a more general approach applicable to a  
90 wide range of reasoning problems.

91 **Planning with LLMs.** Planning, a central ability in intelligent agents, involves generating a series  
92 of actions to achieve a specific goal [40, 7]. Classical planning methods have been widely adopted  
93 in robots and embodied environments [9, 42, 8, 61, 26]. Recently, prompting LLMs to do planning  
94 directly has gained attention and shown potential [24, 23, 53, 13, 35]. SayCan [1], for instance,  
95 combines LLMs with affordance functions to generate feasible plans. Moreover, based on LLMs’  
96 powerful programming ability [37, 29, 36], some recent works first translate natural language  
97 instructions into the executable programming languages, such as Planning Domain Description  
98 Language (PDDL), and runs classical planning algorithms, such as LLM+P [36]. However, code-  
99 based planning is constrained by its narrow domains and the predefined environment, while RAP can  
100 handle open domain problems, including numerical and logical reasoning (see Section 4.2 and 4.3).  
101 More related works on planning and world model are discussed in Appendix A.

## 102 3 Reasoning via Planning (RAP)

103 In this section, we present the Reasoning via Planning (RAP) framework that enables LLMs to  
104 strategically plan a coherent reasoning trace for solving a wide range of reasoning tasks. We first  
105 build the world model by repurposing the LLM with prompting (Section 3.1). The world model  
106 serves as the foundation for deliberate planning, by allowing the LLM to plan ahead and seek out  
107 the expected outcomes in the future. We then introduce the rewards for assessing each state during  
108 reasoning in Section 3.2. Guided by the world model and rewards, the planning with Monte Carlo  
109 Tree Search (MCTS) efficiently explores the vast reasoning space and finds optimal reasoning traces

110 (Section 3.3). Finally, when multiple promising reasoning traces are acquired during planning, we  
111 further introduce an aggregation method in Section 3.4 that yields an integrated result and further  
112 boosts the reasoning performance.

### 113 3.1 Language Model as World Model

114 In general, a world model predicts the next *state* of the reasoning after applying an *action* to the  
115 current state [17, 39]. RAP enables us to instantiate the general concepts of state and action in  
116 different ways depending on the specific reasoning problems at hand. For example, in Blocksworld  
117 (Figure 2 left), it is natural to set a state to describe a configuration of blocks (with natural language),  
118 and an action to be a behavior of moving a block (e.g., “pickup the orange block”). In a math  
119 reasoning problem (Figure 2 middle), we use the state to represent the values of intermediate variables,  
120 and set an action to be a subquestion that drives the reasoning to derive new values (i.e., new state).

121 After the definition of state and action, the reasoning process can thus be described as a Markov  
122 decision process (MDP): given the current state  $s_{t,t=0,1,\dots,T}$ , e.g., the initial state  $s_0$ , the LLM (as a  
123 reasoning agent) samples several actions as the action space, following its generative distribution  
124  $a_t \sim p(a|s_t, c)$ , where  $c$  is a proper prompt (e.g., in-context demonstrations) to steer the LLM for  
125 action generation. The world model then predicts the next state  $s_{t+1}$  of the reasoning. Specifically,  
126 we repurpose the *same* LLM to obtain a state transition distribution  $p(s_{t+1}|s_t, a_t, c')$ , where  $c'$  is  
127 another prompt to guide the LLM to generate a state. For instance, in Blocksworld, the LLM (as the  
128 world model) generates text  $s_{t+1}$  to describe the new configuration of blocks, given the previous state  
129 description  $s_t$  and the action  $a_t$ .

130 Continuing the process results in a reasoning trace, which consists of a sequence of interleaved states  
131 and actions  $(s_0, a_0, s_1, \dots, a_{T-1}, s_T)$ . This differs from the previous reasoning methods, such as  
132 Chain-of-Thought [59], where the intermediate reasoning steps consist of only a sequence of actions,  
133 e.g.,  $(a_0 = \text{“pickup red block”}, a_1 = \text{“stack on yellow block”}, \dots)$  (see comparisons  
134 in Figure 1). Augmenting the reasoning with the (predicted) world states helps the LLM with a  
135 more grounded and coherent inference. Note that the full reasoning trace is simulated by the LLM  
136 itself (as a reasoning agent with an *internal* world model) without interacting with the *external* real  
137 environment. This resembles humans contemplating a possible plan in their minds. The capability of  
138 simulating future states, due to the introduction of the world model, allows us to incorporate principled  
139 planning algorithms to efficiently explore the vast reasoning space as described in Section 3.3.

### 140 3.2 Reward Design

141 During reasoning, we want to assess the feasibility and desirability of each reasoning step, and  
142 guide the reasoning based on the assessment (Section 3.3). The assessment of each reasoning step  
143 (i.e., applying an action  $a_t$  to the state  $s_t$ ) is performed by a *reward* function  $r_t = r(s_t, a_t) \in \mathbb{R}$ .  
144 Similar to the state and action, the reward function can be specified in different ways to accommodate  
145 any knowledge or preferences about the reasoning problem of interest. Here we introduce several  
146 common rewards applicable to different tasks and shown to be effective in our experiments.

147 **Likelihood of the action.** When an action is generated by the LLM conditioning on the in-context  
148 demonstration and the current state, the probability of the specific action reflects the LLM’s preference.  
149 We thus can incorporate the log probability of the action as a reward. This reward reflects the “instinct”  
150 of LLMs as an agent, and can be also used as a prior for which action to explore.

151 **Confidence of the state transition.** State prediction is nontrivial in some problems, e.g., in math  
152 reasoning (Figure 2, middle), given an action (i.e., a subquestion), the world model updates the set of  
153 known variables by answering the subquestion. Since LLMs may make mistakes when answering  
154 these questions, We incorporate the confidence of the state transition (i.e., the answer to a subquestion  
155 in this case) as a reward. Specifically, we sample multiple answers from the language model, and use  
156 the proportion of the most frequent answer as the confidence. A high confidence indicates a reliable  
157 reasoning step, which is worth more exploration in the future.

158 **Self-evaluation by the LLM.** It’s sometimes easier to recognize the errors in reasoning than avoid  
159 generating them in advance. Thus, it’s beneficial to allow the LLM to criticize itself with the question  
160 “Is this reasoning step correct?”, and use the next-word probability of the token “Yes” as  
161 a reward. The reward evaluates LLM’s own estimation of the correctness of reasoning. Note that the  
162 specific problems for self-evaluation can be different depending on the tasks.

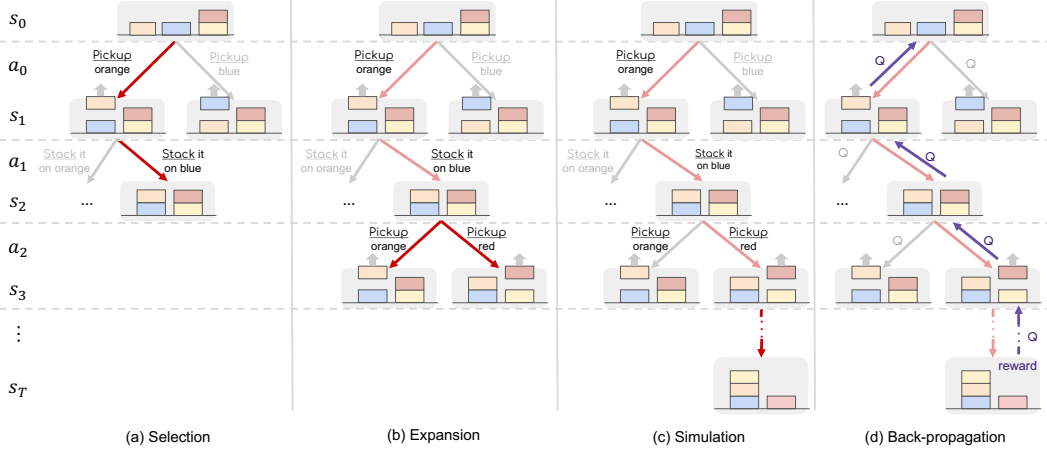


Figure 3: An illustration of the four phases in an iteration in MCTS planning (Section 3.3).

