

THINKING FORWARD AND BACKWARD: EFFECTIVE BACKWARD PLANNING WITH LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models (LLMs) have exhibited remarkable reasoning and planning capabilities. Most prior work in this area has used LLMs to reason through steps from an initial to a goal state or criterion, thereby effectively reasoning in a *forward* direction. Nonetheless, many planning problems exhibit an inherent *asymmetry* such that planning *backward* from the goal is significantly easier — for example, if there are bottlenecks close to the goal. We take inspiration from this observation and demonstrate that this bias holds for LLM planning as well: planning performance in one direction correlates with the planning complexity of the problem in that direction. However, our experiments also reveal systematic biases which lead to poor planning in the backward direction. With this knowledge, we propose a backward planning algorithm for LLMs that first flips the problem and then plans forward in the flipped problem. This helps avoid the backward bias, generate more diverse candidate plans, and exploit asymmetries between the forward and backward directions in planning problems — we find that combining planning in both directions with self-verification improves the overall planning success rates by 4-24% in three planning domains. Code: [anonymous.repo](#).

1 INTRODUCTION

Large Language Models (LLMs) are increasingly capable in reasoning tasks — they can perform commonsense reasoning in a broad set of contexts (Kojima et al., 2022), reason about abstract patterns in data (Mirchandani et al., 2023), and learn to reason from human feedback (Liang et al., 2024). Such capabilities also open up the possibility of LLMs performing long-horizon planning (Valmeekam et al., 2024), where the model needs to reason about how the initial state and final goal of the problem can be connected through intermediate steps. Most existing work has explored such problems by asking the model to reason in the *forward* direction, i.e., generating intermediate steps from the initial state to the final goal. However, in many planning problems, there is an inherent *asymmetry*: generating the correct last steps leading to the goal can be much easier than generating the correct steps from the beginning. This leads to the question: *can LLMs plan better if they also reason in the backward direction?*

As an example, consider a robot navigating in a room (Fig. 1): if the final goal is the bedroom at the end of a long and narrow hallway, it is natural to connect the bedroom with the beginning of the hallway first in the plan, and then search for the path that connects the hallway to the initial state. In this example, there is a “bottleneck” that causes the asymmetry: the number of possibilities when planning backward from the goal is constrained by the bottleneck (hallway), while the possibilities fan out quickly when planning from the start. Such bottlenecks are ubiquitous in planning problems; for example, in proving mathematical theorems (Loveland, 2016), there may be many possible steps to start a proof of a theorem, but the final steps can be much more closely related to the theorem statement and thus easier to choose. In our experiments, we quantify this bottleneck effect by comparing the exact number of search steps used by common forward-search algorithm (e.g., Breadth-First Search) when applied in the forward and backward directions on the underlying graph of the planning problem (Fig. 1).

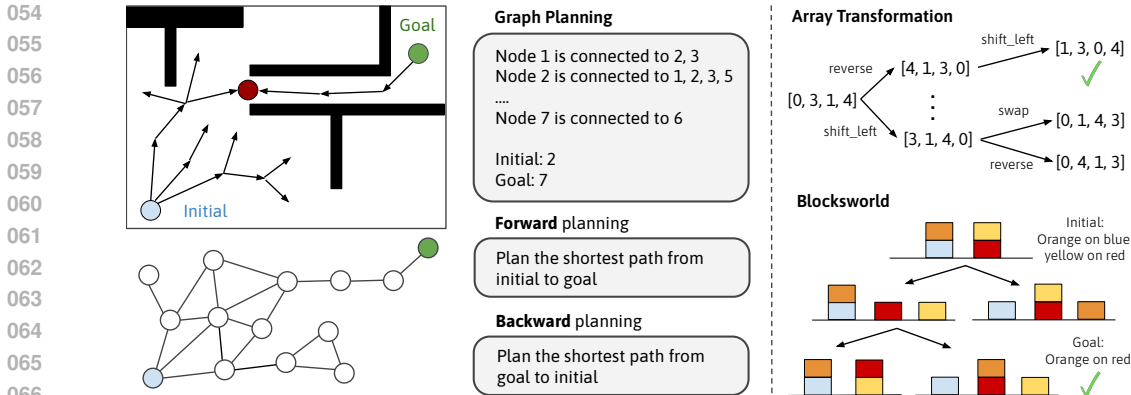


Figure 1: (Left) In planning problems such as navigation, it can often be easier to plan in the backward direction, especially when a “bottleneck” state (red node) exists and it is easier to find it from the goal rather than from the initial state. We consider GRAPH PLANNING, where the LLM needs to plan the shortest path from the initial node to the goal. (Right) Two other planning domains: ARRAY TRANSFORMATION, where an integer array is manipulated through different functions to the desired array, and BLOCKSWORLD, where stacks of blocks are re-oriented to the goal state.

The goals of this paper are twofold. First, we perform extensive experiments on three classical planning problems (Section 2) and examine whether LLM-based planning is significantly impacted by the asymmetry described above, i.e., if LLMs achieve higher planning success rates in the direction (forward/backward) where fewer computations are needed. Second, we study whether LLMs can exploit bottleneck effects by effectively planning backward (Section 3). Our results provide strong evidence for the first hypothesis. However, unfortunately, we find that LLMs exhibit a systematic bias of performing worse when planning backward; this may be attributed to the forward autoregressive nature of LLM output generation, as well as biases from the training dataset.

Addressing the backward reasoning bias, we propose a simple solution (Section 4): many problems can be transformed such that the original goal becomes the initial state and the original initial state becomes the goal — we ask LLMs to first “flip” the problem, and then plan in the forward direction (corresponding to the original backward direction). Given a planning problem, we ask LLMs to sample possible plans in the forward direction of both the original problem and the flipped one, and then self-verify (Stechly et al., 2024) all the plans before choosing the final one. We find that this simple setup allows LLMs to generate many more diverse candidate plans and exploit the bottleneck structure of the problems, improving the overall success rate in multiple planning domains by 4-24% (Section 5). Perhaps surprisingly, we also present evidence that in certain settings LLMs can also *reason whether to flip* the problem or not by examining the problem structure.

2 FORMULATION AND PLANNING DOMAINS

We define a (text-based) planning problem as $P = (\mathcal{S}, \mathcal{A}_\mathcal{S}, s_0, g, T_{\max}, f)$, where \mathcal{S} is the space of possible states, $\mathcal{A}_\mathcal{S} = \{\mathcal{A}_s\}_{s \in \mathcal{S}}$ includes the space of possible actions (planned steps) at each state $s \in \mathcal{S}$, $s_0 \in \mathcal{S}$ is the initial state, $g \in \mathcal{S}$ is the goal, T_{\max} is the maximum allowable steps, and f is a ground-truth solution verifier that determines if a plan solves P , i.e., whether a plan connects s_0 to g , and possibly, is of optimal length. The generated plan $A = (a_0, \dots, a_T)_{T \leq (T_{\max}-1)}$ is verified by f based on these rules. Each action a_t converts state s_t to s_{t+1} , and we denote this as $s_t \xrightarrow{a_t} s_{t+1}$.

We denote pre-trained LLMs with parameters θ as p_θ . Instead of the usual token-level generation, for convenience we consider LLMs generating the next step in a plan by sampling $a_t \sim p_\theta(a_t | s_0, g, \{a_0, \dots, a_{t-1}\}, q)$, where q denotes the rest of the text prompt (e.g., few-shot exemplars and instructions). For convenience, from now on we also assume that q includes text descriptions of s_0, g , and (a_0, \dots, a_{t-1}) . During planning, LLMs may also generate other intermediaries such as the estimated states along the plan $(\hat{s}_1, \dots, \hat{s}_t)$, which might help the LLM reason about the next step given the current estimated state.

Next we introduce the three different planning domains (Fig. 1) considered in our work. Appendix A provides more details on the design of these domains.

Graph Planning. In GRAPH PLANNING, each problem consists of a graph with N nodes (labeled with a number $1, \dots, N$). Each pair of nodes is connected with an edge with probability ρ ; we consider both undirected and directed graphs. The initial state s_0 and goal g are two different nodes. In order to vary difficulty, randomly generated graphs are filtered such that the shortest path from s_0 to g has length K . In this domain, we consider a plan correct if it is optimal, i.e., shortest among paths between s_0 and g . We use a text-based *incident representation* for graphs (e.g., “Node 1 is connected to 2, 3”), which has been shown to be effective in the context of LLMs solving graph problems (Fatemi et al., 2023).

