Large Language Models Still Can't Plan (A Benchmark for LLMs on Planning and Reasoning about Change)

Anonymous Author(s)
Affiliation
Address
email

Abstract

Recent advances in large language models (LLMs) have transformed the field of natural language processing (NLP). From GPT-3 to PaLM, the state-of-the-art performance on natural language tasks is being pushed forward with every new large language model. Along with natural language abilities, there has been a significant interest in understanding whether such models exhibit reasoning capabilities with the use of reasoning benchmarks. However, even though results are seemingly positive, these benchmarks prove to be simplistic in nature and the performance of LLMs on these benchmarks cannot be used as evidence to support, many a times outlandish, claims being made about LLMs' reasoning capabilities. Further, these only represent a very limited set of simple reasoning tasks and we need to look at more sophisticated reasoning problems if we are to measure the true limits of such LLM-based systems. Motivated by this, we propose an extensible assessment framework to test the capabilities of LLMs on reasoning about actions and change, a central aspect of human intelligence. We provide multiple test cases that are more involved than any of the previously established benchmarks and each test case evaluates a different aspect of reasoning about actions and change. Results on GPT-3 (davinci), Instruct-GPT3 (text-davinci-002) and BLOOM (176B), showcase subpar performance on such reasoning tasks.

1 Introduction

It would be no exaggeration to say that transformer-based large language models (LLMs) have revolutionized the field of natural language processing (NLP). Kicked off by the advances presented by the GPT-x models developed by OpenAI [23], these types of language models currently provide state-of-the-art performance in many of the standard NLP tasks. The latest version of the system, GPT-3 [4] uses about 175 billion parameters and was trained over an extremely large natural language training corpus, consisting of, among other things, excerpts from Wikipedia. Triggered by GPT-3, a plethora of other large language models, which are different variants of the transformer architecture [31], have been developed. Some of the most powerful models are PaLM [5], GLaM [7], Megatron-Turing NLG [26], Meta-OPT [35], Gopher [24], LaMDA [30] and Chinchilla [10]. PaLM currently provides state-of-the-art performance in NLP tasks such as natural language translation, predicting long-range text dependencies and even translation to structured representations [5, 19].

Although LLMs were originally developed mostly to do language completion tasks, with no guarantees about the completion beyond its coherence, there have been increasing claims that they can do other tasks including explaining jokes [14]. Of particular interest is the thread of efforts that aim to evaluate (and showcase) the LLM's ability to do reasoning tasks. For example, there are many claims centered around the fact that GPT-3 may possess some form of reasoning ability [16]. Such sources

generally assume that because the model learned from large amounts of real-world text, it may have acquired some approximation of simple reasoning. This sparked interest in evaluating the large language models on various reasoning tasks including common-sense reasoning [29, 25, 8], logical reasoning [27], and even ethical reasoning [13]. The macro-tenor of the drumbeat of these works has been suggesting that LLM's are indeed capable of doing such kinds of reasoning [15, 33, 5].

Thus, in this paper, we want to look at the ability of large language models to do reasoning about actions and change involving common-sense planning tasks. We develop a suite of benchmarks, based on the kinds of domains employed in the International Planning Competition, to test these capabilities. Our focus on this specific task is spurred by not only the fact that reasoning about actions and change is a core aspect of human intelligence, but also that it is required for many of the tasks considered as potential applications of LLMs including automatic code generation, moral and even deontological reasoning [9].

We are not the first to point out the need to perform such analyses of the reasoning capabilities of GPT-3 like LLMs. For example, [18] performed an analysis of GPT-3's reasoning capabilities on some example tasks, including different commonsense reasoning tasks varying from biological reasoning to arithmetic reasoning. However, the goal of this paper is fundamentally distinct from these earlier works in multiple ways. Firstly, we are not merely trying to point out a few example cases where GPT-3 fails but rather help establish an assessment framework for evaluating these systems' capabilities to perform reasoning about actions and change. While this paper reports the results of testing vanilla GPT-3, Instruct-GPT3 [20] and BLOOM [1], one could in the future use this framework to analyse other LLMs that may be fine-tuned for such tasks. Secondly, through this framework, we are also trying to eliminate the subjective aspect of analysis that forms the core part of many of these earlier efforts. Instead, we automate and perform the analyses in a mechanistic way by leveraging automated planning models and tools to generate the queries and validate the system's answers. Finally, we propose a curriculum for evaluating reasoning about actions and change, wherein we identify a set of related but distinct reasoning tasks, that are central for an agent to successfully perform reasoning about actions and change and introduce a framework for developing code which is meant to auto-generate the possible queries for each. The conclusion from our evaluation on GPT-3 and BLOOM is that LLMs are still far from achieving acceptable performance on common planning/reasoning tasks which pose no issues for humans to do. In particular, our benchmarks do not require complex long-chain combinatorial reasoning and can be easily understood by lay people. Therefore, our contributions are twofold, we

1. Provide an extensible suite of benchmarks for evaluating the planning/reasoning capabilities of LLMs
2. Show that GPT-3 and BLOOM have dismal performance on these benchmarks.

Additionally, we perform a preliminary user study that establishes a human baseline for coming up with plans in a simple planning domain to compare the LLMs' performance. Our hope is that these benchmarks spur other researchers to evaluate out-of-the-box or fined-tuned versions of other LLMs on these planning tasks.

2 Current Reasoning Benchmarks for LLMs

To the best of our knowledge, current available benchmarks are insufficient to make substantive claims about LLMs' ability to reason. The authors of [24] had pointed out that even language models at the scale of 100B or more parameters had struggled on reasoning tasks that require slow and multi-step reasoning, but the more recent works [15, 33, 32] have showcased and claimed that LLMs can do various reasoning tasks with a decent performance. These claims have been based on several reasoning benchmarks (consisting of arithmetic, common-sense and symbolic reasoning tasks), where the reasoning tasks are relatively simple and require shallow reasoning. Datasets like GSM8K [6], AQUA [17] and SVAMP [21] have simple math word problems which are used for evaluating arithmetic reasoning while datasets like CommonsenseQA [29] and StrategyQA [8], which have generic multiple choice and binary yes/no questions respectively, are used for evaluating common sense reasoning. There are several logical reasoning tasks in BIG-BENCH [27] and there are two symbolic reasoning tasks, Last Letter Concatenation and Coin Flip [33], on which LLMs have been evaluated. However, these tasks are simple in nature and do not provide insight into the reasoning capabilities of LLMs. As LLMs have been able to perform well on such tasks, there has

a) A carnival snack booth made $50 selling popcorn each day. It made three times as much selling cotton candy. For a 5-day activity, the booth has to pay $30 rent and $75 for the cost of the ingredients. How much did the booth earn for 5 days after paying the rent and the cost of ingredients?