163 **Task-specific heuristics.** We can also flexibly plug-in other diverse task-specific heuristics into the  
 164 reward function. For example, in plan generation for Blocksworld, we compare the predicted current  
 165 state of blocks with the goal to calculate a reward (Section 4.1). The reward encourages the plan of  
 166 movements to actively pace towards the target.

### 167 3.3 Planning with Monte Carlo Tree Search

168 The world model (Section 3.1) and rewards (Section 3.2) enable LLMs to reason with advanced  
 169 planning algorithms, where we adopt Monte Carlo Tree Search (MCTS) [31, 12], a powerful planning  
 170 algorithm that strategically explores the space of reasoning trees, and strikes a proper balance between  
 171 exploration and exploitation to find a good reasoning trace efficiently.

172 MCTS builds a reasoning tree iteratively, where each node represents a state, and each edge represents  
 173 an action and the transition from the current state to the next state after applying the action (Figure 1).  
 174 To guide the LLM agent to expand and explore the most promising nodes of the tree, the algorithm  
 175 maintains a state-action value function  $Q : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ , where  $Q(s, a)$  estimates the *expected*  
 176 *future reward* of taking action  $a$  in state  $s$ . That is, we assess the potential of a node (or a reasoning  
 177 step) by looking ahead and anticipating the reward in future trajectories starting from this node.  
 178 This fundamentally differs from the current reasoning methods that generate a reasoning trace  
 179 autoregressively from left to right without accounting for the future.

180 Specifically, as illustrated in Figure 3, the MCTS planning performs four operations in each iteration  
 181 to expand the tree and update  $Q$  values, i.e., *selection*, *expansion*, *simulation*, and *back-propagation*.  
 182 The process continues until a specified computational budget (e.g., the number of iterations) is  
 183 reached, and the resulting reasoning traces are acquired from the tree, as we articulated later. The  
 184 pseudo-code and more implementation details are presented in Algorithm 1 and Appendix B.

185 **Selection.** The first phase selects a portion of the existing tree that is most promising for further  
 186 expansion in the next phase. Specifically, starting from the root node (i.e., initial state  $s_0$ ), at each  
 187 level of the tree, the algorithm selects a child node as the next node. The phase finishes when a  
 188 leaf node of the current tree is reached. Figure 3(a) highlights the selected path in red. To balance  
 189 between exploration (of less-visited nodes) and exploitation (of high-value nodes), we use the well-  
 190 known *Upper Confidence bounds applied to Trees (UCT)* algorithm [31] to select each child node.  
 191 Specifically, at node  $s$ , we select the action (which leads to a transition to a child node) in the tree by  
 192 considering both the  $Q$  value (for exploitation) and uncertainty (for exploration):

$$a^* = \arg \max_{a \in A(s)} \left[ Q(s, a) + w \sqrt{\frac{\ln N(s)}{N(c(s, a))}} \right], \quad (1)$$

193 where  $N(s)$  is the number of times node  $s$  has been visited in previous iterations, and  $c(s, a)$  is the  
 194 child node of applying  $a$  in state  $s$ . Therefore, the less a child node was visited before (i.e., the more  
 195 uncertain about this child node), the higher the second term in the equation. The weight  $w$  controls  
 196 the balance between exploration and exploitation.

197 **Expansion.** This phase expands the tree by adding new child nodes to the leaf node selected above.  
198 Specifically, given the state of the leaf node, we use the LLM (as agent) to sample  $d$  possible  
199 actions (e.g., subquestions in math reasoning), and then use the LLM (as world model) to predict the  
200 respective next state, resulting in  $d$  child nodes. Note that if the leaf node selected above is a terminal  
201 node (the end of a reasoning chain) already, we will skip expansion and jump to back-propagation.

202 **Simulation.** This phase simulates the reasoning chain to the end in order to estimate the expected  
203 future rewards ( $Q$  values). Specifically, starting from the current node  $s$ , we iteratively select an  
204 action following a *roll-out policy* and use the world model to predict the next state. The roll-out  
205 process continues until a terminal state is reached. While the design of the roll-out policy is flexible,  
206 in our experiments, we generate  $d$  candidate actions and pick the one with the largest local reward  
207  $a' = \max_{a'} r(s, a)$ . In practice, as the roll-out process will evaluate the reward function for multiple  
208 nodes, for efficiency, we discard the computationally expensive components in  $r$  (for example, the  
209 reward of the state transition confidence requires sampling the answer multiple times), and use the  
210 resulting light-weight reward function for selecting actions during simulation.

211 **Back-propagation.** Once we reach a terminal state in the above phases, we obtain a reasoning path  
212 from the root node to the terminal node. We now back-propagate the rewards on the path to update  
213 the  $Q$  value of each state-action pair along the path. That is,  $Q(s, a)$  is updated by aggregating the  
214 rewards in all future steps of node  $s$ . We may adopt the aggregation method according to the nature  
215 of different tasks and reward design, as discussed in Section 4.

216 As mentioned earlier, once a predetermined number of MCTS iterations is reached, we terminate  
217 the algorithm and select final reasoning trace from the constructed tree. There could be various  
218 ways for the selection. One approach is to start from the root node and iteratively choose the action  
219 with the highest  $Q$  value until reaching a terminal. Alternatively, one can directly select the path  
220 from the iterations that yielded the highest reward, or opt to choose the leaf node (and the respective  
221 root-to-leaf path) that has been visited the most. In practice, we observed that the second strategy  
222 often yields the best results.

### 223 3.4 RAP-Aggregation: Aggregating Multiple Reasoning Outputs

224 Ensemble-based methods, such as self-consistency CoT [58], can effectively improve performance  
225 by aggregating multiple valid reasoning traces. Therefore, for problems, such as math reasoning  
226 (Section 4.2) where only the final answer is required, RAP could produce multiple traces and answers  
227 from different MCTS iterations, which will be aggregated to produce the final answer. We refer to  
228 such a mechanism as RAP-Aggregation. Note that problems like plan generation or logical inference  
229 require a complete reasoning trace as output; thus, RAP-Aggregation will not be applied.

230 More importantly, there is a concern that some incorrect reasoning steps may appear in the early stage  
231 of multiple iterations, thus polluting the aggregation. As a result, we further devise a new weighting  
232 strategy for aggregating candidate answers. Specifically, for each candidate answer, we accumulate  
233 the reward of each reasoning step in the answer’s reasoning traces. We choose the answer with the  
234 highest accumulative reward as the final aggregated answer.

## 235 4 Experiments

236 In this section, we demonstrate the flexibility and effectiveness of our RAP framework by applying it to  
237 a wide range of problems, including plan generation in an embodied environment (4.1), mathematical  
238 reasoning for solving math word problems (4.2), and logical reasoning for verifying hypotheses (4.3).  
239 The subsequent sections demonstrate how the world model formulation in RAP enables a versatile  
240 design of the state and action, catering to various reasoning contexts. We also discuss the choice of  
241 reward in Appendix C.

242 We primarily compare RAP with chain-of-thought (CoT) [59], and its variants like least-to-most  
243 prompting [66] as baselines. We also consider ensembling multiple reasoning paths if applicable (also  
244 known as self-consistency [58]). Moreover, we compare RAP with GPT-4 [44] when computation  
245 resources allow. By default, we use the LLaMA-33B model [56] as the base LLM for both our  
246 methods and baselines, with a sampling temperature of 0.8. All prompts are shown in Appendix D.

### 247 4.1 Plan Generation

248 The plan generation task aims to produce a sequence of actions to achieve a given goal, possibly with  
249 additional constraints. The ability to generate plans is important for intelligent embodied agents,

Table 1: Results on Blocksworld.  $\text{RAP}^{(10)}$  and  $\text{RAP}^{(20)}$  refer to our method where the iteration number is set to 10 and 20, respectively. “pass@10” is a relaxed metric, where 10 plans are sampled for each test case, and the test case regarded as solved if at least one plan is successful. For all other settings including RAP, only a single plan is evaluated.

Method	2-step	4-step	6-step
CoT	0.17	0.02	0.00
CoT - pass@10	0.23	0.07	0.00
CoT (GPT-4)	0.50	0.63	0.40
$\text{RAP}^{(10)}$	1.00	0.86	0.26
$\text{RAP}^{(20)}$	<b>1.00</b>	<b>0.88</b>	<b>0.42</b>

e.g. household robots [46]. This task has also been widely used to evaluate the reasoning ability of LLMs given their challenging requirements of long-horizon reasoning, e.g., Blocksworld is a classic problem, where an agent is asked to rearrange the blocks into stacks in a particular order.