Array Transformation. In ARRAY TRANSFORMATION, the LLM is asked to convert an initial array of integers of length N to a goal array, e.g., from $[0, 1, 3]$ to $[1, 0, 3, 1, 0, 3]$, using a set of array functions. We consider functions including: *repeat* to repeat the array once, *cut* to cut the array in half if the first half matches the second half, *shift_left* and *shift_right* to shift the array to the left by one (e.g., $[0, 1, 3] \rightarrow [1, 3, 0]$) or to the right, *reverse* to reverse the array, and *swap* to swap the first and last elements. Notice that each function has an inverse. We generate the problems by randomly sampling an initial array and a set of K functions which are applied to obtain the final array. We consider a plan correct if it is valid, i.e., if planned functions convert the initial array to the goal. This domain can also be seen as a form of probabilistic context-free grammar (PCFG) (Sipser, 1996).

Blocksworld. BLOCKSWORLD is part of the PlanBench benchmark introduced by Valmeekam et al. (2024) for examining the overall planning capabilities of LLMs. In each problem, there are four blocks colored red, yellow, blue, or orange. The initial and goal states involve different possible stacks of the blocks, and the possible actions include *unstack* <color> block from <color> block, *pick up* <color> block (from table), *put* <color> block on <table>/<color> block. We consider a plan correct if it is valid, i.e., if the planned steps move the initial stack to the goal.

3 LLMs CANNOT PLAN BACKWARD AS EFFECTIVELY AS FORWARD

As described in Section 1, many planning problems have an asymmetry that makes a given direction (forward vs. backward) easier for classical planning algorithms such as breadth-first search (BFS). Here, we examine whether the planning performance of LLMs is similarly impacted, and if the LLM’s performance in a given planning direction can be predicted by the number of search steps required by BFS in that direction. If this is the case, it opens up the possibility of improving LLM-based planning by planing backward when the backward direction is favorable. Below we first introduce the backward planning algorithm we use for LLMs before showing the experimental results.

3.1 BACKWARD PLANNING

During backward planning, the LLM is given the same initial and goal states, but the prompt q asks it to generate the plan A in the reversed order. For example, consider an instance of GRAPH PLANNING where the shortest path from Node 1 to Node 3 is $(1, 5, 2, 3)$; then, the LLM should first output the backward plan as $(3, 2, 5, 1)$ and reverse the order before it proposes the plan to the verifier.

As described in Section 1, backward planning can be useful when the search space from the goal is smaller than from the initial state. We can quantify the asymmetry between forward and backward planning by computing the number of search steps used by algorithms such as BFS in either direction.




Figure 2: Sampling in forward or backward direction and then self-verifying the plans in ARRAY TRANSFORMATION.

3.2 ALGORITHM: SAMPLE FORWARD (BACKWARD) THEN SELF-VERIFY

While the basic version of LLM planning involves asking it to generate a single plan, we consider a more robust version where the LLM first samples multiple candidate solutions in one direction, and then self-verifies them to choose the final solution; then the verifier f checks the final plan. In both stages, the LLM is shown few-shot exemplars to familiarize itself with the problem and to learn to verify the candidate plans step by step. The overall algorithm is illustrated in Fig. 2. Given the maximum number of attempts M , the sampling temperature of the LLM is set to be 0 for the first attempt, and nonzero for the later ones. Self-verification works differently depending on the domain: in GRAPH PLANNING, since the solution needs to be optimal, we save all the unique candidate plans $\{A_j\}$ after all M attempts, and then present all of them to the LLM and ask it to verify them as well as to choose the optimal one; in ARRAY TRANSFORMATION and BLOCKSWORLD, since the solution does not need to be optimal, the LLM verifies each candidate plan as soon as it is sampled, and either stops when it deems one plan correct or after all attempts. For all the experiments, we use the GPT-4o model from OpenAI unless noted otherwise.

3.3 RESULTS

LLM plans better in the easier direction. Here we run experiments with both directed and undirected graphs in GRAPH PLANNING, and ask the LLM to plan either forward or backward. We also compute the number of search steps used by BFS in both directions. Fig. 3 shows the planning success rates achieved by forward and backward planning at different levels of forward/backward difficulty (as quantified by BFS computations). For either direction, the success rate is generally higher when the number of computations is lower in that direction. This finding suggests that LLM planning is akin to a forward-search algorithm such as BFS in terms of the difficulty of planning in a given direction, and thus can be potentially improved by planning backward when it is easier. We also find similar results with ARRAY TRANSFORMATION shown in Appendix B.

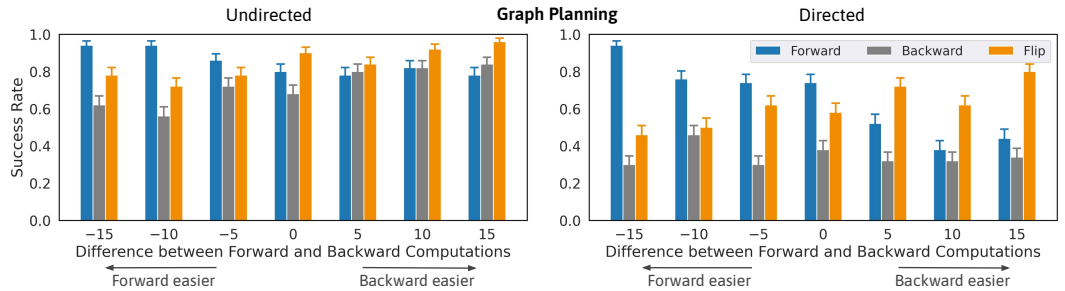


Figure 3: Success rates achieved by forward and backward planning in GRAPH PLANNING vs. difference between forward and backward BFS computations. The LLM plans better in the direction of fewer computations needed, but the forward direction outperforms backward in general.

LLM plans worse in the backward direction.

Next, we study the effectiveness of backward planning by examining whether the LLM can achieve the same level of planning success in the backward direction as compared to forward. We find that this is not the case. Fig. 3 shows that the backward success rate is consistently lower than forward. We also calculate the average planning success rates for all four settings (from experiments in Section 5) — as shown in Table 1, the LLM consistently plans worse in the backward direction. We conjecture that this bias may be attributed to the forward (i.e., left to right) autoregressive nature of LLM output generation, as well as biases from the training dataset. Next we propose a solution to such backward bias and allow LLMs to plan effectively in the backward direction (Flip in Fig. 3).

Domain	Forward	Backward
Graph Planning (undirected)	82.5±2.7%	76.7±3.0%
Graph Planning (directed)	38.5±3.4%	29.7±3.2%
Array Transformation	67.5±3.3%	62.2±3.4%
Blocksworld	39.5±3.5%	20.5±2.9%

Table 1: Average planning success rate in forward vs. backward direction in each setting.

4 PROPOSED SOLUTION: PLAN BACKWARD BY FLIPPING THE PROBLEM

If the LLM cannot plan backward as well as forward, how can it effectively exploit the backward direction in planning? We propose a simple solution: *flip* the problem such that the original goal becomes the new initial state and vice versa. The LLM can then plan in the forward direction for the flipped problem, which corresponds to the backward direction of the original problem. This avoids the bias of weak LLM planning in the backward direction. However, there are a few subtleties that must be taken care of when flipping the planning problem.

Change of the state-dependent action space. In some cases, \mathcal{A}_S needs to be adjusted for the flipped problem. For example, with a directed graph, an edge from Node 1 to Node 3 in the original problem corresponds to an edge from Node 3 to Node 1 in the flipped problem — Node 3 might not be reachable from Node 1 in the flipped problem. In contrast, there is no change needed for undirected graphs, ARRAY TRANSFORMATION, and BLOCKSWORLD. With directed graphs, we prompt the LLM to first generate the new text representation of the graph (Fig. 4 right), and then generate a plan for the flipped problem.

Flipping back the plan. After the plan A' for the flipped problem is generated, steps within it need to be reversed in order and also often “flipped back” to generate a forward plan for the original problem:

$$\text{Flip back the plan: } A' = \{a'_0, \dots, a'_T\} \mapsto A = \{a_T, \dots, a_0\}, \text{ where } s_t \xrightarrow{a'_t} s_{t+1}, s_{t+1} \xrightarrow{a_t} s_t. \quad (1)$$

We assume that each action $a \in \mathcal{A}_s$ can be inverted, i.e., if $s \xrightarrow{a} s'$, there exists $a' \in \mathcal{A}_{s'}$ such that $s' \xrightarrow{a'} s$. In practice, we prompt the LLM to generate the corresponding a_t after it generates a'_t . Fig. 4 left shows an example for BLOCKSWORLD: the action “put yellow on red” is flipped back to “unstack yellow from red”, after the order of the plan is flipped.