b) Kaleb was collecting cans for recycling . On Saturday he filled number0 bags up and on Sunday he filled number1 more bags . If each bag had number2 cans in it, how many cans did he pick up total ?

c) Two friends plan to walk along a 43-km trail, starting at opposite ends of the trail at the same time. If Friend P's rate is 15% faster than Friend Q's, how many kilometers will Friend P have walked when they pass each other?
A)21, B)21.5, C)22, D)22.5, E)23

d) What would someone wear to protect themselves from a cannon?
A. Body armor, B. tank, C. hat, D. …

e) Is it normal to find parsley in multiple sections of the grocery store?
(Yes/No)

f) Alice, Bob, Claire, Dave, Eve, Fred, and Gertrude are playing a game. At the start of the game, they are each holding a ball: Alice has a green ball, Bob has a white ball, Claire has a yellow ball, Dave has a pink ball, Eve has a orange ball, Fred has a black ball, and Gertrude has a brown ball. \n\nAs the game progresses, pairs of players trade balls. First, Bob and Gertrude swap balls. Then, Fred and Claire swap balls. Then, Dave and Gertrude swap balls. Then, Bob and Gertrude swap balls. Then, Alice and Claire swap balls. Then, Gertrude and Claire swap balls. Finally, Eve and Claire swap balls. At the end of the game, Bob has the…

g) Yesterday was April 30, 2021. What is the date today in MM/DD/YYYY?

h) Take the last letters of the words in “Lady Gaga” and concatenate them.

i) A coin is heads up. Maybelle flips the coin. Shalonda does not flip the coin. Is the coin still heads up?

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

Pick up a block
Unstack a block from on top of another block
Put down a block
Stack a block on top of another block

I have the following restrictions on my actions:
I can only pick up or unstack one block at a time.
I can only pick up or unstack a block if my hand is empty.
I can 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 other blocks on top of it and if the block is not picked up.
I can only unstack a block from on top of another block if the block I am unstacking was really on top of the other block.
I can only unstack a block from on top of another block if the block I am unstacking is clear.
Once I pick up or unstack a block, I am holding the block.
I can only put down a block that I am holding.
I can only stack a block on top of another block if I am holding the block being stacked.
I can only stack a block on top of another block if the block onto which I am stacking the block is clear.
Once I put down or stack a block, my hand becomes empty.

[STATEMENT]
As initial conditions I have that, the red block is clear, the blue block is clear, the yellow block is clear, the hand is empty, the blue block is on top of the orange block, the red block is on the table, the orange block is on the table and the yellow block is on the table.
My goal is to have that the orange block is on top of the blue block.

My plan is as follows:

[PLAN]
unstack the blue block from on top of the orange block
put down the blue block
pick up the orange block
stack the orange block on top of the blue block
[PLAN END]

[STATEMENT]
As initial conditions I 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.
My goal is to have that the orange block is on top of the red block.

My plan is as follows:

[PLAN]

Figure 1: Example prompts from existing benchmarks (left) compared to an average-length prompt from one of our test cases (right). (a) GSM8k [6], (b) SvAMP [21], (c) AQuA [17], (d) CommonSenseQA [29], (e) StrategyQA [8], (f) Tracking Shuffled Objects [27], (g) Date Understanding [27], (h) Last Letter Concatenation [33], (i) Coin Flip [33], (j) Goal directed reasoning on blocksworld.

been a lot more triumphalism about their reasoning capabilities, which is currently being echoed in the community.

LLMs can also be used for heuristic guidance with no optimality or completeness guarantees, as long as the guidance is used by a sound planner underneath. For example, in Say-Can [3], LLMs have been used as scoring models, which can be seen as providing planning heuristics, for the actions that the embodied robot can execute. In our work, we are looking at the inherent emergent planning capabilities of the LLMs, if any. The assessment framework we propose in this paper consists of multiple tasks each of which evaluate a certain aspect of reasoning about actions and change. Even though we use a simple common-sense planning domain, the prompts for our tasks are relatively larger and more complex than any of the prompts in the current benchmarks (as shown in Figure 1). This is due to two main reasons; first, the prompts for our tasks are showcased as few-shot examples, whereas the prompts in the current benchmarks are zero-shot. Our prompts provide an instance and an example completion and then ask for a completion on a new instance. Second, we use the domain description to explicitly constrain the possible actions. In many everyday scenarios, we are often asked to take into consideration unforeseen limitations and constraints. Our explicit domain description allows us to introduce such challenges, and forces the LLMs to go beyond merely repeating possible information about the domain they may have come across in its training data. The ability to be conditional on the prompt is critical for the general systems to be customized for the specific domain of interest.

Our results, contrary to the current claims, show that even simple common-sense planning tasks (which are, anecdotally, easy for humans) are far beyond the current capabilities of LLMs. This is actually not surprising given that LLMs only provide the most likely text completion for a given prompt with no real guarantees on reasoning metrics. There have been works that tried to evaluate whether LLMs can generate plans in common-sense domains [12] but in those works they had

evaluated on the world knowledge that LLMs already possess and do not provide any information about the domain in the prompt. Before we delve into the details of our framework, we will first establish the required background to get a better understanding of the kind of reasoning capability that we are focusing on.

