

On the Limit of Language Models as Planning Formalizers

Cassie Huang Li Zhang

Drexel University

{Cassie.Huang}@drexel.edu {Harry.Zhang}@drexel.edu

Abstract

Large Language Models have been found to create plans that are neither executable nor verifiable in grounded environments. An emerging line of work demonstrates success in using the LLM as a formalizer to generate a formal representation of the planning domain in some language, such as Planning Domain Definition Language (PDDL). This formal representation can be deterministically solved to find a plan. We systematically evaluate this methodology while bridging some major gaps. While previous work only generates a partial PDDL representation, given templated, and therefore unrealistic environment descriptions, we generate the complete representation given descriptions of various naturalness levels. Among an array of observations critical to improve LLMs' formal planning abilities, we note that most large enough models can effectively formalize descriptions as PDDL, outperforming those directly generating plans, while being robust to lexical perturbation. As the descriptions become more natural-sounding, we observe a decrease in performance and provide detailed error analysis.¹

1 Introduction

Large Language Models (LLMs) can make *informal plans*, such as suggesting ideas for parties or giving general advice on immigration. However, most users, let alone automated agents like robots, would not be able to actually execute those plans step-by-step to fruition – either to organize parties or acquire visas – without significant prior knowledge or external help. This inability to make executable plans lies in LLMs' inability of grounding and formal reasoning (Liu et al., 2023b; Pan et al., 2023; Zhang et al., 2023). Cutting-edge research in the community has evaluated LLMs' ability to make *formal plans* in grounded environments, such

as textual simulations, where all objects and actions represent actualities in the real world. Therefore, any resulting plan that formally involves those objects and actions would be executable and verifiable by nature. Although formal planning has been desirable in the history of AI (Weld, 1999), recent work found that even state-of-the-art LLMs are unable to generate formal plans (Silver et al., 2024; Valmeekam et al., 2024a; Stechly et al., 2024).

Instead of using the LLM as a planner to generate the plan directly, an alternative line of work uses the LLM as a *formalizer*. Here, the LLM generates a formal representation of a planning domain, for example in PDDL, based on some natural language descriptions of the environment. This representation can then be fed into a solver to find the plan deterministically (see Figure 1). Previous work has shown feasibility and success of LLM-as-formalizer over LLM-as-planner in various domains (Lyu et al., 2023; Xie et al., 2023; Liu et al., 2023a; Zhang et al., 2024a,c; Zhu et al., 2024), as LLMs are more capable of information extraction than formal reasoning (Zhang et al., 2024b). However, the above work has two major shortcomings. First, to simplify the task and evaluation, most have only attempted to generate a partial PDDL representation while assuming the rest is provided, which is often unrealistic in real life. Second, the language used to describe the environments is often artificially templated and structured, leading to potential overestimation of models' ability.

This paper explores the limit of LLM-as-formalizer devoid of the above two simplifications. We use LLMs to generate the entirety of a PDDL representation, including the domain file and the problem file, given a natural-sounding description of the environment and the task (see Figure 1). On four widely used planning simulations from the International Planning Competition, we benchmark a suite of LLMs on generating PDDL that is both solvable and correct. As the descriptions in these

¹Our code and data can be found at <https://github.com/CassieHuang22/llm-as-pddl-formalizer>.