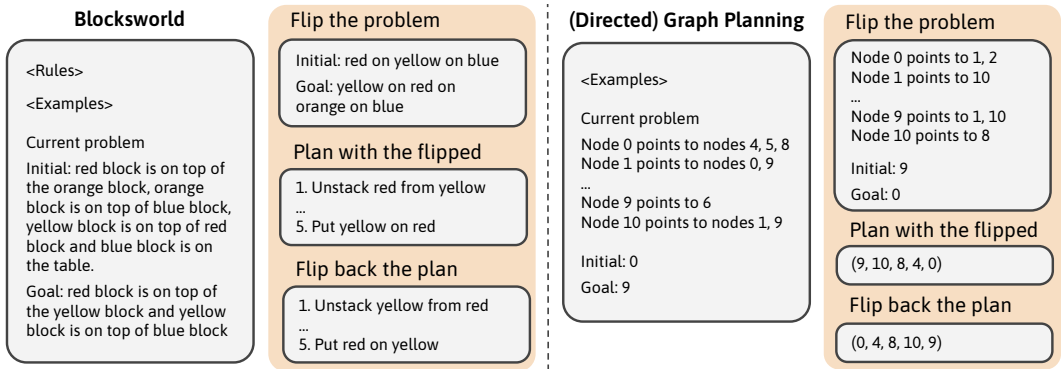


Figure 4: Flipping the problem and then plan in the new forward direction in BLOCKSWORLD (left) and GRAPH PLANNING with directed graph (Right).

Algorithm: Sample forward and flipped plans, then self-verify. Extending the algorithm from Section 3.2 where multiple plans in the forward or backward directions are sampled before self-verifying them, we can now sample in the forward direction of the flipped problem, instead of sampling backward in the original problem. Sampling forward with the original and the flipped problem generates more diverse solutions for self-verification, while avoiding the backward bias and potentially exploiting asymmetries. We can either randomly sample the problem (original or flipped) at each attempt, or prompt the LLM to reason which direction is favorable.

5 EXPERIMENTS WITH FLIPPING THE PROBLEM

With experiments in the three planning domains introduced in Section 2, we investigate the following questions in the corresponding subsections below:

- 270 Q1: Does flipping the problem help improve the planning success rate?
 271
 272 Q2: When does flipping the problem help the most?
 273
 274 Q3: How well can the LLM reason about when to flip the problem?

Baselines. We compare a few different ways to sample candidate solutions before self-verification: (1) **Fwd**: only planning in the forward direction; (2) **Back**: only planning in the backward direction; (3) **Flip**: only planning in the forward direction of the flipped problem; (4) **Fwd-Back**: randomly choose either forward or backward direction and then plan; (5) **Fwd-Flip**: randomly choose either forward direction of the original problem or forward direction of the flipped problem and then plan; (6) **Fwd-Flip-Reason**: have the LLM reason whether to plan forward in the original or flipped problem and then plan. We expect our proposed **Fwd-Flip** to outperform Fwd, Back, Flip, Fwd-Flip.

5.1 BACKWARD PLANNING WITH THE FLIPPED PROBLEM HELPS IMPROVE SUCCESS RATE

In order to test whether flipping the problem can help planning success rate, we run extensive experiments with the three planning domains and compare Fwd, Back, Flip, Fwd-Back, and Fwd-Flip. We use a maximum $M = 6$ attempts for all experiments. Table 2 shows each result averaged over 200 trials. Flip outperforms Back and matches Fwd performance; we also see the similar trends in Fig. 3. In addition, Fwd-Flip consistently leads to the highest planning success rate, improving by 4-24% over Fwd, except for one setting where all baselines achieve close to 100% success. These results corroborate that flipping the problem mitigates the backward bias.

Domain				Fwd	Back	Flip	Fwd-Back	Fwd-Flip
Graph Planning	Directed?	ρ	N					
	No	0.2	12	82.5±2.7%	74.5±3.1%	84.5±2.6%	91.0±2.0%	93.5±1.7%
	No	0.2	10	89.5±2.2%	78.0±2.9%	88.0±2.3%	93.0±1.8%	96.5±1.3%
	No	0.3	10	90.5±2.1%	77.0±3.0%	91.0±2.0%	93.0±1.8%	94.5±1.6%
	Yes	0.2	12	62.5±3.4%	29.0±3.2%	67.0±3.3%	76.0±3.0%	80.0±2.8%
	Yes	0.2	10	73.0±3.1%	36.5±3.4%	73.5±3.1%	79.5±2.9%	88.5±2.3%
Array Transformation	Functions							
	shift, repeat, cut			99.0±0.7%	98.5±0.9%	99.0±0.7%	100.0±0.0%	99.5±0.5%
	shift, reverse, swap			50.0±3.5%	36.0±3.4%	52.0±3.5%	46.0±3.5%	56.0±3.5%
Blocksworld	repeat, cut, reverse, swap			53.5±3.5%	52.0±3.5%	53.0±3.5%	54.5±3.5%	56.5±3.5%
		-		39.5±3.5%	20.5±2.9%	34.5±3.4%	27.0±3.1%	48.5±3.5%

Table 2: Planning success rate averaged over 200 trials for the five methods in the three planning domains. Flip matches Fwd. Fwd-Flip generally achieves the highest success rate.

Fwd-Flip exploits asymmetries in the problems. One of the main motivations of planning backward is that many planning problems have a bottleneck structure: there exists a less connected part of the state space near the initial state or the goal — this makes it easier to plan from the end closer to the bottleneck. In Fig. 5, we calculate the difference between forward vs. backward computations for BFS in GRAPH PLANNING, bin them, and find the success rate of each bin for Fwd and Fwd-Flip — the plot shows the average success rate over the three directed graph settings (shades show the minimum and maximum over the three). We find that Fwd tends to perform worse when the forward computations needed are higher, meaning that it cannot plan as effectively when the forward direction is more difficult. In contrast, Fwd-Flip maintains a similar level of success regardless of forward vs. backward planning difficulty.




Figure 5: Fwd-Flip plans well even when the forward direction is difficult.

Flipping the problem generates more diverse candidate solutions. The improved performance of Fwd-Flip over Fwd can also be partly attributed to the generation of more diverse solutions. Fig. 6

compares Fwd vs. Fwd-Back vs. Fwd-Flip in ARRAY TRANSFORMATION and shows both the average number of unique candidate solutions generated (top) and the planning success rate (bottom) vs. the number of planning attempts, M . Fwd-Flip generates a higher number of unique candidates, and improves the planning success rate by a large margin. Fig. 7 (left) shows an example where Fwd fails to solve the block task due to a persistent error (trying to move a block before “unstacking” it, which is not allowed by the rules of BLOCKSWORLD) even with significant sampling temperature (0.4), while Fwd-Flip generates a different solution and avoids the error.

5.2 IMPROVEMENT FROM FLIPPING THE PROBLEM VARIES

Effect of backward success rate without flipping. While in Section 5.1 we have shown that flipping the problem helps the LLM exploit asymmetries between forward and backward planning, we tend to find that the performance difference between Fwd-Flip and Fwd-Back varies substantially in ARRAY TRANSFORMATION (Table 2). When functions repeat and cut are used, Back does not perform significantly worse than Fwd — we believe this is due to the backward step being easily recognizable by Back (Fig. 7 right) when the initial and goal array sizes differ (since repeat and cut change them). In this case, Fwd-Flip does not improve much over Fwd-Back. With functions shift_left|right, reverse, swap that do not change array sizes, we see a much more significant backward bias (53.5% vs. 36% between Fwd and Back) and there is a bigger margin between Fwd-Flip and Fwd-Back (56% vs. 46%).




Figure 6: Fwd-Flip generates diverse solutions and improves success rate.

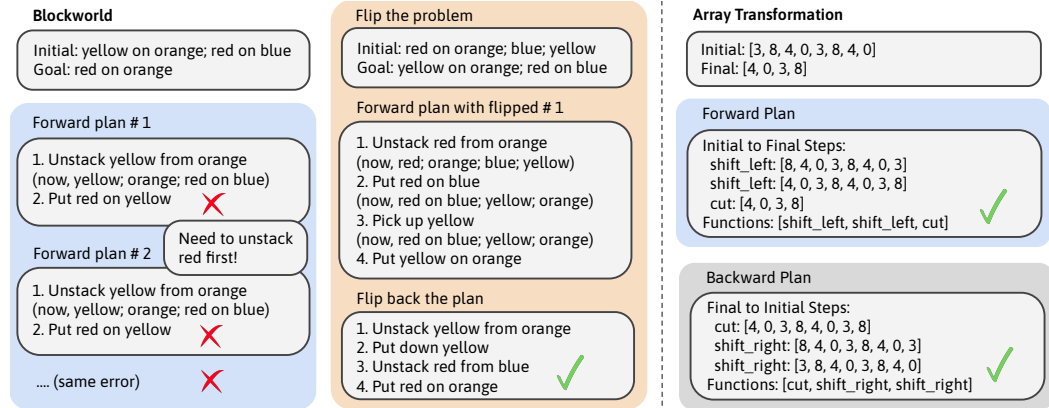


Figure 7: (Left) Flipping the problem helps LLM avoid persistent error in one direction. (Right) With functions cut and repeat used in ARRAY TRANSFORMATION that change the array size, it can become obvious that one of them has to be used, which makes backward planning easier.