**Task setup.** To explore the viability of the RAP framework for plan generation tasks, we adapt and evaluate RAP on the Blocksworld benchmark [50]. We define a **state** as the current orientation of the blocks and an **action** as an instruction that moves blocks. Specifically, an action is composed of one of the 4 verbs (i.e., STACK, UNSTACK, PUT, and PICKUP) and manipulated objects. For the action space, we generate the currently valid actions given the domain restrictions on actions and the current orientation of the blocks. To transit between states, we take the current action and query the LLM to predict the state changes to the relevant blocks. We then update the current state by adding the new block conditions and removing the conditions that are no longer true. Once a state has met all of the conditions listed in the goal or the depth limit of the tree is reached, we terminate the associated node.

To assess the quality of actions within this domain, we use two separate **rewards**. First, we prompt the LLM with some example test cases along with their solutions, and then calculate the log probability of the action given the current state (“*Likelihood of action*” reward in Section 3.2), denoted as  $r_1$ . This reward reflects the intuition of the LLM as the reasoning agent. It’s typically indicative when there are few steps left to the goal, while not as reliable for a distant goal. Additionally, we compare the new state after performing an action with the goal and provide a reward,  $r_2$ , scaling with the number of conditions met (“*Task-specific heuristics*” reward). Specifically, when all the conditions are met, we assign a super large reward to make sure this plan will be selected as the solution.

**Results.** We use test cases from the Blocksworld dataset [57] and group them by solvable steps, resulting in 30 cases solvable with 2 steps, 57 cases with 4 steps, and 114 cases with 6 steps. There are at most 5 blocks in each test case. As the baseline method, we prompt the LLM with 4 test cases with corresponding solutions, and ask it to generate a plan for a new question. This setting is the same as one described in Valmeekam et al. [57], and we denote it as Chain-of-Thought (CoT) for brevity. For RAP, the same prompt is shown to help LLMs calculate  $r_1$ .

As shown in Table 1, CoT with LLaMA-33B can only generate successful plans for a few 2-step cases, and completely fails on harder problems. RAP substantially improves over CoT by nearly solving all problems within 4 steps, and a part of 6-step problems, achieving an average success rate of 64%. It’s worth noting that the searching space of 6-step problems can be as large as  $5^6$ , while our algorithm can find a successful plan 42% of the time within 20 iterations. Even more, our framework allows LLaMA-33B to outperform GPT-4 by 33% relative improvement [44], which is known to have much stronger reasoning ability [6].

We further present a case study of comparing the reasoning paths from Cot and RAP. As illustrated in Figure 4, we find the improvement can be mainly attributed to the following reasons: (1) By maintaining the world state during reasoning, RAP can recognize valid actions for the current state, avoiding generating illegal plans. (2) RAP is capable of backtracking and trying out other solutions when the first intuition from the LLM doesn’t work. Specifically, CoT attempts to achieve the second goal, i.e. “orange on red”, and achieve that with the first two steps. However, accomplishing the second goal first would prevent the first goal from being satisfied. On the contrary, even though RAP makes the same mistakes in the first iterations, our framework drives the agent to explore other possible paths (as described in Section 3.3) and finally generate a successful plan. (3) When calculating  $r_t$ , we can feed only the current state to the LLM and hide the history. E.g., in the case

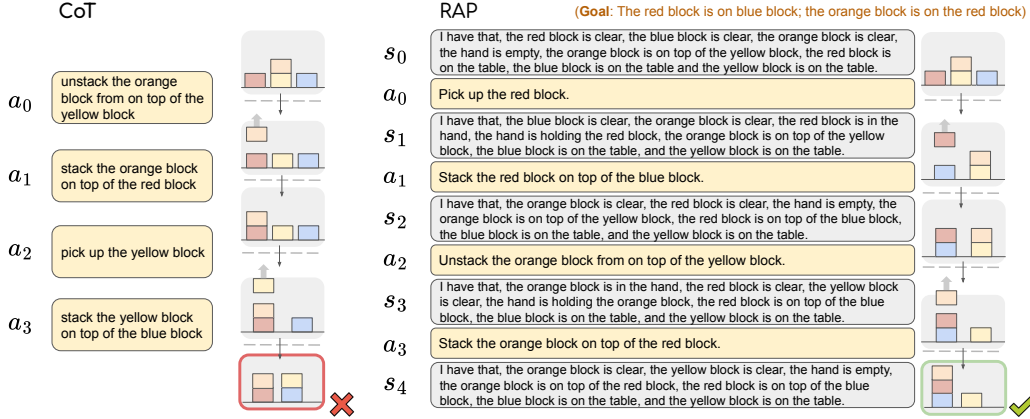


Figure 4: Comparing reasoning traces in Blocksworld from CoT (left) and RAP (right).

of Figure 4, to calculate the reward for  $a_2$ , the LLM is provided with a “new” test case, in which  $s_2$  is the initial state. This significantly lowers the difficulties of the last few steps, and saves more iterations for harder decisions of the first few steps.

## 4.2 Math Reasoning

**Task setup.** Numerical reasoning tasks, such as GSM8k [11], often include a description and a final question. To arrive at the answer to the final question, it is necessary to undertake multi-step mathematical calculations based on the problem’s context. It is thus natural to decompose the final question into a sequence of smaller sub-questions (Figure 2, right). To adapt RAP, we define a **state** as the values of intermediate variables, while an **action** is to propose an incremental sub-question about a new intermediate variable. The world model then responds to the sub-question using the intermediate variables and the problem description, adding the new intermediate variable value into the next state. We combine the self-evaluation by LLM  $r_{t,1}$  and the confidence of state transition  $r_{t,2}$  using weighted geometric mean  $r_t = r_{t,1}^\alpha * r_{t,2}^{1-\alpha}$  as the **reward** function. This reward encourages more relevant and useful sub-questions. To account for the impact of the reasoning path’s length on the reward, we compute the **Q value** by using the maximum of average rewards in future steps.

$$Q^*(s_t, a_t) = \max_{s_t, a_t, r_t, \dots, s_l, a_l, r_l, s_{l+1}} \text{avg}(r_t, \dots, r_l). \quad (2)$$

As a related work, Least-to-Most prompting [66] shares a similar idea to us in sub-question decomposition, but they generate sub-questions all at once. On the contrary, RAP considers each action  $a_t$  based on the current state  $s_t$ , which enables more informed decisions.

**Results.** We evaluate our framework on GSM8k, a dataset of grade high school math problems. We also evaluate the base model with CoT prompting [59], Least-to-Most prompting [66], and their self-consistency [58] variants, as the baselines. We use the same 4-shot examples demonstrations for both our framework and the baselines.

As shown in Table 2, our RAP framework answers 48.8% of the problems correctly, outperforming both the Chain-of-Thought and the Least-to-Most prompting with self-consistency. All methods use the same set of examples as the in-context demonstration. Notably, this result is achieved when RAP selects only one reasoning trace based on the reward. The introduction of RAP-Aggregate further improves the accuracy by  $\sim 3\%$ . We also calculate the accuracy with different numbers of iterations in MCTS and self-consistency samples in baselines, as illustrated in Figure 2. We find that across all numbers of iterations/samples, RAP-Aggregation outperforms baselines consistently, which indicates that when only a few iterations/samples are allowed, our framework is significantly better at finding reliable reasoning paths with the guide of reward.

## 4.3 Logical Reasoning

**Task setup.** A logical reasoning task (e.g. PrOntoQA [47]) typically provides a set of *facts* and *logical rules*, and a model is required to verify if a *hypothesis fact* is true or false by applying the logical rules to the given facts, as illustrated in Figure 2. These tasks not only require the correct final answer (true/false), but also a detailed proof demonstrating the result. To apply our framework,



Table 2: Results on GSM8k. The superscripts indicate the number of samples or iterations.

Method	Accuracy (%)
Chain-of-Thought	29.4
+ SC <sup>(10)</sup>	46.8
Least-to-Most	25.5
+ SC <sup>(10)</sup>	42.5
RAP <sup>(1)</sup>	40.0
RAP <sup>(10)</sup>	48.6
+ aggr	<b>51.6</b>




Figure 5: The performance of RAP and baselines on GSM-8K, with different numbers of sampled paths or iterations.