3 Background

Given that we are interested in investigating the basic reasoning about actions and change problem, we want to look at the most fundamental planning formalism first, namely the goal-directed deterministic planning problem. Colloquially referred to as *classical planning problem*, these problem classes can be mathematically represented by using the tuple $\mathcal{P} = \langle \mathcal{D}, \mathcal{I}, \mathcal{G} \rangle$. $\mathcal{D}$ is referred to as the problem domain, I is the initial state and G is the goal specification. The state-space for the planning problem is defined by the possible truth assignment over the predicates. The domain is again defined by the tuple $\mathcal{D} = \langle \mathcal{F}, \mathcal{O} \rangle$. $\mathcal{F}$ corresponds to the set of fluents, i.e., the state variable used to define the state space and each fluent corresponds to a predicate with some arity, and $\mathcal{A}$ correspond to the set of actions that can be performed as part of the planning problem. Each action $a_i[\mathcal{V}] \in \mathcal{A}$ (where a_i is the operator label and $\mathcal{V}$ is the variable used by the operator and each variable could be mapped to an object), can be further defined by two components, the precondition $prec[\mathcal{V}]$ which describes *when* an action can be executed and the effects $eff[\mathcal{V}]$ which defines *what happens* when an action is executed. We will assume that $prec[\mathcal{V}]$ consists of a set of predicates defined over the variables $\mathcal{V}$. An action is assumed to be executable only if its preconditions are met, i.e, the predicates in the precondition hold in the given state. The effect set $eff[\mathcal{V}]$ is further defined by the tuple $\langle add[\mathcal{V}], del[\mathcal{V}] \rangle$, where $add[\mathcal{V}]$ or add effects is the set of predicates that will be set true by the action and $del[\mathcal{V}]$ or delete effects is the set of predicates that will be set false by the action. An action is said to be grounded if we replace each of the variables with an object, else it is referred to as a lifted domain model (we use a similar convention to differentiate between lifted and grounded predicates).

Below we have provided a snippet of an action description from a popular benchmark problem called Blocksworld for action corresponding to picking up a block.

```
(:action pickup
  :parameters (?ob)
  :precondition (and (clear ?ob) (on-table ?ob) (arm-empty))
  :effect (and (holding ?ob) (not (clear ?ob)) (not (on-table ?ob))
               (not (arm-empty))))
```

The parameter line provides the possible variables, in this case *?ob*, which can stand for possible blocks. The precondition says that you can only pick up a block if it is clear (i.e. predicate *(clear ?ob)* is true for the block), the block is on the table and the arm is empty. The effects tell you that after you execute the action, the predicate *(holding ?ob)* becomes true and the block will no longer be considered *clear*, and *on-table*. Finally, the arm will no longer be considered *empty*. A solution to a planning problem is called a plan, and corresponds to a sequence of actions that once executed in the initial state would lead to a state where the goal specification is true. The actions may additionally be associated with cost, in these cases, one could also talk about optimal plans, i.e., a plan π is called an optimal one if no plan exists that is less costly than π.

The above description presents one of the simpler classes of planning models and can be extended in multiple ways including allowing for object typing (including type hierarchy), more complex forms of preconditions and conditional effects, not to mention supporting richer classes of planning formalisms.

4 Assessment Architecture

Our basic test framework consists of two categories of components, the domain-independent ones, provided as part of the framework, and the domain-dependent components which need to be developed for each new domain we test.

Domain-independent component: the domain-independent component is built around a planner and a plan verification component that takes various planning problems and crafts test instances corresponding to various curriculum items. This component provides the mechanism to verify the

solutions generated by the LLM. The current method is going to operate almost exclusively on symbolic models (specifically ones specified using PDDL [2]) and other structured inputs compatible with such representations. The domain-dependent component would be responsible for translating outputs generated by the LLM into forms that can be used by the system.

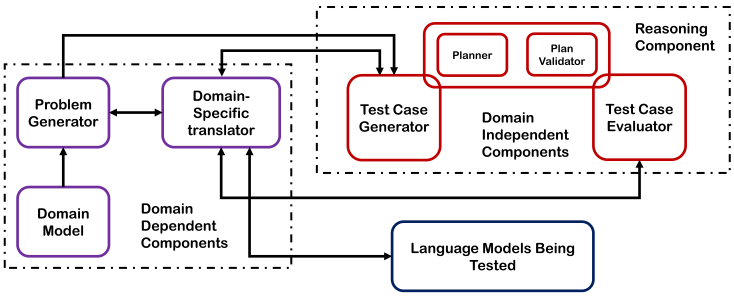

Figure 2: The diagrammatic overview of the overall test framework. Broadly, our system consists of a domain-specific component that allows the generation of various instances of the specific PDDL planning problems and the translation of the PDDL information to text and back. The domain-independent component is responsible for generating the various specific test instances that will be fed into the LLM and to verify the output generated by the LLM.

Domain-dependent component: the domain-dependent component consists of three main parts. A lifted domain model file, that describes the various actions that may be available to solve any given planning problem, the various predicates that could be used to describe the various relationships over the objects that may be present at a given problem instance of the domain, and the various types of objects that may be part of the given problem. The domain model is lifted because it does not refer to the actual objects that may be part of the problem, but instead, the actions are defined independently of the exact objects it may influence.

Problem generator: a planning problem consists of a description of the set of specific objects that are part of the specific planning problem, the initial state (described in terms of the truth values of the various predicates), and a goal description. A valid solution consists of a sequence of actions that can drive the system state to a state that satisfies the goal condition. The role of the problem generator is therefore to generate random problem instances consisting of various objects, initial states, and goals. These problems become the basis of generating the various test cases that we will be using throughout the framework. Any distributional requirements we hope to use in the tests could be built into this problem generator.

Translator: the translator converts the symbolic model information to natural language text and *vice versa* . In particular, we are interested in developing a mechanism to translate state information and plans into natural language. For the current testbed (described below), we developed a template-based mechanism to achieve this. In particular, we provide a natural language template for each predicate and each action, and we form texts of states and plans by concatenating these individual strings. In terms of parsing natural language text back into structured forms, the particular task we are interested in is converting plans generated by the LLM back into plan forms that can be used by plan validator tools like [11]. Since we use our prompts to shape the LLM's output, we require each action in the plan to be listed on a different line. Then, we can parse the exact action and arguments of the action by either using template-based matching or by assuming that the verb in the sentence corresponds to the action and each noun corresponds to an object which forms a parameter of the action (then mapping it to a possible action). The domain-independent component is responsible for generating the content for the various prompts that would be generated as part of the different test cases and for validating the output generated by the LLM.

As discussed earlier, the component primarily works on formal representations of the problems, so it relies on the translator component to convert any information it generates to natural language or to convert natural language information back to formal representations. For each test case, we mainly rely on a domain-independent planner and a plan validator to generate the relevant information or to validate the output provided by the LLM. In each case, there is a test-case-specific component that uses the problems provided by the problem generator component to craft specific test-case content. In

the next section, we go over each test case and the specific technique we use to generate the contents for the test case.