Effect of LLM capability. We also hypothesize that the effect of planning in the flipped problem also depends on the inherent capability of the LLM. First, regardless of the choice of the LLM, Fwd-Flip should consistently improve the success rate. In Fig. 8 we run GPT-3.5-turbo and GPT-4-turbo besides GPT-4o with a directed graph setting and in BLOCKSWORLD, and we see that Fwd-Flip consistently outperforms Fwd. We also find that the design choices of the algorithm can affect the performance, and we highlight two aspects below:

- **Reliability of self-verification:** while we allow the LLM to generate multiple candidates (in either direction), it can only present a single solution after it self-verifies the candidates. Hence, the LLM needs to reliably self-verify the candidate plans to achieve a high success rate. In addition to the success rate, Fig. 8 shows the self-verification error rate. We find that the self-verification error reduces as the LLM becomes more capable. In BLOCKSWORLD, we find that GPT-3.5-turbo always self-verifies its sampled plan to be correct, leading to close to zero success rate.

378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431




Figure 8: Planning performance of different LLMs with Fwd and Fwd-Flip in (directed) GRAPH PLANNING and BLOCKSWORLD. Flipping the problem helps all LLMs plan better, and better LLMs reduce errors when executing the algorithm.

- Reliability of flipping the state-dependent action space: in Section 4 we described how flipping the problem requires flipping back the plan and possibly changing the state-dependent action space. While we find that the LLM can flip back the plan itself without error, sometimes it can be error-prone when flipping the action space, affecting the performance of Flip and Fwd-Flip. In GRAPH PLANNING with directed graphs, flipping the problem involves flipping the edges and re-ordering them — Fig. 9 shows an example. Possible errors in re-ordering the edges cause Flip to under-perform Fwd in some of the directed graph settings, while Flip always matches Fwd with undirected graphs.




Figure 9: When flipping the directed graph, we find it is helpful to re-order the flipped edges such that they appear in the order of “Node 0”, “Node 1”, etc., to achieve better success rate. However, this also introduces possible error, causing Flip to under-perform Fwd sometimes and affecting Fwd-Flip.

Nonetheless, as Fig. 8 illustrates, these errors are reduced with better LLMs. Flipping the problem improves success rates significantly even with improvements in the underlying LLM, demonstrating the scalability of our method.

Effect of initial/goal asymmetry. We also notice that in BLOCKSWORLD, Flip under-performs Fwd (Table 2). The reason is that while the initial states of the problems are completely specified (e.g., “yellow block on red block, orange block on blue block”), the goal can be partially specified (e.g., “red block on orange block”). This means that when the initial and goal states are flipped, the LLM needs to first generate a complete flipped initial state from the original goal state — planning with full initial and goal states can be more challenging than with full initial and partial goal states.

Effect of the number of planning attempts. Lastly, we also find that the maximum number of planning attempts, M , can affect the relative performance between Fwd and Fwd-Flip — this is particularly true when Flip under-performs Fwd such as in BLOCKSWORLD due to the initial/goal asymmetry mentioned above. In Fig. 10 we show Fwd vs. Fwd-Flip as M varies from 1 to 6 — Fwd-Flip under-performs Fwd slightly at $M = 1$, but out-performs Fwd with $M > 1$ with increasing improvements.




Figure 10: Fwd-Flip performs better given more planning attempts.

5.3 LLM CAN REASON WHEN TO FLIP IN CERTAIN SETTINGS

With Fwd-Flip we randomly choose between the original and the flipped problem for each candidate solution. We now explore whether the LLM can also reason when to flip the problem — we

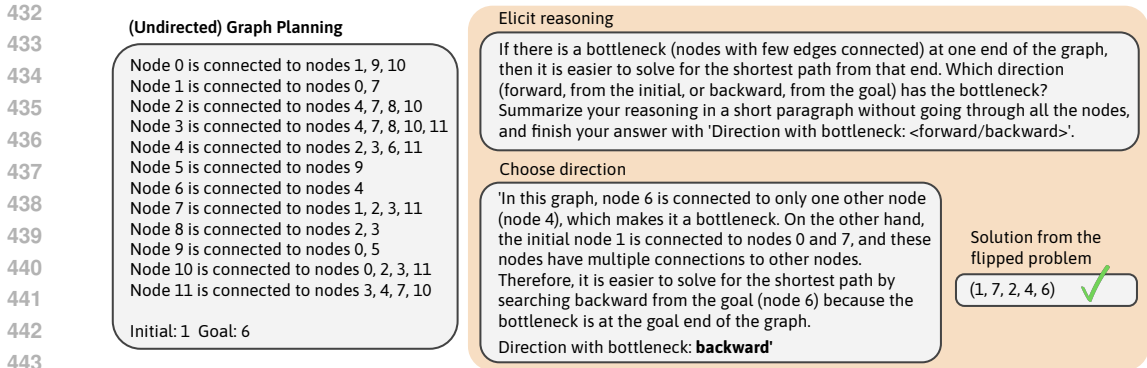


Figure 11: LLM can analyze the graph and choose the easier direction for planning in zero-shot.

hypothesize that the LLM can analyze the problem structure to some extent and reason whether the backward direction can be easier. We run experiments with one setting from each of the planning domains (Table 3); unlike previous experiments, here we use $M = 1$ as we find that the LLM is often certain about the direction. We find that Fwd-Flip-Reason out-performs Fwd and Flip in undirected graphs. Fig. 11 shows an example where the LLM identifies the bottleneck near the goal and chooses to plan backward, leading to the correct answer. We then calculate the rate at which the LLM chooses the direction where the number of computations is lower and find that the LLM can choose the easier direction 78.5% of the time with undirected graphs and 60.5% of the time with directed graphs. However, with directed graphs, the LLM suffers from possible errors when flipping the problem, and thus Fwd-Flip-Reason with $M = 1$ does not perform better than Fwd.

In ARRAY TRANSFORMATION and BLOCKSWORLD, since the problem asymmetry is less clear, we do not expect the LLM to find a better direction in general, and thus Fwd-Flip-Reason does not out-perform Fwd or Flip. Appendix A provides details on the prompts used to elicit LLM reasoning.

Domain	Fwd	Flip	Fwd-Flip-Reason
Graph Planning (Undirected)	79.0±2.9%	78.5±2.9%	87.5±2.3%
Graph Planning (Directed)	50.5±3.5%	46.0±3.5%	50.5±3.5%
Array transformation (shift, reverse, swap)	18.5±2.7%	19.5±2.8%	16.5±2.6%
Blocksworld	33.0±3.3%	21.5±2.9%	27.0±3.1%

Table 3: Planning success rate averaged over 200 trials for the five methods in the three planning domains when the LLM is asked to choose the direction to reason in. The maximum planning attempt M is set to 1.

6 RELATED WORK

Bidirectional planning. It is well known in classical planning (LaValle, 2006) that searching and planning from the backward direction can often reduce the computations needed. Bi-directional search has been incorporated into popular sampling-based planners (e.g., rapidly-exploring random trees) (Jordan & Perez, 2013) and heuristics-based techniques (e.g., A*) (Kuroiwa & Fukunaga, 2020) to improve efficiency.

Planning with LLMs. LLMs have been widely applied in different planning problems (Kambhampati et al.) such as text games (Yao et al., 2024; Wang et al., 2023), robot planning (Huang et al., 2022; Ahn et al., 2022; Ren et al., 2023), scientific experimentation (Wang et al., 2022a), and web navigation (Deng et al., 2024) — their strong planning capabilities originate from pre-training with enormous amounts of text data, which elicits reasoning capabilities (Wei et al., 2022). Beyond simple zero- or few-shot learning, LLMs have also been combined with classical planners (Liu et al., 2023; Silver et al., 2024) and external tools (Zeng et al., 2022) to further boost performance. However, no

486 previous work has systematically studied asymmetries between forward and backward planning for
487 LLMs, or leveraged backward reasoning to improve LLM planning.
488

489 **Backward reasoning in LLMs.** Despite not being applied in planning problems, backward reason-
490 ing has been used with LLMs for self-verifying forward-sampled solutions to math problems (Jiang
491 et al., 2023). Another line of work (Kazemi et al., 2022; Lee & Hwang, 2024) applies backward
492 chaining, which recursively breaks the target into sub-targets based on defined rules, mostly for veri-
493 fying a statement or proof. Our work instead proposes a general solution to planning problems using
494 both forward and backward reasoning, and is motivated by a systematic examination of asymmetries
495 in forward and backward planning for LLMs.
496

497 7 CONCLUSION AND FUTURE WORK 498

499 In this work, we investigate how to enable LLMs to effectively plan in the backward direction
500 from the desired goal to improve the overall planning success rate. First, our experiments reveal
501 consistently worse performance of LLMs when planning backward. To address this, we propose
502 instructing the LLM to flip the problem first and then plan forward in the flipped problem. Combined
503 with self-verification, we find that generating candidate solutions from both directions improves
504 planning success rates by 4-24% over forward-only planning in three different domains.