329 we define the **state** as a fact we are focusing on, analogous to the human’s working memory [3] for  
 330 inference. An **action** is defined as selecting a rule from the fact set. The world model performs a one-  
 331 hop reasoning step to get a new fact as the next state. The **reward** is calculated with Self-evaluation  
 332 (Section 3.2. Specifically, we prompt the LLM with a few examples with their labels to help it better  
 333 understand the quality of reasoning steps. We use the average reward of future steps to update the **Q**  
 334 **function**, the same as Equation (2) for GSM8k.

335 **Results.** We assess the performance of our RAP  
 336 framework on PrOntoQA [47]. We adopt their settings  
 337 of “true” ontology (using real-world knowledge), “ran-  
 338 dom” ordering of rules. We mix the examples requiring  
 339 3, 4, and 5 reasoning hops in a correct proof to prevent  
 340 LLM from memorizing when to finish the reasoning.  
 341 We sample 500 examples from the generation script  
 342 released by Saparov and He [47]. We compare both the  
 343 prediction accuracy of the final answer and the accuracy  
 344 of the entire proof. We do 20 iterations for MCTS and 20 samples for self-consistency in baselines.

Table 3: Results on ProntoQA.

Method	Pred Acc	Proof Acc
CoT	87.8	64.8
CoT + SC	89.8	-
RAP (Ours)	<b>94.2</b>	<b>78.8</b>

345 As the results presented in Table 3, our framework achieves an output accuracy of 94.2% and a proof  
 346 accuracy of 78.8%, surpassing the CoT baseline by 14% proof accuracy and the self-consistency  
 347 CoT baseline by 4.4% prediction accuracy. Such substantial improvements clearly demonstrate  
 348 the effectiveness of RAP in solving logical reasoning problems in the PrOntoQA dataset. That is  
 349 because RAP can effectively recognize the error when a reasoning chain comes to a dead end, and  
 350 propagate the signal back to earlier reasoning steps, with the planning algorithm allowing it to explore  
 351 alternatives to the previous steps (Figure 2). The self-evaluation reward performs as a prior, which  
 352 prioritizes the most promising actions for exploration.

## 353 5 Conclusion

354 In this paper, we present Reasoning via Planning (RAP), a novel LLM reasoning framework that  
 355 equips LLMs with an ability to reason akin to human-like strategic planning. By coupling the LLMs’  
 356 reasoning capabilities with a world model and principled planning via Monte Carlo Tree Search, RAP  
 357 bridges the gap between LLMs and human planning capabilities. Our framework, which repurposes  
 358 the LLM to act as both a world model and a reasoning agent, enables the LLM to simulate states of the  
 359 world and anticipate action outcomes, while achieving an effective balance between exploration and  
 360 exploitation in the vast reasoning space. Extensive experiments on a variety of challenging reasoning  
 361 problems demonstrate RAP’s superiority over several contemporary CoT-based reasoning approaches,  
 362 and even the advanced GPT-4 in certain settings. RAP’s flexibility in formulating rewards, states, and  
 363 actions further proves its potential as a general framework for solving diverse reasoning tasks. We  
 364 posit that RAP, with its innovative melding of planning and reasoning, has the potential to redefine the  
 365 way we approach LLM reasoning - essentially forging a new pathway toward achieving human-level  
 366 strategic thinking and planning in artificial intelligence.

367 **References**

- 368 [1] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David,  
369 Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, et al. Do as i can, not  
370 as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- 371 [2] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos,  
372 Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report.  
373 *arXiv preprint arXiv:2305.10403*, 2023.
- 374 [3] Alan Baddeley. Working memory. *Science*, 255(5044):556–559, 1992.
- 375 [4] Robert Eamon Briscoe. Mental imagery and the varieties of amodal perception. *Pacific*  
376 *Philosophical Quarterly*, 92(2):153–173, 2011.
- 377 [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
378 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are  
379 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- 380 [6] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece  
381 Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general  
382 intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.
- 383 [7] Tom Bylander. The computational complexity of propositional strips planning. *Artificial*  
384 *Intelligence*, 69(1-2):165–204, 1994.
- 385 [8] Eduardo F Camacho and Carlos Bordons Alba. *Model predictive control*. Springer science &  
386 business media, 2013.
- 387 [9] Jaime Carbonell, Oren Etzioni, Yolanda Gil, Robert Joseph, Craig Knoblock, Steve Minton, and  
388 Manuela Veloso. Prodigy: An integrated architecture for planning and learning. *ACM SIGART*  
389 *Bulletin*, 2(4):51–55, 1991.
- 390 [10] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam  
391 Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm:  
392 Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- 393 [11] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
394 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to  
395 solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- 396 [12] Rémi Coulom. Efficient selectivity and backup operators in monte-carlo tree search. In  
397 *Computers and Games: 5th International Conference, CG 2006, Turin, Italy, May 29-31, 2006.*  
398 *Revised Papers 5*, pages 72–83. Springer, 2007.
- 399 [13] Yan Ding, Xiaohan Zhang, Chris Paxton, and Shiqi Zhang. Task and motion planning with  
400 large language models for object rearrangement. *arXiv preprint arXiv:2303.06247*, 2023.
- 401 [14] Wojciech W Gasparski and Tufan Orel. *Designology: Studies on Planning for Action*, volume 1.  
402 Transaction Publishers, 2014.
- 403 [15] Dedre Gentner and Albert L Stevens. *Mental models*. Psychology Press, 2014.
- 404 [16] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. *Ad-*  
405 *vances in neural information processing systems*, 31, 2018.
- 406 [17] David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2018.
- 407 [18] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control:  
408 Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.
- 409 [19] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and  
410 James Davidson. Learning latent dynamics for planning from pixels. In *International conference*  
411 *on machine learning*, pages 2555–2565. PMLR, 2019.

- 412 [20] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with  
413 discrete world models. *arXiv preprint arXiv:2010.02193*, 2020.
- 414 [21] Mark K Ho, David Abel, Carlos G Correa, Michael L Littman, Jonathan D Cohen, and Thomas L  
415 Griffiths. Control of mental representations in human planning. *arXiv e-prints*, pages arXiv-  
416 2105, 2021.
- 417 [22] Jie Huang and Kevin Chen-Chuan Chang. Towards reasoning in large language models: A  
418 survey. *arXiv preprint arXiv:2212.10403*, 2022.
- 419 [23] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as  
420 zero-shot planners: Extracting actionable knowledge for embodied agents. In *International  
421 Conference on Machine Learning*, pages 9118–9147. PMLR, 2022.
- 422 [24] Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng,  
423 Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, et al. Inner monologue: Embodied  
424 reasoning through planning with language models. *arXiv preprint arXiv:2207.05608*, 2022.
- 425 [25] Quentin JM Huys, Neir Eshel, Elizabeth O’Nions, Luke Sheridan, Peter Dayan, and Jonathan P  
426 Roiser. Bonsai trees in your head: how the pavlovian system sculpts goal-directed choices by  
427 pruning decision trees. *PLoS computational biology*, 8(3):e1002410, 2012.
- 428 [26] Yu-qian Jiang, Shi-qi Zhang, Piyush Khandelwal, and Peter Stone. Task planning in robotics:  
429 an empirical comparison of pddl-and asp-based systems. *Frontiers of Information Technology  
430 & Electronic Engineering*, 20:363–373, 2019.
- 431 [27] Philip N Johnson-Laird. Mental models and human reasoning. *Proceedings of the National  
432 Academy of Sciences*, 107(43):18243–18250, 2010.
- 433 [28] Philip Nicholas Johnson-Laird. *Mental models: Towards a cognitive science of language,  
434 inference, and consciousness*. Number 6. Harvard University Press, 1983.
- 435 [29] Ana Jojic, Zhen Wang, and Nebojsa Jojic. Gpt is becoming a turing machine: Here are some  
436 ways to program it. *arXiv preprint arXiv:2303.14310*, 2023.
- 437 [30] Jaehun Jung, Lianhui Qin, Sean Welleck, Faeze Brahman, Chandra Bhagavatula, Ronan Le  
438 Bras, and Yejin Choi. Maieutic prompting: Logically consistent reasoning with recursive  
439 explanations. *arXiv preprint arXiv:2205.11822*, 2022.
- 440 [31] Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *Machine Learning:  
441 ECML 2006: 17th European Conference on Machine Learning Berlin, Germany, September  
442 18-22, 2006 Proceedings 17*, pages 282–293. Springer, 2006.
- 443 [32] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large  
444 language models are zero-shot reasoners. *arXiv preprint arXiv:2205.11916*, 2022.
- 445 [33] Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27.  
446 *Open Review*, 62, 2022.
- 447 [34] Belinda Z Li, Maxwell Nye, and Jacob Andreas. Language modeling with latent situations.  
448 *arXiv preprint arXiv:2212.10012*, 2022.
- 449 [35] Kevin Lin, Christopher Agia, Toki Migimatsu, Marco Pavone, and Jeannette Bohg. Text2motion:  
450 From natural language instructions to feasible plans. *arXiv preprint arXiv:2303.12153*, 2023.
- 451 [36] Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and Peter  
452 Stone. Llm+ p: Empowering large language models with optimal planning proficiency. *arXiv  
453 preprint arXiv:2304.11477*, 2023.
- 454 [37] Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apid-  
455 ianaki, and Chris Callison-Burch. Faithful chain-of-thought reasoning. *arXiv preprint  
456 arXiv:2301.13379*, 2023.