5 Current Curriculum for Testing

Each test case is meant to evaluate a central reasoning about actions and change capability and is tested in the context of a common sense planning domain. Each test case makes use of the few shot query setting of LLM where the LLM is provided a few sample answers to the specific reasoning ability being tested and is asked to respond to a new instance. The exact form of the prompt will depend on the specific test cases, but every instance will start with a description of the lifted planning domain that describes what actions can be executed, their preconditions and their effects.

The current set of test cases includes the following cases

1. Plan Generation - Can the LLM come up with valid plans that will achieve a specific goal?
2. Cost Optimal Planning - Can the LLM come up with plans that are optimal to achieve a specific goal?
3. Reasoning about plan execution - Can the LLM reason about what happens when a plan is executed?
4. Robustness to goal reformulation - Can the LLM recognize the same goal when specified in different ways?
5. Ability to reuse plans - Can the LLM recognize scenarios where it can reuse part or the whole of the original plan to achieve the new goal?
6. Replanning - Can the LLM replan for cases where an unexpected change is reported?
7. Plan Generalization - Can the LLM take specific plans, extract underlying procedural patterns and apply them to a new instance?

Out of the seven test cases, the first two test cases correspond to actual planning problems (i.e. goal-directed planning and cost-optimal planning) and the rest correspond to simpler auxiliary tasks related to reasoning about action and change.

Domain: currently, we ground the test cases in a simple common-sense planning domain, Blocksworld. Blocksworld problems generally consist of a set of blocks, for making it closer to a common sense domain identified with unique colors, placed either on a table or on top of other blocks and the goal is to arrange some of these blocks in a stack in a particular order. The general expectation here would be that one can pick up a block if it is clear, i.e., there are no other blocks on top of that block and you can only stack a block on top of another block if it is clear. The choice of this particular domain is motivated by both the fact that this is a simple common sense domain and is a very popular domain in planning literature, that has a long history of being used to demonstrate various planning challenges. The domain description is included in the beginning of every prompt (see Appendix A.1.1 for the domain description). In the rest of the section, we discuss the structure of the prompt for each of the test cases. We point the reader to Appendix A.2 for an example prompt and the corresponding completion generated by GPT-3 for each of the test cases.

5.1 Plan Generation

Following the lifted domain description, the prompt consists of a few instances of planning problem descriptions (consisting of a description of the initial state, the goal) and the corresponding plan (which ends with a tag, henceforth referred to as the plan-end tag, that denotes the end of the plan) and finally, we end the prompt with a planning problem description. The text generated by the LLM until the plan-end tag is used as a potential candidate for extracting the plan. If the plan-end tag is missing or if the plan cannot be extracted then we ignore that particular instance in our evaluation.

5.2 Optimal Planning

The prompt is quite similar to the one used in the earlier test case with a few changes. We modify the lifted domain description by including a statement that associates a cost with each action (see Appendix A.1.2). To make the concept of action cost better fit into common sense domains, we can

map the cost to more common concepts like the time taken for executing the action or the amount of money that needs to be spent to execute an action. In the case of each problem description, before the plan is presented we need to explicitly mention that the plan is trying to minimize cost (which depending on the scenario might correspond to saying that the plan takes the least amount of time or the plan correspond to the cheapest plan). The result generated by the LLM is evaluated similarly to the previous query, but in addition to checking if the plan is valid, we also check if the cost of the plan corresponds to the optimal plan cost.

5.3 Reasoning about plan execution

Here the objective is not to check whether the LLM can come up with plans, but rather if they can predict the outcome of executing an action. The prompt here again starts with the domain description, but instead of providing planning problems and plans, we provide a state, an action sequence and then questions about the state that would result from executing that action sequence in the provided state. Finally the prompt ends with a new state, a new action sequence, and a question about the resulting state. The LLM is expected to come up with an answer, which is checked by applying a plan executor that will try to identify what state will result from the execution of the current action sequence on the state.

5.4 Robustness to goal formulation

In this test case, we will see if the LLM can recognize goals it has seen before if they are slightly modified. Here the prompt remains the same as the one used for goal-directed reasoning. However, all the example problems have the same initial state, and the last problem provided has not only the same initial state but also the same goal. Here the goal may be obfuscated in a few ways, for example, the goal facts may be reordered or one might only include a subset of the original goal specification (meaning the same plan would still work). We can again use the same evaluation technique as the goal-directed reasoning test case to validate the output.

5.5 Ability to Reuse Plans

In this test case, we are interested in seeing if the LLM can reuse plans or parts of plans that it has seen before. The prompt is again the same as the goal-directed reasoning, but the prompt ends with a problem that can be solved by a prefix of a previously seen plan. We again keep the initial state the same across the example problems shown. The evaluation remains the same as the goal-directed reasoning test case.

5.6 Replanning

Replanning corresponds to the problem where there may be an unexpected event that occurs while executing a plan and the system needs to come up with a new plan in response to the event. Here, we focus on the ability of the LLM to replan when unexpected changes are reported. The prompt here starts with a domain description, then a set of instances where an unexpected event occurred during execution, and a new plan in response to the event. In each instance, a planning problem and a corresponding plan are provided at the beginning, the execution of the plan is described and then an unexpected event is noted (event corresponds to some facts unexpectedly turning true or false) and then a new plan from the changed state is presented. The prompt ends with a new case where the plan after replanning is left out and the LLM is expected to complete. The evaluation involves checking whether the new plan is valid from the changed state. The LLM output is evaluated to be true if the new plan it generates achieves the goals from the unexpectedly changed state.

For the Blocksworld domain, we constrain the unexpected event to be of a specific type. We execute a random prefix of the plan which ensures that some block is held at the end of that prefix. We then change the resulting state by stacking the held block onto another random block which is clear and make the hand empty. This change is reported and the LLM is asked to replan from the changed state.