505 One immediate future direction is to better teach LLMs to reason and plan backward, e.g., by
506 fine-tuning with correct forward and *backward* reasoning traces (Zelikman et al., 2022). We also
507 believe that our framework of combining forward and backward reasoning can be extended to general
508 reasoning problems, e.g., allowing LLMs to generate more diverse reasoning traces from the backward
509 direction (Yao et al., 2024), or, enforcing reasoning self-consistency from both forward and backward
510 directions (Wang et al., 2022b).
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

REFERENCES

- 540
541
542 Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea
543 Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say:
544 Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- 545 Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su.
546 Mind2web: Towards a generalist agent for the web. *Advances in Neural Information Processing*
547 *Systems*, 36, 2024.
- 548 Bahare Fatemi, Jonathan Halcrow, and Bryan Perozzi. Talk like a graph: Encoding graphs for large
549 language models. *arXiv preprint arXiv:2310.04560*, 2023.
- 551 Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot
552 planners: Extracting actionable knowledge for embodied agents. In *International Conference on*
553 *Machine Learning*. PMLR, 2022.
- 554 Weisen Jiang, Han Shi, Longhui Yu, Zhengying Liu, Yu Zhang, Zhenguo Li, and James Kwok.
555 Forward-backward reasoning in large language models for mathematical verification. 2023.
- 556
557 Matthew Jordan and Alejandro Perez. Optimal bidirectional rapidly-exploring random trees. 2013.
- 558 Subbarao Kambhampati, Karthik Valmeekam, Matthew Marquez, and Lin Guan. On the role of
559 large language models in planning, july 2023. tutorial presented at the international conference on
560 automated planning and scheduling (icaps), prague.
- 561 Mehran Kazemi, Najoung Kim, Deepti Bhatia, Xin Xu, and Deepak Ramachandran. Lambda:
562 Backward chaining for automated reasoning in natural language. *arXiv preprint arXiv:2212.13894*,
563 2022.
- 564
565 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large
566 language models are zero-shot reasoners. *Advances in neural information processing systems*,
567 2022.
- 568 Ryo Kuroiwa and Alex Fukunaga. Front-to-front heuristic search for satisficing classical planning.
569 In *IJCAI*, pp. 4098–4105, 2020.
- 570
571 Steven M LaValle. *Planning algorithms*. Cambridge university press, 2006.
- 572 Jinu Lee and Wonseok Hwang. Symba: Symbolic backward chaining for multi-step natural language
573 reasoning. *arXiv preprint arXiv:2402.12806*, 2024.
- 574
575 Jacky Liang, Fei Xia, Wenhao Yu, Andy Zeng, Montserrat Gonzalez Arenas, Maria Attarian, Maria
576 Bauza, Matthew Bennice, Alex Bewley, Adil Dostmohamed, et al. Learning to learn faster from
577 human feedback with language model predictive control. *arXiv preprint arXiv:2402.11450*, 2024.
- 578 Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and Peter Stone.
579 Llm+ p: Empowering large language models with optimal planning proficiency. *arXiv preprint*
580 *arXiv:2304.11477*, 2023.
- 581
582 Donald W Loveland. *Automated theorem proving: A logical basis*. Elsevier, 2016.
- 583
584 Suvir Mirchandani, Fei Xia, Pete Florence, Brian Ichter, Danny Driess, Montserrat Gonzalez Arenas,
585 Kanishka Rao, Dorsa Sadigh, and Andy Zeng. Large language models as general pattern machines.
586 *arXiv preprint arXiv:2307.04721*, 2023.
- 587 Allen Z Ren, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Peng Xu,
588 Leila Takayama, Fei Xia, Jake Varley, et al. Robots that ask for help: Uncertainty alignment for
589 large language model planners. *arXiv preprint arXiv:2307.01928*, 2023.
- 590 Tom Silver, Soham Dan, Kavitha Srinivas, Joshua B Tenenbaum, Leslie Kaelbling, and Michael Katz.
591 Generalized planning in pddl domains with pretrained large language models. In *Proceedings of*
592 *the AAAI Conference on Artificial Intelligence*, 2024.
- 593 Michael Sipser. Introduction to the theory of computation. *ACM Sigact News*, 27(1):27–29, 1996.

594 Kaya Stechly, Karthik Valmeekam, and Subbarao Kambhampati. On the self-verification limitations
595 of large language models on reasoning and planning tasks. *arXiv preprint arXiv:2402.08115*, 2024.
596

597 Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambham-
598 pati. Planbench: An extensible benchmark for evaluating large language models on planning and
599 reasoning about change. *Advances in Neural Information Processing Systems*, 2024.

600 Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlikar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and
601 Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv*
602 *preprint arXiv:2305.16291*, 2023.
603

604 Ruoyao Wang, Peter Jansen, Marc-Alexandre Côté, and Prithviraj Ammanabrolu. Scienceworld: Is
605 your agent smarter than a 5th grader? *arXiv preprint arXiv:2203.07540*, 2022a.

606 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
607 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.
608 *arXiv preprint arXiv:2203.11171*, 2022b.
609

610 Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama,
611 Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models.
612 *arXiv preprint arXiv:2206.07682*, 2022.

613 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan.
614 Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural*
615 *Information Processing Systems*, 2024.

616 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. Star: Bootstrapping reasoning with
617 reasoning. *Advances in Neural Information Processing Systems*, 35:15476–15488, 2022.
618

619 Andy Zeng, Maria Attarian, Brian Ichter, Krzysztof Choromanski, Adrian Wong, Stefan Welker,
620 Federico Tombari, Aveek Purohit, Michael Ryoo, Vikas Sindhwani, et al. Socratic models:
621 Composing zero-shot multimodal reasoning with language. *arXiv preprint arXiv:2204.00598*,
622 2022.
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

A APPENDIX - PLANNING DOMAINS

We provide all codes needed to replicate the experiments in the attached supplementary materials.

A.1 GRAPH PLANNING

Configurations. To generate the desired graphs, we use `gnp_random_graph` function from the `networkx` package, with which we specify N , the number of nodes, ρ , the probability of two nodes are connected with an edge, and whether the edges are directed or not. We also apply rejection sampling to ensure the shortest path involves a total of five nodes. The the initial and goal nodes are sampled randomly from the graph.

Prompts. For both the prompts for sampling the candidate plans and self-verifying them, we randomly generate three similar graphs as few-shot examples (listing 1, listing 2). With directed graph, we also prompt the model to re-order the flipped edges such that the graph still starts with “Node 0 points to ...” instead of “Node 7, 9 points to 0”, if the original problem has “Node 0 points to 7, 9.” (listing 3), which we find improve the planning success with the flipped problem. listing 4 shows the prompt used for eliciting LLM’s preference over the planning directions in zero-shot. We use temperature $T = 0$ for re-ordering, self-verifying, eliciting preference, and the first planning attempt, and $T = 0.5$ for the rest of planning attempts.

Listing 1: Sample prompt used in GRAPH PLANNING for sampling candidate solution. We use “is connected to” for edges in undirected graphs, and “points to” for edges in directed graph. For backward planning, we use “Plan the shortest path from goal to initial node...” as the final instruction.

You will be given an undirected graph search problem with a few examples.

```

673 ** Example 1 **
674 Node 1 is connected to nodes 2, 3, 6, 11
675 Node 2 is connected to nodes 1, 3
676 Node 3 is connected to nodes 1, 2
677 Node 4 is connected to nodes 7
678 Node 6 is connected to nodes 1, 9, 10
679 Node 7 is connected to nodes 4, 8, 9
680 Node 8 is connected to nodes 7, 11
681 Node 9 is connected to nodes 6, 7
682 Node 10 is connected to nodes 6, 11
683 Node 11 is connected to nodes 1, 8, 10
684
685 Initial: 3
686 Goal: 7
687 Shortest Path: (3, 1, 11, 8, 7)
688
689 ** Example 2 **
690 Node 0 is connected to nodes 3, 4, 6, 7, 9
691 Node 1 is connected to nodes 10, 11
692 Node 2 is connected to nodes 7
693 Node 3 is connected to nodes 0, 9, 11
694 Node 4 is connected to nodes 0
695 Node 5 is connected to nodes 10, 11
696 Node 6 is connected to nodes 0
697 Node 7 is connected to nodes 0, 2, 8, 9, 11
698 Node 8 is connected to nodes 7, 11
699 Node 9 is connected to nodes 0, 3, 7
700 Node 10 is connected to nodes 1, 5, 11
701 Node 11 is connected to nodes 1, 3, 5, 7, 8, 10
702
703 Initial: 6
704 Goal: 1
705 Shortest Path: (6, 0, 7, 11, 1)
706
707 ** Example 3 **
708 Node 0 is connected to nodes 1, 9
709 Node 1 is connected to nodes 0, 3, 5
710 Node 2 is connected to nodes 3, 7, 10
711 Node 3 is connected to nodes 1, 2, 6, 11
712 Node 4 is connected to nodes 8
713 Node 5 is connected to nodes 1, 7, 11
714 Node 6 is connected to nodes 3, 11
715 Node 7 is connected to nodes 2, 5
716 Node 8 is connected to nodes 4, 9
717 Node 9 is connected to nodes 0, 8
718 Node 10 is connected to nodes 2
719 Node 11 is connected to nodes 3, 5, 6
720
721 Initial: 8
722 Goal: 3
723 Shortest Path: (8, 9, 0, 1, 3)