- 457 [38] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri  
458 Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, et al. Self-refine: Iterative refinement  
459 with self-feedback. *arXiv preprint arXiv:2303.17651*, 2023.
- 460 [39] Yutaka Matsuo, Yann LeCun, Maneesh Sahani, Doina Precup, David Silver, Masashi Sugiyama,  
461 Eiji Uchibe, and Jun Morimoto. Deep learning, reinforcement learning, and world models.  
462 *Neural Networks*, 2022.
- 463 [40] John McCarthy. Situations, actions, and causal laws. Technical report, STANFORD UNIV CA  
464 DEPT OF COMPUTER SCIENCE, 1963.
- 465 [41] Grégoire Mialon, Roberto Dessì, Maria Lomeli, Christoforos Nalmpantis, Ram Pasunuru,  
466 Roberta Raileanu, Baptiste Rozière, Timo Schick, Jane Dwivedi-Yu, Asli Celikyilmaz, et al.  
467 Augmented language models: a survey. *arXiv preprint arXiv:2302.07842*, 2023.
- 468 [42] Dana S Nau, Tsz-Chiu Au, Okhtay Ilghami, Ugur Kuter, J William Murdock, Dan Wu, and  
469 Fusun Yaman. Shop2: An htn planning system. *Journal of artificial intelligence research*, 20:  
470 379–404, 2003.
- 471 [43] Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin,  
472 David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. Show  
473 your work: Scratchpads for intermediate computation with language models. *arXiv preprint*  
474 *arXiv:2112.00114*, 2021.
- 475 [44] OpenAI. Gpt-4 technical report, 2023.
- 476 [45] Debjit Paul, Mete Ismayilzada, Maxime Peyrard, Beatriz Borges, Antoine Bosselut, Robert  
477 West, and Boi Faltings. Refiner: Reasoning feedback on intermediate representations. *arXiv*  
478 *preprint arXiv:2304.01904*, 2023.
- 479 [46] Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio  
480 Torralba. Virtualhome: Simulating household activities via programs, 2018.
- 481 [47] Abulhair Saparov and He He. Language models are greedy reasoners: A systematic formal  
482 analysis of chain-of-thought. *arXiv preprint arXiv:2210.01240*, 2022.
- 483 [48] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Si-  
484 mon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering  
485 atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
- 486 [49] Jay Schulkin. *Action, perception and the brain: Adaptation and cephalic expression*. Springer,  
487 2012.
- 488 [50] Ramanan Sekar, Oleh Rybkin, Kostas Daniilidis, Pieter Abbeel, Danijar Hafner, and Deepak  
489 Pathak. Planning to explore via self-supervised world models. In *International Conference on*  
490 *Machine Learning*, pages 8583–8592. PMLR, 2020.
- 491 [51] Noah Shinn, Beck Labash, and Ashwin Gopinath. Reflexion: an autonomous agent with  
492 dynamic memory and self-reflection. *ArXiv*, abs/2303.11366, 2023.
- 493 [52] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur  
494 Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al. Mastering  
495 chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint*  
496 *arXiv:1712.01815*, 2017.
- 497 [53] Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay,  
498 Dieter Fox, Jesse Thomason, and Animesh Garg. Progprompt: Generating situated robot task  
499 plans using large language models. *arXiv preprint arXiv:2209.11302*, 2022.
- 500 [54] Edward C Tolman. Cognitive maps in rats and men. *Psychological review*, 55(4):189, 1948.
- 501 [55] Marc Toussaint. Learning a world model and planning with a self-organizing, dynamic neural  
502 system. *Advances in neural information processing systems*, 16, 2003.

- 503 [56] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timo-  
504 thée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open  
505 and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 506 [57] Karthik Valmeekam, Sarath Sreedharan, Matthew Marquez, Alberto Olmo, and Subbarao  
507 Kambhampati. On the planning abilities of large language models (a critical investigation with  
508 a proposed benchmark). *arXiv preprint arXiv:2302.06706*, 2023.
- 509 [58] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, and Denny Zhou. Self-  
510 consistency improves chain of thought reasoning in language models. *arXiv preprint*  
511 *arXiv:2203.11171*, 2022.
- 512 [59] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny  
513 Zhou. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint*  
514 *arXiv:2201.11903*, 2022.
- 515 [60] Sean Welleck, Ximing Lu, Peter West, Faeze Brahman, Tianxiao Shen, Daniel Khashabi, and  
516 Yejin Choi. Generating sequences by learning to self-correct. *arXiv preprint arXiv:2211.00053*,  
517 2022.
- 518 [61] Grady Williams, Paul Drews, Brian Goldfain, James M Rehg, and Evangelos A Theodorou.  
519 Information-theoretic model predictive control: Theory and applications to autonomous driving.  
520 *IEEE Transactions on Robotics*, 34(6):1603–1622, 2018.
- 521 [62] Philipp Wu, Alejandro Escontrela, Danijar Hafner, Pieter Abbeel, and Ken Goldberg. Day-  
522 dreamer: World models for physical robot learning. In *Conference on Robot Learning*, pages  
523 2226–2240. PMLR, 2023.
- 524 [63] Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, Xu Zhao, Min-Yen Kan, Junxian He, and Qizhe  
525 Xie. Decomposition enhances reasoning via self-evaluation guided decoding. *arXiv preprint*  
526 *arXiv:2305.00633*, 2023.
- 527 [64] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik  
528 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. 2023.
- 529 [65] Ping Yu, Tianlu Wang, Olga Golovneva, Badr Alkhamissy, Gargi Ghosh, Mona Diab, and  
530 Asli Celikyilmaz. Alert: Adapting language models to reasoning tasks. *arXiv preprint*  
531 *arXiv:2212.08286*, 2022.
- 532 [66] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale  
533 Schuurmans, Olivier Bousquet, Quoc Le, and Ed Chi. Least-to-most prompting enables  
534 complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*, 2022.
- 535 [67] Xinyu Zhu, Junjie Wang, Lin Zhang, Yuxiang Zhang, Ruyi Gan, Jiaying Zhang, and Yujiu Yang.  
536 Solving math word problem via cooperative reasoning induced language models. *arXiv preprint*  
537 *arXiv:2210.16257*, 2022.