5.7 Plan Generalization

In this test case, we want to evaluate whether LLMs can recognize the underlying pattern in the plans provided in the prompt and reuse it for a new planning problem. The prompt is the same as

Task	Instances correct		
	GPT-3	Instruct-GPT3	BLOOM
Plan Generation We showcase an instance and the respective plan as an example and prompt the machine with a new instance.	3/500 (0.6%)	25/500 (5%)	1/200 (0.5%)
Optimal Planning We showcase an instance, the respective optimal plan and the associated cost as an example and prompt the machine with a new instance.	1/500 (0.2%)	16/500 (3.2%)	0/100 (0%)
Replanning We showcase an instance, the respective plan and present an unexpected change of the state. We then also present a new plan from the changed state. Finally, for a new instance we repeat the same except we ask the machine for the new plan.	28/500 (5.6%)	24/500 (4.8%)	3/100 (3%)
Plan Generalization We showcase an instance and the respective plan as an example and prompt the machine with a new instance. The plans for both the instances can be generated by a fixed program containing loops and conditionals.	33/500 (6.6%)	49/500 (9.8%)	11/100 (11%)
Plan Reuse We showcase an instance and the respective plan as an example and prompt the machine with an instance which requires only a certain prefix of the plan provided in the example.	0/500 (0%)	72/500 (14.4%)	0/100 (0%)
Robustness to Goal Reformulation (Shuffling goal predicates) We showcase an instance and the respective plan as an example and prompt the machine with the same instance but shuffle the ordering of the goals.	387/500 (77.4%)	384/500 (76.8%)	21/100 (21%)
Robustness to Goal Reformulation (Full $\rightarrow$ Partial) We showcase an instance with a fully specified goal state and the respective plan as an example and prompt the machine with the same instance but provide a partially specified goal state.	346/500 (69.2%)	380/500 (76%)	9/100 (9%)
Robustness to Goal Reformulation (Partial $\rightarrow$ Full) We showcase an instance with a partially specified goal state and the respective plan as an example and prompt the machine with the same instance but provide a fully specified goal state.	110/500 (22%)	301/500 (60.2%)	5/100 (5%)

Table 1: LLM Assessment Suite Results on vanilla GPT-3 (davinci), Instruct-GPT3 (text-davinci-002) and BLOOM (176B model)

the goal-directed reasoning case, except that all plans were generated by a fixed program. Here the program may contain loops or conditional statements, but can only solve certain types of problems, that is, the initial state and goals meet certain conditions. Such programs can be thought of as a direct generalization of line plans that we have considered in the rest of the paper [28]. Execution of this program for a specific planning problem generates a sequence of actions. In this case, we will provide some example traces generated from the program and ask LLM to come up with a plan for a new problem that could be solved by it. The evaluation again would be to take the generated plan and see if it is valid for the given problem.

6 Evaluation

6.1 Results on the Blocksworld domain

Our evaluation here primarily focuses on two Large Language Models, GPT-3 and BLOOM. In particular, we evaluated the test framework on the Blocksworld domain. In Table 1, we have presented the results of vanilla GPT-3 (Davinci), Instruct-GPT3 (text-davinci-002), and BLOOM on six of the test cases. For vanilla GPT-3 and Instruct-GPT3, we had tested on 500 instances while with BLOOM, we had tested on 200 instances for plan generation and 100 instances for the rest of the test cases. The experiments with GPT-3 (both the vanilla and instruct versions) took $\sim$30 minutes for each test case (500 instances) while BLOOM took $\sim$36 hours (every 100 instances). The best results (within each model) were for the auxiliary goal reformulation test cases. For these three cases, all that was required for the LLM was to repeat the same plan as the one shown in the example. Even then, vanilla GPT-3 and Instruct-GPT-3 failed to do that for some of the instances in the first two cases

and the majority of the instances in the third case. BLOOM, on the other hand, was poor in all three cases. Coming to the two test cases that correspond to actual planning problems (plan generation and optimal planning), all three models performed poorly with Instruct-GPT3 performing better than the other two. Overall, the performance of these LLMs on our benchmark shows that, as of right now, LLMs are pretty ineffective in reasoning about actions and change. It would be interesting to see the utility of fine-tuning or even that of prompt engineering applied to our experiments.

6.2 Human Baseline for the Blocksworld domain

We have previously mentioned that planning tasks on the blocksworld domain are anecdotally simple enough for humans to perform. To establish this and come up with a baseline to compare LLMs performance, we conducted a preliminary study where we asked 50 participants to come up with a plan for a blocksworld instance picked at random, from the set of 500 instances that we used for the evaluation of LLMs. We presented the same domain description as we did for the LLMs (see Appendix A.1.1) and then primed them with an example instance. Further, we provided them with an interface where they had two phases of interaction. In the first phase, they could write up plans by themselves for the given instance and then in the second phase, translate them (by picking the closest action from a list of grounded actions). The translated plans were used in the back-end and were evaluated in an automated fashion[1]. They went through this procedure first for an example instance (where they were provided with a glimpse of the example solution before using the interface) and then for the actual instance. We provided them with a bonus if they come up with a valid plan.

Out of the 50 participants, 39 of them (78%) came up with a valid plan. Along with validity, we also tested the optimality of their plans even though they were not required to come up with an optimal plan. Out of the 39 participants, 35 (89.7%) participants came up with an optimal plan. These initial results show that the blocksworld domain is a simple enough domain where most humans are able to come up with plans (which are also optimal) while LLMs, on the other hand, showcase subpar performance. A more extensive human subject study is underway to establish a more concrete human baseline for performance comparisons.

7 Conclusion and Future Work

In this paper, we looked at a reasoning assessment suite for large language models (LLMs) that consists of various test cases each evaluating a central aspect of reasoning about actions and change. Our results show that even in simple common-sense planning domains where humans could easily come up with plans, LLMs like GPT-3 and BLOOM seem to display a dismal performance. However, we do not claim that no LLM system could potentially ever perform effective reasoning about actions and change. Our goal is to establish an extensible benchmark where researchers can evaluate current and future large language models. Our assessment suite is still in progress and we look to improve it in multiple ways in the future. Firstly, we plan to include a modified version of the reasoning about plan execution task to ask questions that require a more descriptive answer and provide automated validations for the answers. Secondly, we also plan to look at how fine-tuning would impact the performance of the LLMs on our benchmarks, although as [34] points out that LLMs tend to focus on correlational features even with extensive training, we are not expecting a major improvement. Finally, we plan to add an evaluation metric that also considers partial correctness of plans. This benchmark can also be extended to other domains, either to other common-sense domains (like Virtual Home [22]) or to specialized ones. It would also be interesting to probe these models to come up with explanations for the generated plans and see if these make sense to humans in the loop. In conclusion, we hope that this benchmark encourages other researchers to test the capabilities of their systems across different LLM models [5, 7, 26, 35, 24, 30, 10] and even those that are finetuned for such tasks.

References

[1] BigScience Large Open-science Open-access Multilingual Language Model. `https://huggingface.co/bigscience/bloom`.