```

```

702
703 ** Current problem **
704 Node 0 is connected to nodes 5
705 Node 1 is connected to nodes 2, 4, 9
706 Node 2 is connected to nodes 1, 5, 7
707 Node 3 is connected to nodes 6, 7
708 Node 4 is connected to nodes 1
709 Node 5 is connected to nodes 0, 2, 6, 11
710 Node 6 is connected to nodes 3, 5
711 Node 7 is connected to nodes 2, 3, 10
712 Node 8 is connected to nodes 9
713 Node 9 is connected to nodes 1, 8
714 Node 10 is connected to nodes 7
715 Node 11 is connected to nodes 5
716
717 Initial: 6
718 Goal: 4
719
720 Plan the shortest path from initial to goal node for the this **undirected** graph. Follow the format 'Shorest
721 Path: (...)' and do not output anything else.

```

Listing 2: Sample prompt used in GRAPH PLANNING for self-verifying the candidate solutions.

```

716
717 You will be given a directed graph search problem with a few examples.
718
719 ** Example 1 **
720 Node 0 points to nodes 1, 8, 10, 11
721 Node 1 points to nodes 0, 3, 4
722 Node 2 points to nodes 3
723 Node 3 points to nodes 4, 6
724 Node 4 points to nodes 2, 5, 7, 9
725 Node 5 points to nodes 9
726 Node 6 points to nodes 0, 7, 11
727 Node 7 points to nodes 5, 10
728 Node 8 points to nodes 0, 11
729 Node 9 points to nodes 7, 11
730 Node 10 points to nodes 4, 7, 8, 9
731 Node 11 points to nodes
732
733 Initial: 6
734 Goal: 9
735
736 Which one is the correct shortest path?
737 A. (6, 0, 1, 4, 7, 10, 9)
738 B. (6, 7, 5, 9)
739 Checking each options step by step:
740 A: check 6 to 0, 6 points to [0, 7, 11], 0 in [0, 7, 11]? True; check 0 to 1, 0 points to [1, 8, 10, 11], 1 in
741 [1, 8, 10, 11]? True; check 1 to 4, 1 points to [0, 3, 4], 4 in [0, 3, 4]? True; check 4 to 7, 4 points
742 to [2, 5, 7, 9], 7 in [2, 5, 7, 9]? True; check 7 to 10, 7 points to [5, 10], 10 in [5, 10]? True;
743 check 10 to 9, 10 points to [4, 7, 8, 9], 9 in [4, 7, 8, 9]? True - valid path of length 7
744 B: check 6 to 7, 6 points to [0, 7, 11], 7 in [0, 7, 11]? True; check 7 to 5, 7 points to [5, 10], 5 in [5,
745 10]? True; check 5 to 9, 5 points to [9], 9 in [9]? True - valid path of length 4
746 Valid options: A with length length 7, B with length length 4. Thus the correct shortest option is B
747
748 ** Example 2 **
749 Node 0 points to nodes 7, 8, 11
750 Node 1 points to nodes 3, 4, 5, 10
751 Node 2 points to nodes
752 Node 3 points to nodes 4, 8, 9
753 Node 4 points to nodes 3, 8
754 Node 5 points to nodes 1, 8, 11
755 Node 6 points to nodes 0, 1, 3, 11
756 Node 7 points to nodes 1, 2, 6, 9, 11
757 Node 8 points to nodes 5
758 Node 9 points to nodes 6
759 Node 10 points to nodes 0, 2, 5, 6, 7
760 Node 11 points to nodes 0, 2, 3, 4, 9
761
762 Initial: 7
763 Goal: 11
764
765 Which one is the correct shortest path?
766 A. (7, 6, 10, 0, 8, 5, 11)
767 B. (7, 11)
768 C. (7, 1, 9, 6, 0, 11)
769 Checking each options step by step:
770 A: check 7 to 6, 7 points to [1, 2, 6, 9, 11], 6 in [1, 2, 6, 9, 11]? True; check 6 to 10, 6 points to [0, 1,
771 3, 11], 10 in [0, 1, 3, 11]? False - invalid path
772 B: check 7 to 11, 7 points to [1, 2, 6, 9, 11], 11 in [1, 2, 6, 9, 11]? True - valid path of length 2
773 C: check 7 to 1, 7 points to [1, 2, 6, 9, 11], 1 in [1, 2, 6, 9, 11]? True; check 1 to 9, 1 points to [3, 4,
774 5, 10], 9 in [3, 4, 5, 10]? False - invalid path
775 Valid options: B with length length 2. Thus the correct shortest option is B
776
777 ** Example 3 **
778 Node 0 points to nodes 2, 4, 8, 9
779 Node 1 points to nodes 3, 7, 8, 10
780 Node 2 points to nodes 8, 10
781 Node 3 points to nodes 1, 4, 8, 9, 10
782 Node 4 points to nodes 7, 9
783 Node 5 points to nodes 3, 7

```

```

756 Node 6 points to nodes 5, 11
757 Node 7 points to nodes 2, 10
758 Node 8 points to nodes 0, 3, 7
759 Node 9 points to nodes 10
760 Node 10 points to nodes 0, 4
761 Node 11 points to nodes 3, 7
762
763 Initial: 2
764 Goal: 0
765
766 Which one is the correct shortest path?
767 A. (2, 10, 0)
768 B. (2, 8, 3, 9, 7, 0)
769 C. (2, 10, 0)
770 D. (2, 10, 3, 1, 10, 0)
771 Checking each options step by step:
772 A: check 2 to 10, 2 points to [8, 10], 10 in [8, 10]? True; check 10 to 0, 10 points to [0, 4], 0 in [0, 4]?
773 True - valid path of length 3
774 B: check 2 to 8, 2 points to [8, 10], 8 in [8, 10]? True; check 8 to 3, 8 points to [0, 3, 7], 3 in [0, 3, 7]?
775 True; check 3 to 9, 3 points to [1, 4, 8, 9, 10], 9 in [1, 4, 8, 9, 10]? True; check 9 to 7, 9 points
776 to [10], 7 in [10]? False - invalid path
777 C: check 2 to 10, 2 points to [8, 10], 10 in [8, 10]? True; check 10 to 0, 10 points to [0, 4], 0 in [0, 4]?
778 True - valid path of length 3
779 D: check 2 to 10, 2 points to [8, 10], 10 in [8, 10]? True; check 10 to 3, 10 points to [0, 4], 3 in [0, 4]?
780 False - invalid path
781 Valid options: A with length length 3, C with length length 3. Thus the correct shortest option is C
782
783 ** Current problem **
784 Node 0 points to nodes 2, 3, 4, 8, 10
785 Node 1 points to nodes 2, 6, 7
786 Node 2 points to nodes 4, 5
787 Node 3 points to nodes 2
788 Node 4 points to nodes 2, 5, 6, 9, 11
789 Node 6 points to nodes 2, 3, 4, 8
790 Node 7 points to nodes 9
791 Node 9 points to nodes 10, 11
792 Node 10 points to nodes 0, 8, 9
793 Node 11 points to nodes 1, 4, 9
794
795 Initial: 3
796 Goal: 1
797 Which one is the correct shortest path?
798 A. (3, 2, 4, 11, 1)
799 Remember the graph is directed. Follow the exact same format as the examples and check each options step by
800 step. Begin with 'Checking each options step by step:'
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999

```

Listing 3: Sample prompt used in GRAPH PLANNING for re-ordering the flipped directed graph. The example here is the same for all problems