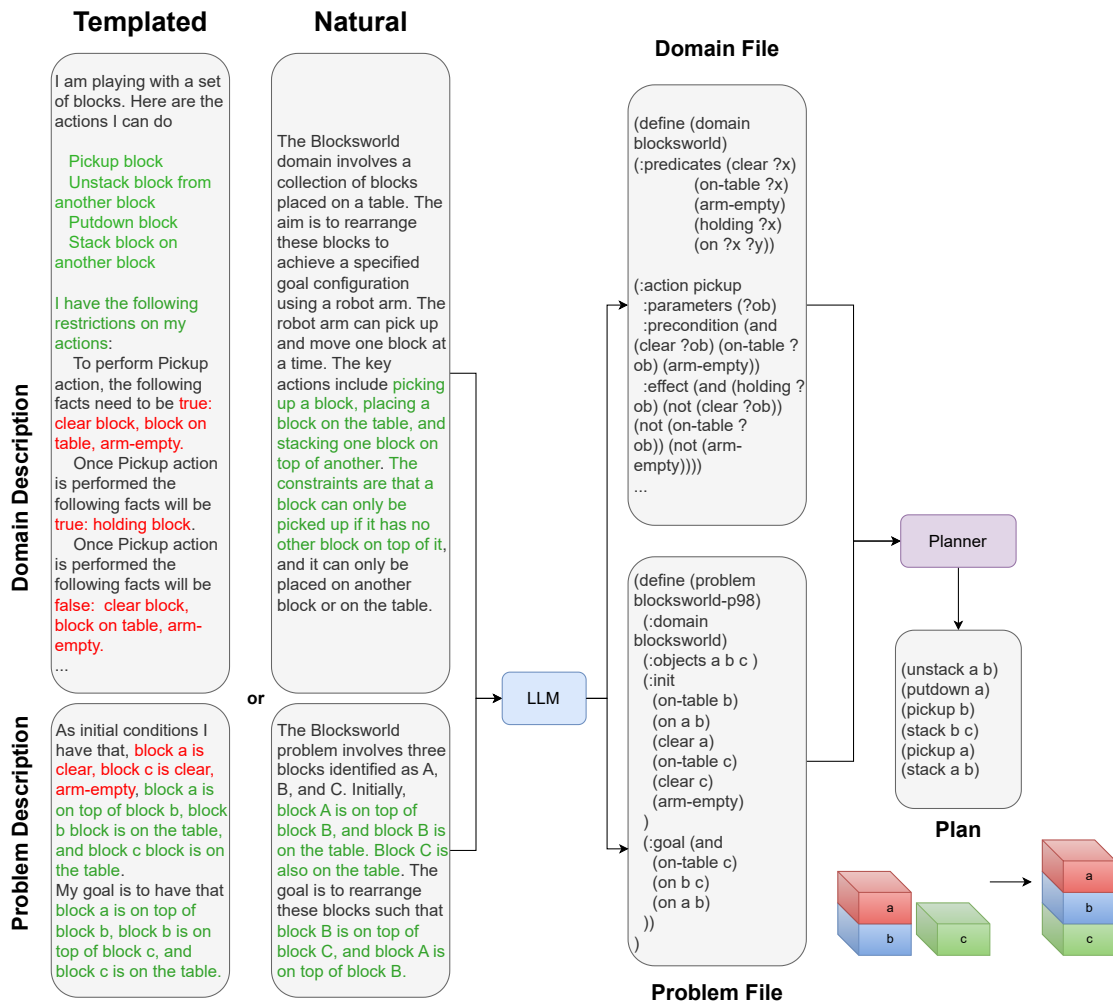


Figure 1: LLM-as-formalizer uses natural language descriptions to generate the Domain and Problem File in PDDL, then these are given to a planner to find a plan. We explore the effect of natural-ness of the language in the description, by giving the model both templated and natural descriptions. Examples of Domain Descriptions and Problem Descriptions from the BlocksWorld Domain are shown. The green text displays what the two examples have in common (listing all possible actions and restrictions) and the red text displays text that is not considered natural. The “Templated” text corresponds to the “Heavily Templated” version discussed in Section 4.

datasets are templated, we also construct model-generated, human-validated descriptions that are natural-sounding to different levels.

Our work is the first to systematically analyze state-of-the-art LLMs’ ability to use the trending methodology of LLM-as-formalizer on the highly challenging task of formal planning. We put forward an array of observations that will benefit future efforts. Discussed in detail in Section 5, our key findings are as follows.

- Strongest models like gpt-4o, o3-mini, or DeepSeek-R1 can not only plan but also formalize entire PDDL, while other models like Gemma-2 or Llama-3.1 (up to 405B) cannot

do either.