[1] We had also manually evaluated the plans that they wrote in case they made a mistake during the translation phase

[2] Constructions Aeronautiques, Adele Howe, Craig Knoblock, ISI Drew McDermott, Ashwin Ram, Manuela Veloso, Daniel Weld, David Wilkins SRI, Anthony Barrett, Dave Christianson, et al. PDDL| The Planning Domain Definition Language. *Technical Report, Tech. Rep.*, 1998.

[3] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.

[4] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.

[5] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.

[6] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

[7] Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. GLaM: Efficient Scaling of Language Models with Mixture-of-Experts. *arXiv preprint arXiv:2112.06905*, 2021.

[8] Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of the Association for Computational Linguistics*, 9:346–361, 2021.

[9] Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch Critch, Jerry Li Li, Dawn Song, and Jacob Steinhardt. Aligning ai with shared human values. In *International Conference on Learning Representations*, 2021.

[10] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.

[11] Richard Howey, Derek Long, and Maria Fox. VAL: Automatic plan validation, continuous effects and mixed initiative planning using PDDL. In *16th IEEE International Conference on Tools with Artificial Intelligence*, pages 294–301. IEEE, 2004.

[12] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. *arXiv preprint arXiv:2201.07207*, 2022.

[13] Liwei Jiang, Jena D. Hwang, Chandrasekhar Bhagavatula, Ronan Le Bras, Maxwell Forbes, Jon Borchardt, Jenny Liang, Oren Etzioni, Maarten Sap, and Yejin Choi. Delphi: Towards Machine Ethics and Norms. *ArXiv*, abs/2110.07574, 2021.

[14] Subbarao Kambhampati. Language Imitation Games and the Arrival of Broad and Shallow AI. https://cacm.acm.org/blogs/blog-cacm/256068-language-imitation-games-and-the-arrival-of-broad-and-shallow-ai/fulltext, October 2021.

[15] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large Language Models are Zero-Shot Reasoners. *arXiv preprint arXiv:2205.11916*, 2022.

[16] Belinda Z Li, Maxwell Nye, and Jacob Andreas. Implicit representations of meaning in neural language models. *arXiv preprint arXiv:2106.00737*, 2021.

[17] Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. Program induction by rationale generation: Learning to solve and explain algebraic word problems. *arXiv preprint arXiv:1705.04146*, 2017.

[18] Gary Marcus and Ernest Davis. Experiments testing GPT-3's ability at commonsense reasoning: results. https://cs.nyu.edu/ davise/papers/GPT3CompleteTests.html.

[19] Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. GPT3-to-plan: Extracting plans from text using GPT-3. *arXiv preprint arXiv:2106.07131*, 2021.

[20] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*, 2022.

[21] Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are NLP Models really able to Solve Simple Math Word Problems? *arXiv preprint arXiv:2103.07191*, 2021.

[22] Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. Virtualhome: Simulating household activities via programs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8494–8502, 2018.

[23] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. 2018.

[24] Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*, 2021.

[25] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8732–8740, 2020.

[26] Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. Using deepspeed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990*, 2022.

[27] Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*, 2022.

[28] Siddharth Srivastava, Shlomo Zilberstein, Neil Immerman, and Hector Geffner. Qualitative numeric planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 25, pages 1010–1016, 2011.

[29] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937*, 2018.

[30] Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. LaMDA: Language Models for Dialog Applications. *arXiv preprint arXiv:2201.08239*, 2022.

[31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[32] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.

[33] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*, 2022.

[34] Honghua Zhang, Liunian Harold Li, Tao Meng, Kai-Wei Chang, and Guy Van den Broeck. On the paradox of learning to reason from data. *arXiv preprint arXiv:2205.11502*, 2022.

[35] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.

A Appendix

A.1 Domain description included in the prompts

A.1.1 Domain included in the prompt

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

Pick up a block
Unstack a block from on top of another block
Put down a block
Stack a block on top of another block

I have the following restrictions on my actions:
I can only pick up or unstack one block at a time.
I can only pick up or unstack a block if my hand is empty.
I can 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 other blocks on top
     of it and if the block is not picked up.
I can only unstack a block from on top of another block if the block I
     am unstacking was really on top of the other block.
I can only unstack a block from on top of another block if the block I
     am unstacking is clear.
Once I pick up or unstack a block, I am holding the block.
I can only put down a block that I am holding.
I can only stack a block on top of another block if I am holding the
    block being stacked.
I can only stack a block on top of another block if the block onto
    which I am stacking the block is clear.
Once I put down or stack a block, my hand becomes empty.
====================================================================
```

Listing 1: Blocksworld Domain Description

A.1.2 Domain included in the prompt (for Optimal planning)

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

Pick up a block. It takes 1 minute to pick up a block.
Unstack a block from on top of another block. It takes 1 minute to
    unstack a block from on top of another block.
Put down a block. It takes 1 minute to put down a block.
Stack a block on top of another block. It takes 1 minute to stack a
    block on top of another block.

I have the following restrictions on my actions:
I can only pick up or unstack one block at a time.
I can only pick up or unstack a block if my hand is empty.
I can 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 other blocks on top
     of it and if the block is not picked up.
I can only unstack a block from on top of another block if the block I
     am unstacking was really on top of the other block.
I can only unstack a block from on top of another block if the block I
     am unstacking is clear.
Once I pick up or unstack a block, I am holding the block.
I can only put down a block that I am holding.
I can only stack a block on top of another block if I am holding the
    block being stacked.
I can only stack a block on top of another block if the block onto
    which I am stacking the block is clear.
```

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

Listing 2: Blocksworld Domain Description

A.2 Example prompts and completion by GPT-3

Below we present the first instance of our 500 instances for each of the tasks and the corresponding completion by GPT-3 as an example.