You will be asked to re-order a directed graph.

```

** Example **
Nodes 8 points to node 0
Nodes 4, 10 points to node 1
Nodes 5 points to node 2
Nodes 1, 9, 11 points to node 3
Nodes 5, 6, 11 points to node 4
Nodes points to node 5
Nodes 2, 9 points to node 6
Nodes 1, 10 points to node 7
Nodes 2, 4, 10 points to node 8
Nodes 10 points to node 9
Nodes 2, 3, 7 points to node 10
Nodes 1, 2 points to node 11
Full procedure:
1. List all directed edges
8 -> 0
4 -> 1
10 -> 1
5 -> 2
1 -> 3
9 -> 3
11 -> 3
5 -> 4
6 -> 4
11 -> 4
2 -> 6
9 -> 6
1 -> 7
10 -> 7
2 -> 8
4 -> 8
10 -> 8
10 -> 9
2 -> 10
3 -> 10
7 -> 10
1 -> 11
2 -> 11

```

```

810 2. Group the edges for each node
811 0 ->
812 1 -> 3, 7, 11
813 2 -> 6, 8, 10, 11
814 3 -> 10
815 4 -> 1, 8
816 5 -> 2, 4
817 6 -> 4
818 7 -> 10
819 9 -> 3, 6
820 10 -> 1, 7, 8, 9
821 11 -> 3, 4
822 3. Convert the edges into the text format
823 Node 1 points to nodes 3, 7, 11
824 Node 2 points to nodes 6, 8, 10, 11
825 Node 3 points to node 10
826 Node 4 points to nodes 1, 8
827 Node 5 points to nodes 2, 4
828 Node 6 points to node 4
829 Node 7 points to node 10
830 Node 8 points to node 0
831 Node 9 points to nodes 3, 6
832 Node 10 points to nodes 1, 7, 8, 9
833 Node 11 points to nodes 3, 4
834
835 ** Current Graph **
836 Nodes 2, 3, 4, 8, 10 points to node 0
837 Nodes 2, 6, 7 points to node 1
838 Nodes 4, 5 points to node 2
839 Nodes 2 points to node 3
840 Nodes 2, 5, 6, 9, 11 points to node 4
841 Nodes 2, 3, 4, 8 points to node 6
842 Nodes 9 points to node 7
843 Nodes 10, 11 points to node 9
844 Nodes 0, 8, 9 points to node 10
845 Nodes 1, 4, 9 points to node 11
846
847 Remember the edges are directed. Please re-order this directed graph with the exact same full procedure as the
848 example. Follow the same format and do not output anything else.

```

Listing 4: Sample prompt used in GRAPH PLANNING for eliciting LLM’s preference over the planning directions in zero-shot.

```

837 You will be given an undirected graph search problem with a few examples. You will decide which search
838 direction is easier to solve for the shortest path from the initial to the goal.
839
840 ** Current problem **
841 Node 0 is connected to nodes 7, 10
842 Node 1 is connected to nodes 4, 5, 11
843 Node 2 is connected to nodes 8
844 Node 3 is connected to nodes 6, 11
845 Node 4 is connected to nodes 1, 6, 9
846 Node 5 is connected to nodes 1, 6, 7, 11
847 Node 6 is connected to nodes 3, 4, 5
848 Node 7 is connected to nodes 0, 5, 9
849 Node 8 is connected to nodes 2
850 Node 9 is connected to nodes 4, 7, 10
851 Node 10 is connected to nodes 0, 9
852 Node 11 is connected to nodes 1, 3, 5
853
854 Initial: 0 Goal: 3
855
856 If there is a bottleneck (nodes with few edges connected) at one end of the graph, then it is easier to solve
857 for the shortest path from that end. Which direction (forward, from the initial, or backward, from the
858 goal) has the bottleneck? Summarize your reasoning in a short paragraph without going through all the
859 nodes, and finish your answer with 'Direction with bottleneck: <forward/backward>'.

```

A.2 ARRAY TRANSFORMATION

Configurations. There are six possible functions used: `repeat`, `cut`, `shift_left`, `shift_right`, `reverse`, and `swap`. Depending on the set of functions used (e.g., `{repeat, cut, shift_left, shift_right}`), we first sample a random array of size 4 as the initial array, sample a random set of three functions, and then apply these functions to the initial array to get the goal one — if `cut` is sampled, we then first invert all the functions, reverse the order, apply the functions, and then reverse the initial and goal arrays. We limit that `repeat` can appear only once among the three to ensure the goal array is not too long.

Prompts. The prompts used for sampling candidate solutions and self-verifying them are shown in [listing 5](#), [listing 6](#), and [listing 7](#). [listing 8](#) shows the prompt used for eliciting LLM’s preference

864 over the planning directions in zero-shot. We use temperature $T = 0$ for self-verifying, eliciting
 865 preference, and the first planning attempt, and $T = 0.5$ for the rest of planning attempts.
 866

867 Listing 5: Header used in the prompt for ARRAY TRANSFORMATION. The functions shown depend
 868 on the set of functions used.

869 A random sequence of three of the below functions transform the initial array into the final array.
 870 Given initial and final, output the sequence of transformations.

```
871 def reverse(x):
872     # reverse the sequence
873     return x[::-1]
874
875 def shift_left(x):
876     # shift the sequence to the left by one
877     return x[1:] + x[:1]
878
879 def shift_right(x):
880     # shift the sequence to the right by one
881     return x[-1] + x[:-1]
882
883 def swap(x):
884     # swap the first and last elements
885     return x[-1:] + x[1:-1] + x[0:1]
886
887 def repeat(x):
888     # repeat the sequence once
889     return x + x
890
891 def cut(x):
892     # cut the sequence in half
893     assert x[:len(x) // 2] == x[len(x) // 2:]
894     return x[:len(x) // 2]
```

887 Listing 6: Sample prompt used in ARRAY TRANSFORMATION for planning candidate solutions,
 888 excluding the header.

```
889 ***** Examples:
890 Initial: [4, 3, 7, 4, 4, 3, 7, 4]
891 Final: [7, 4, 4, 3]
892 Initial to Final Steps:
893   cut: [4, 3, 7, 4]
894   shift_right: [4, 4, 3, 7]
895   shift_right: [7, 4, 4, 3]
896 Functions: [cut, shift_right, shift_right]
897
898 Initial: [8, 0, 0, 2]
899 Final: [0, 2, 8, 0, 0, 2, 8, 0]
900 Initial to Final Steps:
901   shift_left: [0, 0, 2, 8]
902   repeat: [0, 0, 2, 8, 0, 0, 2, 8]
903   shift_left: [0, 2, 8, 0, 0, 2, 8, 0]
904 Functions: [shift_left, repeat, shift_left]
905
906 Initial: [2, 9, 6, 5]
907 Final: [6, 5, 2, 9, 6, 5, 2, 9]
908 Initial to Final Steps:
909   shift_right: [5, 2, 9, 6]
910   shift_right: [6, 5, 2, 9]
911   repeat: [6, 5, 2, 9, 6, 5, 2, 9]
912 Functions: [shift_right, shift_right, repeat]
913
914 ***** Current problem:
915 Initial: [4, 0, 0, 5]
916 Final: [0, 5, 4, 0, 0, 5, 4, 0]
917 Please solve with the exact same format. Do not repeat the problem.
```

907 Listing 7: Sample prompt used in ARRAY TRANSFORMATION for self-verifying the candidate
 908 solutions, excluding the header.
 909

```
910 ***** Examples:
911 Initial: [3, 0, 8, 8, 3, 0, 8, 8]
912 Desired Final: [8, 8, 0, 3]
913 Functions: [shift_left, cut, shift_right]
914 Verify initial to final steps:
915   shift_left: [3, 0, 8, 8, 3, 0, 8, 8][1:] + [3, 0, 8, 8, 3, 0, 8, 8][:-1] -> [0, 8, 8, 3, 0, 8, 8, 3]
916   cut: [0, 8, 8, 3, 0, 8, 8, 3] half -> [0, 8, 8, 3] and [0, 8, 8, 3] equal -> [0, 8, 8, 3]
917   shift_right: [[0, 8, 8, 3][:-1] + [0, 8, 8, 3][:-1]] -> [3, 0, 8, 8]
918   actual final: [3, 0, 8, 8], desired final: [8, 8, 0, 3], does not match
919   Incorrect
920
921 Initial: [1, 5, 7, 2]
922 Desired Final: [7, 2, 1, 5, 7, 2, 1, 5]
923 Functions: [shift_right, shift_right, repeat]
924 Verify initial to final steps:
```

```

918 shift_right: [[1, 5, 7, 2][-1]] + [1, 5, 7, 2][:-1] -> [2, 1, 5, 7]
919 shift_right: [[2, 1, 5, 7][-1]] + [2, 1, 5, 7][:-1] -> [7, 2, 1, 5]
920 repeat: [7, 2, 1, 5] + [7, 2, 1, 5] -> [7, 2, 1, 5, 7, 2, 1, 5]
921 actual final: [7, 2, 1, 5, 7, 2, 1, 5], desired final: [7, 2, 1, 5, 7, 2, 1, 5], match
921 Correct

922 Initial: [5, 5, 0, 2, 5, 5, 3, 2]
923 Desired Final: [4, 2, 0, 5, 5]
923 Functions: [shift_left, cut, shift_right]
924 Verify initial to final steps:
924   shift_left: [5, 5, 0, 2, 5, 5, 3, 2][1:] + [5, 5, 0, 2, 5, 5, 3, 2][:1] -> [5, 0, 2, 5, 5, 3, 2, 5]
925   cut: [5, 0, 2, 5, 5, 3, 2, 5] half -> [5, 0, 2, 5] and [5, 3, 2, 5] not equal -> cut failed
926 Incorrect

927 Initial: [2, 5, 9, 5]
928 Desired Final: [2, 5, 9, 5, 2, 5, 9, 5]
928 Functions: [shift_right, repeat, shift_left]
929 Verify initial to final steps:
929   shift_right: [[2, 5, 9, 5][-1]] + [2, 5, 9, 5][:-1] -> [5, 2, 5, 9]
930   repeat: [5, 2, 5, 9] + [5, 2, 5, 9] -> [5, 2, 5, 9, 5, 2, 5, 9]
931   shift_left: [5, 2, 5, 9, 5, 2, 5, 9][1:] + [5, 2, 5, 9, 5, 2, 5, 9][:1] -> [2, 5, 9, 5, 2, 5, 9, 5]
932   actual final: [2, 5, 9, 5, 2, 5, 9, 5], desired final: [2, 5, 9, 5, 2, 5, 9, 5], match
932   Correct

933 ***** Current problem:
933 Initial: [4, 0, 0, 5]
934 Final: [0, 5, 4, 0, 0, 5, 4, 0]
934 Functions: [shift_left, repeat, shift_left]
935 Please verify initial to final steps with the exactly same format. Do not repeat the problem.