---

**Algorithm 1** RAP-MCTS

---

**Require:** Initial state  $s_0$ , state transition probability function  $p_\theta$ , reward function  $r_\theta$ , action generator  $p_\phi$ , number of generated actions  $d$ , depth limit  $L$ , number of roll-outs  $N$ , and exploration weight  $w$

- 1: Initialize memory of actions  $A : \mathcal{S} \mapsto \mathcal{A}$ , children  $c : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$  and rewards  $r : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$
- 2: Initialize the state-action value function  $Q : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$  and visit counter  $N : \mathcal{S} \mapsto \mathbb{N}$
- 3: **for**  $n \leftarrow 0, \dots, N - 1$  **do**
- 4:    $t \leftarrow 0$
- 5:   **while**  $N(s_t) > 0$  **do** ▷ Selection
- 6:      $N(s_t) \leftarrow N(s_t) + 1$
- 7:      $a_t \leftarrow \arg \max_{a \in A(s_t)} \left[ Q(s_t, a) + w \sqrt{\frac{\ln N(s_t)}{N(c(s_t, a))}} \right]$
- 8:      $r_t = r(s_t, a_t), s_{t+1} \leftarrow c(s_t, a_t)$
- 9:      $t \leftarrow t + 1$
- 10:   **end while**
- 11:   **while**  $s_t$  is not a terminal state  $\wedge t \leq L$  **do**
- 12:     **for**  $i \leftarrow 1, \dots, d$  **do** ▷ Expansion
- 13:       Sample  $a_t^{(i)} \sim p_\phi(a | s_t), s_{t+1}^{(i)} \sim p_\theta(s_t, a_t^{(i)})$ , and  $r_t^{(i)} \sim r_\theta(s_t, a_t^{(i)})$
- 14:       Update  $A(s_t) \leftarrow \{a_t^{(i)}\}_{i=1}^d, c(s_t, a_t^{(i)}) \leftarrow s_{t+1}^{(i)}$ , and  $r(s_t, a_t) \leftarrow r_t^{(i)}$
- 15:     **end for**
- 16:      $a_{t+1} \leftarrow \arg \max_{a \in A(s_t)} r(s_t, a_t)$  ▷ Simulation
- 17:      $r_t \leftarrow r(s_t, a_t), s_{t+1} \leftarrow c(s_t, a_t)$
- 18:      $t \leftarrow t + 1$
- 19:   **end while**
- 20:   **for**  $t' \leftarrow t, \dots, 0$  **do** ▷ Back propagation
- 21:     Update  $Q(s_{t'}, a_{t'})$  with  $\{r_{t'}, r_{t'+1}, \dots, r_t\}$
- 22:   **end for**
- 23: **end for**

---

## 538 A Related work: Planning and World Model

539 Recent years have witnessed successful applications of planning algorithms [50], such as Alp-  
540 haZero [52], and MuZero [48]. These algorithms are typically based on tree-structured search and are  
541 designed to effectively maintain the balance of exploration and exploitation. Knowledge of transition  
542 dynamics is the prerequisite for planning, and recent research on model-based reinforcement learning  
543 proposed to learn a world model (or dynamics model) to plan or assist policy learning. To improve  
544 sample efficiency, previous research attempts to learn a world model from offline trajectories, and  
545 directly learn a policy within the world model [16, 17]. With latent imagination in a world model, RL  
546 agents can be trained to solve long-horizon tasks [18, 20]. Besides, the world model is also shown to  
547 be helpful to physical robot learning [62]. In this paper, we use LLMs as world models and apply a  
548 planning algorithm to search for a reasoning path. This is similar in spirit to model predictive control  
549 [8]. Compared with previous works, our framework uses general LLMs as the world model and can  
550 be adapted to a wide range of open-domain reasoning tasks.

## 551 B MCTS Planning

552 We adapt MCTS to search for the optimal reasoning path (Algorithm 1). Compared with traditional  
553 applications of MCTS, we are faced with a large reasoning space, and the heavy computational cost of  
554 LLMs. Thus, we made several modifications to the classic MCTS in our implementation: (1) For open  
555 domain problems, e.g., math problems, it’s impossible to enumerate all actions (subquestions), so we  
556 reduce the action space by sampling a fixed number of potential actions from LLMs, conditioned on  
557 a prompt of the current state and in-context demonstration. (2) In the selection phase, if there are  
558 actions that haven’t been visited before, we estimate the Q value with lightweight local rewards, e.g.,  
559 self-evaluation reward, and then select the action with UCT. This provides prior knowledge for the  
560 exploration, which is crucial given the limited iteration budgets.

Table 4: Ablation study on Blocksworld.  $R_1$  is action likelihood reward,  $R_2$  is task-specific reward, and  $R_3$  is self-evaluation reward.

$R_1$	$R_2$	$R_3$	Success
✓	✓	✗	0.88
✓	✓	✓	0.91
✓	✗	✗	0.46
✗	✓	✗	0.21
✗	✗	✓	0.14
✗	✗	✗	0.02

Table 5: Ablation study on GSM8k (first 300 examples).  $R_1$  is state transition confidence reward,  $R_2$  is action likelihood reward, and  $R_3$  is self-evaluation reward.

$R_1$	$R_2$	$R_3$	RAP <sup>(1)</sup>	RAP <sup>(10)</sup>	+aggr
✓	✗	✓	0.410	0.450	0.503
✓	✗	✗	0.350	0.447	0.490
✓	✓	✗	0.373	0.423	0.443

## 561 C Reward Choice

562 Similar to the reward design in RL research, the rewards in LLM reasoning also require some  
 563 curation and design. In our main experiments, we choose the combination of rewards in our current  
 564 experiments based on specific task heuristics and our exploratory experiments. To understand the  
 565 effects of the reward choice for LLM reasoning, we supplement comprehensive experiments on  
 566 rewards for plan generation (Table 4) and math reasoning (Table 5). Note that, in both tables, the first  
 567 row indicates the setting we use in the main experiments.

568 **Experiment results.** As shown in Table 4, the combination of action likelihood and task-specific  
 569 reward (row 1) can significantly outperform the single reward baselines (row 3, 4, 5). Interestingly,  
 570 adding the self-evaluation reward can further improve the performance slightly (row 2). Furthermore,  
 571 as the results on the first 300 samples of GSM8k shown in Table 5, we can see adding either action  
 572 likelihood (row 3) or self-evaluation (row 1) on top of confidence reward (row 2) can boost the  
 573 RAP performance of only using confidence reward (row 1) with one iteration, but action likelihood  
 574 reward downgrades the accuracy with more iterations. The self-evaluation reward leads to the best  
 575 performance overall. This indicates the importance of self-evaluation reward in guiding reasoning as  
 576 an effective and computationally efficient prior to exploration.

577 **On self-evaluation and action likelihood reward.** The rewards of self-evaluation and action likeli-  
 578 hood are of particular interest, as they can be applied to a wide range of reasoning tasks. Generally,  
 579 the best usage and combination with other rewards require empirical design and understanding of the  
 580 task nature, and their effectiveness can vary significantly across different tasks. Here, we provide  
 581 some intuitions behind the reward choices:

582 (a) For the problems in which one reasoning step is short and structured, the action likelihood can  
 583 be very indicative. Otherwise, it may be disturbed by unimportant tokens and become unreliable.  
 584 For instance, a single step within the Blocksworld domain typically adheres to specific patterns  
 585 (e.g., PICK/PUT/STACK a block...), rendering the action likelihood indicative. However, in the math  
 586 domain, a reasoning step is expressed in natural language sentences, allowing for greater freedom  
 587 and potentially introducing noise.

588 (b) For the problems where it’s easier to recognize some errors afterward than avoid them during  
 589 generation, self-evaluation emerges as a helpful mechanism for enhancing reasoning accuracy. In  
 590 mathematical reasoning, LLMs may struggle to generate a correct reasoning step in the first place,  
 591 but the detection of calculation or logic errors is more feasible. In Blocksworlds, however, assessing  
 592 the quality of a candidate action is not straightforward and still requires multi-step reasoning. This  
 593 characteristic diminishes the accuracy of the self-evaluation reward, making it less helpful especially  
 594 given that likelihood already provides a good intuition for search.

## 595 **D Prompt**

### 596 **D.1 Plan Generation**

597 We show the prompt to calculate the action likelihood for RAP below. The same prompt is also  
598 applied in CoT baseline. <init\_state> and <goals> would be instantiated by the problem to  
599 solve.

600 I am playing with a set of blocks where I need to arrange the blocks into  
601 stacks. Here are the actions I can do

602

603 Pick up a block

604 Unstack a block from on top of another block

605 Put down a block

606 Stack a block on top of another block

607

608 I have the following restrictions on my actions:

609 I can only pick up or unstack one block at a time.

610 I can only pick up or unstack a block if my hand is empty.

611 I can only pick up a block if the block is on the table and the block is  
612 clear. A block is clear if the block has no other blocks on top of it  
613 and if the block is not picked up.