A.2.1 Plan generation

```
[STATEMENT]
As initial conditions I have that, the red block is clear, the blue
    block is clear, the yellow block is clear, the hand is empty, the
    blue block is on top of the orange block, the red block is on the
    table, the orange block is on the table and the yellow block is on
     the table.
My goal is to have that the orange block is on top of the blue block.

My plan is as follows:

[PLAN]
unstack the blue block from on top of the orange block
put down the blue block
pick up the orange block
stack the orange block on top of the blue block
[PLAN END]

[STATEMENT]
As initial conditions I 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.
My goal is to have that the orange block is on top of the red block.

My plan is as follows:

[PLAN]
--------- GPT3 response ---------
unstack the yellow block from on top of the orange block
put down the yellow block
pick up the orange block
stack the orange block on top of the red block
[PLAN END]
--------- Extracted plan ---------
(unstack yellow orange)
(put-down yellow)
(pick-up orange)
(stack orange red)

-------- Ground truth plan ---------
(unstack yellow orange)
(put-down yellow)
(pick-up orange)
(stack orange red)

==============================SUCCESS==============================
```

Listing 3: Plan Generation

A.2.2 Optimal Planning

```
[STATEMENT]
```

```
As initial conditions I have that, the red block is clear, the blue
    block is clear, the yellow block is clear, the hand is empty, the
    blue block is on top of the orange block, the red block is on the
    table, the orange block is on the table and the yellow block is on
     the table..
My goal is to have that the orange block is on top of the blue block.
    I want to minimize the time taken to achieve my goal.
My plan is as follows:

[PLAN]
unstack the blue block from on top of the orange block
put down the blue block
pick up the orange block
stack the orange block on top of the blue block
[PLAN END]
The total time to execute the plan is 4 minutes.

[STATEMENT]
As initial conditions I 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..
My goal is to have that the orange block is on top of the red block. I
     want to minimize the time taken to achieve my goal.
My plan is as follows:

[PLAN]
--------- GPT3 response ---------
unstack the yellow block from on top of the orange block
put down the yellow block
pick up the orange block
stack the orange block on top of the red block
[PLAN END]
The total time to execute the plan is 4 minutes.
--------- Extracted plan ---------
(unstack yellow orange)
(put-down yellow)
(pick-up orange)
(stack orange red)

-------- Ground truth plan ---------
(unstack yellow orange)
(put-down yellow)
(pick-up orange)
(stack orange red)

----------------Optimal Plan----------------
---------Correct cost output by LLM---------
================================SUCCESS================================
```

Listing 4: Optimal planning

A.2.3 Replanning

```
[STATEMENT]
As initial conditions I have that, the red block is clear, the blue
    block is clear, the yellow block is clear, the hand is empty, the
    blue block is on top of the orange block, the red block is on the
    table, the orange block is on the table and the yellow block is on
     the table.
My goal is to have that the orange block is on top of the blue block.
My plan is as follows:

[PLAN]
unstack the blue block from on top of the orange block
```

```
put down the blue block
pick up the orange block
stack the orange block on top of the blue block
[PLAN END]

During execution, an unexpected event has occurred.
After executing the action "pick up the orange block" at step 3 in the
    plan, the following facts unexpectedly became true:
the orange block is clear,
the hand is empty,
the orange block is on top of the red block.

The following facts became unexpectedly false:
the red block is clear,
the hand is currently holding orange block.

After re-planning from the new state, the plan is as follows:
[PLAN]
unstack the orange block from on top of the red block
stack the orange block on top of the blue block
[PLAN END]

[STATEMENT]
As initial conditions I 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.
My goal is to have that the orange block is on top of the red block.
My plan is as follows:

[PLAN]
unstack the yellow block from on top of the orange block
put down the yellow block
pick up the orange block
stack the orange block on top of the red block
[PLAN END]

During execution, an unexpected event has occurred.
After executing the action "pick up the orange block" at step 3 in the
    plan, the following facts unexpectedly became true:
the orange block is clear,
the hand is empty,
the orange block is on top of the red block.

The following facts became unexpectedly false:
the red block is clear,
the hand is currently holding orange block.

After re-planning from the new state, the plan is as follows:
[PLAN]
--------- GPT3 response ---------
unstack the orange block from on top of the red block
put down the orange block
stack the orange block on top of the yellow block
[PLAN END]
--------- Extracted plan ---------
(unstack orange red)
(put-down orange)
(stack orange yellow)

-------- Ground truth plan ---------

==============================FAILURE==============================
```

Listing 5: Replanning when an unexpected event occurs and changes the state.

A.2.4 Plan Generalization

```
[STATEMENT]
As initial conditions I have that, the blue block is clear, the yellow
     block is clear, the magenta block is clear, the black block is
    clear, the cyan block is clear, the green block is clear, the
    violet block is clear, the gold block is clear, the hand is empty,
     the blue block is on the table, the yellow block is on the table,
     the magenta block is on the table, the black block is on the
    table, the cyan block is on the table, the green block is on the
    table, the violet block is on the table and the gold block is on
    the table.
My goal is to have that the blue block is on top of the cyan block,
    the yellow block is on top of the blue block, the magenta block is
     on top of the green block, the black block is on top of the
    yellow block, the cyan block is on top of the gold block, the
    green block is on top of the black block and the violet block is
    on top of the magenta block.

My plan is as follows:

[PLAN]
pick up the cyan block
stack the cyan block on top of the gold block
pick up the blue block
stack the blue block on top of the cyan block
pick up the yellow block
stack the yellow block on top of the blue block
pick up the black block
stack the black block on top of the yellow block
pick up the green block
stack the green block on top of the black block
pick up the magenta block
stack the magenta block on top of the green block
pick up the violet block
stack the violet block on top of the magenta block
[PLAN END]

[STATEMENT]
As initial conditions I have that, the red block is clear, the blue
    block is clear, the black block is clear, the cyan block is clear,
     the green block is clear, the violet block is clear, the silver
    block is clear, the hand is empty, the red block is on the table,
    the blue block is on the table, the black block is on the table,
    the cyan block is on the table, the green block is on the table,
    the violet block is on the table and the silver block is on the
    table.
My goal is to have that the red block is on top of the blue block, the
     blue block is on top of the silver block, the black block is on
    top of the red block, the cyan block is on top of the green block,
     the violet block is on top of the black block and the silver
    block is on top of the cyan block.

My plan is as follows:

[PLAN]
--------- GPT3 response ---------
pick up the red block
stack the red block on top of the blue block
pick up the blue block
stack the blue block on top of the silver block
pick up the black block
stack the black block on top of the red block
pick up the cyan block
stack the cyan block on top of the green block
```

```
pick up the violet block
stack the violet block on top of the black block
pick up the silver block
stack the silver block on top of the cyan block
[PLAN END]
--------- Extracted plan ---------
(pick-up red)
(stack red blue)
(pick-up blue)
(stack blue silver)
(pick-up black)
(stack black red)
(pick-up cyan)
(stack cyan green)
(pick-up violet)
(stack violet black)
(pick-up silver)
(stack silver cyan)

-------- Ground truth plan ---------
(pick-up cyan)
(stack cyan green)
(pick-up silver)
(stack silver cyan)
(pick-up blue)
(stack blue silver)
(pick-up red)
(stack red blue)
(pick-up black)
(stack black red)
(pick-up violet)
(stack violet black)

==============================FAILURE==============================
```