```

Listing 8: Sample prompt used in ARRAY TRANSFORMATION for eliciting LLM’s preference over the planning directions in zero-shot, excluding the header.

```

939 ***** Current problem:
940 Initial: [6, 4, 0, 4]
941 Final: [0, 4, 4, 6]

942 The problem can be solved either in the forward direction (from initial to final), or by flipping the problem
943 first (final becomes initial, initial becomes final) and then solving in the new forward direction.
944 Which direction would you like to solve in? Think about possible bottleneck where fewer search steps are
945 needed. Summarize your reasoning in a short paragraph without going through the intermediate steps and
946 arrays, and finish your answer with 'Direction with bottleneck: <forward/flipped>'.

```

946 A.3 BLOCKSWORLD

948 **Configurations.** We use the problems from the `task_1_plan_generation` task (validity) in the PlanBench benchmark (Valmeekam et al., 2024) without modifying them.

951 **Prompts.** listing 9 shows the original prompt from the benchmark. However, we find LLM often struggle to plan by following the examples in the original prompt, often mistaking the correct initial state. Instead, we use a two-step approach where LLM first summarizes the initial conditions into short format (e.g., “yellow on red; blue; orange”) (listing 10), and then plan — during planning, LLM generates intermediate states in short form to help it reason the next steps (listing 11). listing 12 shows the prompt used for eliciting LLM’s preference over the planning directions in zero-shot. We use temperature $T = 0$ for self-verifying, eliciting preference, and the first planning attempt, and $T = 0.4$ for the rest of planning attempts.

959 Listing 9: Original prompt (header, example, and current problem) from the PlanBench benchmark used in BLOCKSWORLD. Notice that the goal state can be partial.

```

961 You will play with a set of blocks where you need to arrange the blocks into stacks.

962 [POSSIBLE ACTIONS]
963 Pick up a block
964 Unstack a block from on top of another block
965 Put down a block
966 Stack a block on top of another block

967 [RULES]
968 Only pick up or unstack one block at a time.
969 Only pick up or unstack a block if hand is empty.
970 Only pick up a block if the block is on the table and the block is clear. A block is clear if the block has no
971 other blocks on top of it and if the block is not picked up.
972 Only unstack a block from on top of another block if the block being unstacked was really on top of the other
973 block.
974 Only unstack a block from on top of another block if the block being unstacked is clear.
975 Once you pick up or unstack a block, you are holding the block.
976 Only put down a block that you are holding.
977 Only stack a block on top of another block if you are holding the block being stacked.

```

972 Only stack a block on top of another block if the block onto which you are stacking the block is clear.
 973 Once you put down or stack a block, your hand becomes empty.
 974 Once you stack a block on top of a second block, the second block is no longer clear.

975 [[EXAMPLE]]
 976 As initial conditions I have that, the red block is clear, the blue block is clear, the yellow block is clear,
 977 the hand is empty, the blue block is on top of the orange block, the red block is on the table, the
 978 orange block is on the table and the yellow block is on the table.
 979 My goal is to have that the orange block is on top of the blue block.
 980 My plan is as follows:
 981 [PLAN]
 982 unstack the blue block from on top of the orange block
 983 put down the blue block
 984 pick up the orange block
 985 stack the orange block on top of the blue block
 986 [PLAN END]

987 [STATEMENT]
 988 As initial conditions you have that, the red block is clear, the yellow block is clear, the hand is empty, the
 989 red block is on top of the blue block, the yellow block is on top of the orange block, the blue block
 990 is on the table and the orange block is on the table.
 991 Your goal is to have that the orange block is on top of the red block.

987 Listing 10: Sample prompt used in BLOCKSWORLD to summarize the initial and goal state, excluding
 988 the header.

989 [[EXAMPLE]]

990 [STATEMENT]
 991 As initial conditions you have that, the red block is clear, the yellow block is clear, the hand is empty, the
 992 red block is on top of the blue block, the yellow block is on top of the orange block, the blue block
 993 is on the table and the orange block is on the table.
 994 Your goal is to have that the orange block is on top of the red block.

995 First you summarize the init state and goal:

996 [PLAN]
 997 init state (each clause is a stack): red on blue; yellow on orange
 998 goal: orange on red
 999 [PLAN END]

1000 [[CURRENT PROBLEM]]

1001 [STATEMENT]
 1002 As initial conditions you have that, the blue block is clear, the hand is empty, the blue block is on top of
 1003 the orange block, the orange block is on top of the yellow block, the yellow block is on top of the red
 1004 block and the red block is on the table.
 1005 Your goal is to have that the red block is on top of the orange block and the yellow block is on top of the
 1006 red block.

1007 Please follow the format and generate the init state and goal for the current problem. Make sure the stacks
 1008 are combined if they should, e.g., 'red on blue; yellow on red' should be combined as 'yellow on red on
 1009 blue'.

1006 Listing 11: Sample prompt used in BLOCKSWORLD to plan the steps.

1008 [[EXAMPLE]]

1009 [STATEMENT]
 1010 init state (each clause is a stack): orange on red on blue; yellow
 1011 goal: red on blue; yellow on orange.
 1012 Your plan is as follows:

1013 [PLAN]
 1014 unstack the orange block from on top of the red block
 1015 (orange on hand; red on blue; yellow)
 1016 put down the orange block
 1017 (red on blue; yellow; orange)
 1018 pick up the yellow block
 1019 (yellow on hand; red on blue; orange)
 1020 stack the yellow block on top of the orange block
 1021 (red on blue; yellow on orange) Goal satisfied
 1022 [PLAN END]

1023 [[CURRENT PROBLEM]]

1024 [STATEMENT]
 1025 init state (each clause is a stack): yellow on red on orange; blue
 1026 goal: blue on orange on yellow on red

1027 Please follow the format and generate your plan for the current problem. Start with [PLAN]

1024 Listing 12: Sample prompt used in BLOCKSWORLD to elicit LLM's preference over the planning
 1025 directions in zero-shot, excluding the header.

1026

[STATEMENT]

1027

As initial conditions you have that, the red block is clear, the yellow block is clear, the hand is empty, the red block is on top of the blue block, the yellow block is on top of the orange block, the blue block is on the table and the orange block is on the table.

1028

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

1029

1030

The problem can be solved either in the forward direction (from initial condition to goal), or by flipping the problem first (goal becomes initial, initial becomes goal) and then solving in the new forward direction. Which direction would you like to solve in? Think about possible bottleneck where fewer search steps are needed. Summarize your reasoning in a short paragraph without going through the intermediate steps and states, and finish your answer with 'Direction with bottleneck: <forward/flipped >'.

1031

1032

1033

1034

1035

1036

B APPENDIX - ADDITIONAL RESULTS

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079



Figure 12: Success rates achieved by forward and backward planning in ARRAY TRANSFORMATION vs. difference between forward and backward BFS computations. In general LLM plans better in the direction of fewer computations needed, but the forward direction outperforms backward. In the last set of experiments, we find Back works well by recognizing arrays that should be repeated or cut, and thus there is no visible difference between Fwd and Back.