614 I can only unstack a block from on top of another block if the block I am  
615 unstacking was really on top of the other block.

616 I can only unstack a block from on top of another block if the block I am  
617 unstacking is clear.

618 Once I pick up or unstack a block, I am holding the block.

619 I can only put down a block that I am holding.

620 I can only stack a block on top of another block if I am holding the block  
621 being stacked.

622 I can only stack a block on top of another block if the block onto which I  
623 am stacking the block is clear.

624 Once I put down or stack a block, my hand becomes empty.

625

626 [STATEMENT]

627 As initial conditions I have that, the red block is clear, the yellow  
628 block is clear, the hand is empty, the red block is on top of the blue  
629 block, the yellow block is on top of the orange block, the blue block  
630 is on the table and the orange block is on the table.

631 My goal is to have that the orange block is on top of the red block.

632

633 My plan is as follows:

634

635 [PLAN]

636 unstack the yellow block from on top of the orange block

637 put down the yellow block

638 pick up the orange block

639 stack the orange block on top of the red block

640 [PLAN END]

641

642 [STATEMENT]

643 As initial conditions I have that, the orange block is clear, the yellow  
644 block is clear, the hand is empty, the blue block is on top of the red  
645 block, the orange block is on top of the blue block, the red block is  
646 on the table and the yellow block is on the table.

647 My goal is to have that the blue block is on top of the red block and the  
648 yellow block is on top of the orange block.

649

650 My plan is as follows:

651



```

652 [PLAN]
653 pick up the yellow block
654 stack the yellow block on top of the orange block
655 [PLAN END]
656
657 [STATEMENT]
658 As initial conditions I have that, the red block is clear, the blue block
659     is clear, the orange block is clear, the hand is empty, the blue block
660     is on top of the yellow block, the red block is on the table, the
661     orange block is on the table and the yellow block is on the table.
662 My goal is to have that the blue block is on top of the orange block and
663     the yellow block is on top of the red block.
664
665 My plan is as follows:
666
667 [PLAN]
668 unstack the blue block from on top of the yellow block
669 stack the blue block on top of the orange block
670 pick up the yellow block
671 stack the yellow block on top of the red block
672 [PLAN END]
673
674 [STATEMENT]
675 As initial conditions I have that, the red block is clear, the blue block
676     is clear, the yellow block is clear, the hand is empty, the yellow
677     block is on top of the orange block, the red block is on the table,
678     the blue block is on the table and the orange block is on the table.
679 My goal is to have that the orange block is on top of the blue block and
680     the yellow block is on top of the red block.
681
682 My plan is as follows:
683
684 [PLAN]
685 unstack the yellow block from on top of the orange block
686 stack the yellow block on top of the red block
687 pick up the orange block
688 stack the orange block on top of the blue block
689 [PLAN END]
690
691 [STATEMENT]
692 As initial conditions I have that, <initial_state>
693 My goal is to have that <goals>.
694
695 My plan is as follows:
696
697 [PLAN]

```

698 For the next state prediction with the world model, we apply the prompts conditioned on the last  
699 action. Here we show the prompt to update the state after a “pick up” action as an example. Again,  
700 <state> and <action> would be instantiated with the current state and action.

```

701 I am playing with a set of blocks where I need to arrange the blocks into
702     stacks. Here are the actions I can do
703
704 Pick up a block
705 Unstack a block from on top of another block
706 Put down a block
707 Stack a block on top of another block
708
709 I have the following restrictions on my actions:

```

710 I can only pick up or unstack one block at a time.  
711 I can only pick up or unstack a block if my hand is empty.  
712 I can only pick up a block if the block is on the table and the block is  
713 clear. A block is clear if the block has no other blocks on top of it  
714 and if the block is not picked up.  
715 I can only unstack a block from on top of another block if the block I am  
716 unstacking was really on top of the other block.  
717 I can only unstack a block from on top of another block if the block I am  
718 unstacking is clear. Once I pick up or unstack a block, I am holding  
719 the block.  
720 I can only put down a block that I am holding.  
721 I can only stack a block on top of another block if I am holding the block  
722 being stacked.  
723 I can only stack a block on top of another block if the block onto which I  
724 am stacking the block is clear. Once I put down or stack a block, my  
725 hand becomes empty.  
726  
727 After being given an initial state and an action, give the new state after  
728 performing the action.  
729  
730 [SCENARIO 1]  
731 [STATE 0] I have that, the white block is clear, the cyan block is clear,  
732 the brown block is clear, the hand is empty, the white block is on top  
733 of the purple block, the purple block is on the table, the cyan block  
734 is on the table and the brown block is on the table.  
735 [ACTION] Pick up the brown block.  
736 [CHANGE] The hand was empty and is now holding the brown block, the brown  
737 block was on the table and is now in the hand, and the brown block is  
738 no longer clear.  
739 [STATE 1] I have that, the white block is clear, the cyan block is clear,  
740 the brown block is in the hand, the hand is holding the brown block,  
741 the white block is on top of the purple block, the purple block is on  
742 the table and the cyan block is on the table.  
743  
744 [SCENARIO 2]  
745 [STATE 0] I have that, the purple block is clear, the cyan block is clear,  
746 the white block is clear, the hand is empty, the white block is on top  
747 of the brown block, the purple block is on the table, the cyan block  
748 is on the table and the brown block is on the table.  
749 [ACTION] Pick up the cyan block.  
750 [CHANGE] The hand was empty and is now holding the cyan block, the cyan  
751 block was on the table and is now in the hand, and the cyan block is  
752 no longer clear.  
753 [STATE 1] I have that, the cyan block is in the hand, the white block is  
754 clear, the purple block is clear, the hand is holding the cyan block,  
755 the white block is on top of the brown block, the purple block is on  
756 the table and the brown block is on the table.  
757  
758 [SCENARIO 3]  
759 [STATE 0] <state>  
760 [ACTION] <action>  
761 [CHANGE]

## 762 D.2 Math Reasoning

763 We show the prompt of RAP for math reasoning as below. The prompt is used for both action proposal  
764 and next state prediction. After instantiate <question>, we append a prefix Question 5.1 to the  
765 prompt, so that we can sample the first action with the LLM. The future actions are sampled similarly,  
766 except that all previous sub-questions and sub-answers need to be appended to the prompt, following

767 the formats of in-context demonstration. The next state prediction, i.e., answering the sub-question,  
768 works in the same way.

769 Given a question, please decompose it into sub-questions. For each  
770 sub-question, please answer it in a complete sentence, ending with  
771 "The answer is". When the original question is answerable, please  
772 start the subquestion with "Now we can answer the question: ".  
773

774 Question 1: Four years ago, Kody was only half as old as Mohamed. If  
775 Mohamed is currently twice as 30 years old, how old is Kody?

776 Question 1.1: How old is Mohamed?

777 Answer 1.1: He is currently  $30 * 2 = 60$  years old. The answer is 60.

778 Question 1.2: How old was Mohamed four years ago?

779 Answer 1.2: Four years ago, he must have been  $60 - 4 = 56$  years old. The  
780 answer is 56.

781 Question 1.3: How old was Kody four years ago?

782 Answer 1.3: Kody was half as old as Mohamed four years ago. Thus, Kody was  
783  $56 / 2 = 28$  years old. The answer is 28.

784 Question 1.4: Now we can answer the question: How old is Kody?

785 Answer 1.4: She is currently  $28 + 4 = 32$  years old. The answer is 32.  
786

787 Question 2: On a moonless night, three fireflies danced in the evening  
788 breeze. They were joined by four less than a dozen more fireflies  
789 before two of the fireflies flew away. How many fireflies remained?

790 Question 2.1: How many fireflies joined?

791 Answer 2.1: The fireflies were joined by four less than a dozen more  
792 fireflies, which are  $12 - 4 = 8$  fireflies. The answer is 8.

793 Question 2.2: Now we can answer the question: How many fireflies remained?

794 Answer 2.2: Three fireflies were dancing originally. They were joined by 8  
795 fireflies before two of them flew away. So there were  $3 + 8 - 2 = 9$   
796 remaining. The answer is 9.  
797

798 Question 3: Ali has four \$10 bills and six \$20 bills that he saved after  
799 working for Mr. James on his farm. Ali gives her sister half of the  
800 total money he has and uses  $3/5$  of the remaining amount of money to  
801 buy dinner. Calculate the amount of money he has after buying the  
802 dinner.