Listing 6: Plan generalization

A.2.5 Plan Reuse

```
[STATEMENT]
As initial conditions I have that, the red block is clear, the blue
    block is clear, the yellow block is clear, the hand is empty, the
    blue block is on top of the orange block, the red block is on the
    table, the orange block is on the table and the yellow block is on
     the table.
My goal is to have that the orange block is on top of the blue block.
My plan is as follows:

[PLAN]
unstack the blue block from on top of the orange block
put down the blue block
pick up the orange block
stack the orange block on top of the blue block
[PLAN END]

[STATEMENT]
As initial conditions I have that, the red block is clear, the blue
    block is clear, the yellow block is clear, the hand is empty, the
    blue block is on top of the orange block, the red block is on the
    table, the orange block is on the table and the yellow block is on
     the table.
My goal is to have that the red block is clear, the orange block is
    clear, the yellow block is clear, the hand is currently holding
    blue block, the red block is on the table, the orange block is on
    the table and the yellow block is on the table.
```

```
My plan is as follows:

[PLAN]
--------- GPT3 response ---------
unstack the blue block from on top of the red block
put down the blue block
pick up the red block
stack the red block on top of the orange block
[PLAN END]

-------- Ground truth plan ---------
(unstack blue orange)

==============================FAILURE==============================
```

Listing 7: Using a certain prefix of the plan provided in the example

A.2.6 Robustness to Goal Reformulation (Shuffling goal predicates)

```
[STATEMENT]
As initial conditions I have that, the red block is clear, the blue
    block is clear, the yellow block is clear, the hand is empty, the
    blue block is on top of the orange block, the red block is on the
    table, the orange block is on the table and the yellow block is on
     the table.
My goal is to have that the orange block is on top of the blue block.

My plan is as follows:

[PLAN]
unstack the blue block from on top of the orange block
put down the blue block
pick up the orange block
stack the orange block on top of the blue block
[PLAN END]

[STATEMENT]
As initial conditions I have that, the red block is clear, the blue
    block is clear, the yellow block is clear, the hand is empty, the
    blue block is on top of the orange block, the red block is on the
    table, the orange block is on the table and the yellow block is on
     the table.
My goal is to have that the orange block is on top of the blue block.

My plan is as follows:

[PLAN]
--------- GPT3 response ---------
unstack the blue block from on top of the orange block
put down the blue block
pick up the orange block
stack the orange block on top of the blue block
[PLAN END]
--------- Extracted plan ---------
(unstack blue orange)
(put-down blue)
(pick-up orange)
(stack orange blue)

-------- Ground truth plan ---------
(unstack blue orange)
(put-down blue)
(pick-up orange)
(stack orange blue)

```

```
=================== (GOAL ORDERING CHANGE) SUCCESS ===================
```

Listing 8: Changing the order of goal predicates

A.2.7 Robustness to Goal Reformulation (Full→Partial)

```
[STATEMENT]
As initial conditions I have that, the red block is clear, the blue
    block is clear, the yellow block is clear, the hand is empty, the
    blue block is on top of the orange block, the red block is on the
    table, the orange block is on the table and the yellow block is on
     the table.
My goal is to have that the orange block is on top of the blue block,
    the red block is clear, the red block is on the table, the blue
    block is on the table, the hand is empty, the yellow block is on
    the table, the orange block is clear and the yellow block is clear
    .

My plan is as follows:

[PLAN]
unstack the blue block from on top of the orange block
put down the blue block
pick up the orange block
stack the orange block on top of the blue block
[PLAN END]

[STATEMENT]
As initial conditions I have that, the red block is clear, the blue
    block is clear, the yellow block is clear, the hand is empty, the
    blue block is on top of the orange block, the red block is on the
    table, the orange block is on the table and the yellow block is on
     the table.
My goal is to have that the orange block is on top of the blue block.

My plan is as follows:

[PLAN]
--------- GPT3 response ---------
unstack the blue block from on top of the orange block
put down the blue block
pick up the orange block
stack the orange block on top of the blue block
[PLAN END]
--------- Extracted plan ---------
(unstack blue orange)
(put-down blue)
(pick-up orange)
(stack orange blue)

-------- Ground truth plan ---------
(unstack blue orange)
(put-down blue)
(pick-up orange)
(stack orange blue)

====================== (Full->Partial) SUCCESS ======================
```

Listing 9: Providing an example with fully specified goal and asking to complete for the same example with a partially specified goal

A.2.8 Robustness to Goal Reformulation (Partial→Full)

```
[STATEMENT]
```

```
As initial conditions I have that, the red block is clear, the blue
    block is clear, the yellow block is clear, the hand is empty, the
    blue block is on top of the orange block, the red block is on the
    table, the orange block is on the table and the yellow block is on
     the table.
My goal is to have that the orange block is on top of the blue block.

My plan is as follows:

[PLAN]
unstack the blue block from on top of the orange block
put down the blue block
pick up the orange block
stack the orange block on top of the blue block
[PLAN END]

[STATEMENT]
As initial conditions I have that, the red block is clear, the blue
    block is clear, the yellow block is clear, the hand is empty, the
    blue block is on top of the orange block, the red block is on the
    table, the orange block is on the table and the yellow block is on
     the table.
My goal is to have that the orange block is on top of the blue block,
    the red block is clear, the red block is on the table, the blue
    block is on the table, the hand is empty, the yellow block is on
    the table, the orange block is clear and the yellow block is clear
    .

My plan is as follows:

[PLAN]
--------- GPT3 response ---------
unstack the blue block from on top of the orange block
put down the blue block
pick up the orange block
stack the orange block on top of the blue block
unstack the red block from on top of the orange block
put down the red block
pick up the yellow block
stack the yellow block on top of the red block
[PLAN END]
--------- Extracted plan ---------
(unstack blue orange)
(put-down blue)
(pick-up orange)
(stack orange blue)
(unstack red orange)
(put-down red)
(pick-up yellow)
(stack yellow red)

-------- Ground truth plan ---------
(unstack blue orange)
(put-down blue)
(pick-up orange)
(stack orange blue)

====================== (Partial->Full) FAILURE ======================
```

Listing 10: Providing an example with partially specified goal and asking to complete for the same example with a fully specified goal