803 Question 3.1: How much money does Ali have in total?

804 Answer 3.1: Ali has four \$10 bills and six \$20 bills. So he has  $4 * 10 + 6$   
805  $* 20 = 160$  dollars. The answer is 160.

806 Question 3.2: How much money does Ali give to his sister?

807 Answer 3.2: Ali gives half of the total money he has to his sister. So he  
808 gives  $160 / 2 = 80$  dollars to his sister. The answer is 80.

809 Question 3.3: How much money does Ali have after giving his sister the  
810 money?

811 Answer 3.3: After giving his sister the money, Ali has  $160 - 80 = 80$   
812 dollars left. The answer is 80.

813 Question 3.4: How much money does Ali use to buy dinner?

814 Answer 3.4: Ali uses  $3/5$  of the remaining amount of money to buy dinner.  
815 So he uses  $80 * 3/5 = 48$  dollars to buy dinner. The answer is 48.

816 Question 3.5: Now we can answer the question: How much money does Ali have  
817 after buying the dinner?

818 Answer 3.5: After buying the dinner, Ali has  $80 - 48 = 32$  dollars left.  
819 The answer is 32.  
820

821 Question 4: A car is driving through a tunnel with many turns. After a  
822 while, the car must travel through a ring that requires a total of 4  
823 right-hand turns. After the 1st turn, it travels 5 meters. After the  
824 2nd turn, it travels 8 meters. After the 3rd turn, it travels a little  
825 further and at the 4th turn, it immediately exits the tunnel. If the

826 car has driven a total of 23 meters around the ring, how far did it  
827 have to travel after the 3rd turn?  
828 Question 4.1: How far did the car travel except for the 3rd turn?  
829 Answer 4.1: It travels 5 meters after the 1st, 8 meters after the 2nd, and  
830 0 meters after the 4th turn. It's a total of  $5 + 8 + 0 = 13$  meters.  
831 The answer is 13.  
832 Question 4.2: Now we can answer the question: How far did the car have to  
833 travel after the 3rd turn?  
834 Answer 4.2: The car has driven a total of 23 meters around the ring. It  
835 travels 13 meters except for the 3rd turn. So it has to travel  $23 - 13$   
836  $= 10$  meters after the 3rd turn. The answer is 10.  
837  
838 Question 5: <question>

### 839 D.3 Logical Reasoning

840 We show the prompt for action proposal, action likelihood calculation, and next state prediction.  
841 <fact> and <query> would be instantiated with the problem.

842 Given a list of facts, and a current claim, output one possible fact as  
843 the next step. Be sure to copy the exact sentences in the facts. Do  
844 not change any wording. Do not create your own words.

845  
846 Facts 1: Each lepidopteran is an insect. Each arthropod is a protostome.  
847 Every animal is multicellular. Protostomes are invertebrates. Each  
848 whale is bony. Each painted lady is a butterfly. Invertebrates are  
849 animals. Butterflies are lepidopterans. Each insect is six-legged.  
850 Every insect is an arthropod. Arthropods are not bony.

851 Query 1: True or false: Sally is not bony.

852 Claim 1.1: Sally is an insect.

853 Next 1.1: Each insect is six-legged.

854 Claim 1.2: Sally is a butterfly.

855 Next 1.2: Butterflies are lepidopterans.

856 Claim 1.3: Sally is a lepidopteran.

857 Next 1.3: Each lepidopteran is an insect.

858 Claim 1.4: Sally is not bony.

859 Next 1.4: Finish.

860 Claim 1.5: Sally is an arthropod.

861 Next 1.5: Arthropods are not bony.

862 Claim 1.6: Sally is a painted lady.

863 Next 1.6: Each painted lady is a butterfly.

864

865 Facts 2: Prime numbers are natural numbers. Every Mersenne prime is not  
866 composite. Imaginary numbers are not real. Every real number is a  
867 number. Natural numbers are integers. Every real number is real. Every  
868 Mersenne prime is a prime number. Natural numbers are positive. Prime  
869 numbers are not composite. Integers are real numbers.

870 Query 2: True or false: 127 is not real.

871 Claim 2.1: 127 is real.

872 Next 2.1: Finish.

873 Claim 2.1: 127 is a natural number.

874 Next 2.1: Natural numbers are integers.

875 Claim 2.2: 127 is a prime number.

876 Next 2.2: Prime numbers are natural numbers.

877 Claim 2.3: 127 is a real number.

878 Next 2.3: Every real number is real.

879 Claim 2.4: 127 is a Mersenne prime.

880 Next 2.4: Every Mersenne prime is a prime number.

881 Claim 2.5: 127 is an integer.

882 Next 2.5: Integers are real numbers.  
883  
884 Facts 3: Lepidopterans are insects. Every animal is multicellular. Each  
885 insect is an arthropod. Each invertebrate is an animal. Insects are  
886 six-legged. Arthropods are small. Arthropods are invertebrates. Each  
887 butterfly is a lepidopteran. Whales are not small.  
888 Query 3: True or false: Polly is not small.  
889 Claim 3.1: Polly is an arthropod.  
890 Next 3.1: Arthropods are small.  
891 Claim 3.2: Polly is an insect.  
892 Next 3.2: Each insect is an arthropod.  
893 Claim 3.3: Polly is small.  
894 Next 3.3: Finish.  
895 Claim 3.4: Polly is a lepidopteran.  
896 Next 3.4: Lepidopterans are insects.  
897  
898 Facts 4: Every cat is a feline. Mammals are vertebrates. Bilaterians are  
899 animals. Vertebrates are chordates. Carnivores are mammals. Mammals  
900 are not cold-blooded. Each chordate is a bilaterian. Every feline is a  
901 carnivore. Snakes are cold-blooded. Animals are not unicellular. Every  
902 carnivore is not herbivorous.  
903 Query 4: True or false: Fae is not cold-blooded.  
904 Claim 4.1: Fae is a feline.  
905 Next 4.1: Every feline is a carnivore.  
906 Claim 4.2: Fae is not cold-blooded.  
907 Next 4.2: Finish.  
908 Claim 4.2: Fae is a mammal.  
909 Next 4.2: Mammals are not cold-blooded.  
910 Claim 4.3: Fae is a cat.  
911 Next 4.3: Every cat is a feline.  
912 Claim 4.4: Fae is a carnivore.  
913 Next 4.4: Carnivores are mammals.  
914  
915 Facts 5: Prime numbers are prime. Real numbers are numbers. Every integer  
916 is a real number. Real numbers are not imaginary. Mersenne primes are  
917 prime numbers. Complex numbers are imaginary. Each prime number is a  
918 natural number. Natural numbers are positive. Each Mersenne prime is  
919 prime. Each natural number is an integer.  
920 Query 5: True or false: 7 is imaginary.  
921 Claim 5.1: 7 is not imaginary.  
922 Next 5.1: Finish.  
923 Claim 5.1: 7 is a natural number.  
924 Next 5.1: Each natural number is an integer.  
925 Claim 5.2: 7 is a prime number.  
926 Next 5.2: Each prime number is a natural number.  
927 Claim 5.3: 7 is a real number.  
928 Next 5.3: Real numbers are not imaginary.  
929 Claim 5.4: 7 is an integer.  
930 Next 5.4: Every integer is a real number.  
931  
932 Facts 6: Spiders are not six-legged. Insects are six-legged. Insects are  
933 arthropods. Every animal is not unicellular. Invertebrates are  
934 animals. Lepidopterans are insects. Every arthropod is segmented.  
935 Arthropods are invertebrates. Every butterfly is a lepidopteran.  
936 Stella is a butterfly.  
937 Query 6: True or false: Stella is six-legged.  
938 Claim 6.1: Stella is an insect.  
939 Next 6.1: Insects are six-legged.  
940 Claim 6.2: Stella is a lepidopteran.

941 Next 6.2: Lepidopterans are insects.  
942 Claim 6.3: Stella is a butterfly.  
943 Next 6.3: Every butterfly is a lepidopteran.  
944 Claim 6.4: Stella is six-legged.  
945 Next 6.4: Finish.  
946  
947 Facts 7: <fact>  
948 Query 7: <query>