- LLM-as-formalizer often outperforms LLM-as-planner for many model-data combinations, though not always.
- A more human-like, natural description of the planning environment creates more challenges for models.
- LLM-as-formalizer is much more robust to long-tail lexical distribution than LLM-as-planner.
- Weaker models commonly succumb to syntax errors, while all models face semantic errors.

2 Task: Formal Planning with PDDL

Formal planning (or classical planning) with PDDL involves a domain file ($\mathbb{D}\mathbb{F}$) and problem file ($\mathbb{P}\mathbb{F}$) (Figure 1). The $\mathbb{D}\mathbb{F}$ describes general properties in a planning domain that hold true across problems, while the $\mathbb{P}\mathbb{F}$ describes specific configurations of each problem instance. Concretely, the $\mathbb{D}\mathbb{F}$ defines all available actions for a state-based environment as well as predicates that represent the properties of object types. Each action definition contains the name of the action, parameters and semantics. The semantics of an action includes the preconditions which describe the necessary world states where the action is valid to execute, and effects which describe how the states change after the action is executed. The $\mathbb{P}\mathbb{F}$ defines the involved objects, initial states, and goal states. These two files are then given to a deterministic planner which will algorithmically search for a plan. Such a plan is a series of executable, instantiated actions that sequentially transition the world states from initial to goal. In the AI community, classical planning has been deemed an effective approach to solve real-world users’ problems, as the process is precise, explainable, verifiable, and deterministic.

However, formal planning demands a well-crafted $\mathbb{D}\mathbb{F}$ and $\mathbb{P}\mathbb{F}$ pair. In a real-world planning scenario, an average user would not describe their environment and problem with PDDL, but more likely with a textual domain description ($\mathbb{D}\mathbb{D}$) and problem description ($\mathbb{P}\mathbb{D}$). While the $\mathbb{D}\mathbb{D}$ and $\mathbb{P}\mathbb{D}$ technically describe all the necessary information from the $\mathbb{D}\mathbb{F}$ and $\mathbb{P}\mathbb{F}$ in natural language, information that can be easily inferred by humans is omitted as the language in the descriptions gets more natural. For example, when describing the action “picking up a block”, a very natural $\mathbb{D}\mathbb{D}$ would omit the effect that our hand is no longer empty. These descriptions are then translated into PDDL, which can then be used to find a plan when given to a solver. Hence, we focus on the textual flavor of formal planning: given a $\mathbb{D}\mathbb{D}$ and $\mathbb{P}\mathbb{D}$, the model outputs a successful plan with regard to the $\mathbb{D}\mathbb{F}$ and $\mathbb{P}\mathbb{F}$ that are withheld from the model.

3 Methodology: LLM-as-Formalizer

To tackle the task above, recent work involving LLMs diverged into two methodologies. The first, **LLM-as-planner**, directly uses LLMs to generate a plan based on the $\mathbb{D}\mathbb{D}$ and $\mathbb{P}\mathbb{D}$. The second, **LLM-as-formalizer**, uses LLMs to recover the $\mathbb{D}\mathbb{F}$ and

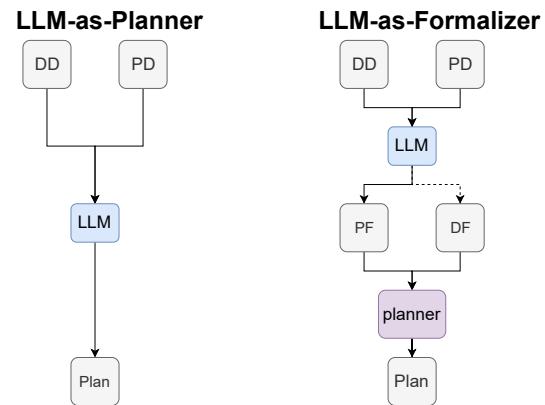


Figure 2: Methodologies for using LLMs in planning. LLM-as-Planner generates the plan directly, while LLM-as-Formalizer translates the $\mathbb{D}\mathbb{D}$ and $\mathbb{P}\mathbb{D}$ into PDDL. Previous work like Liu et al. (2023a) use the LLM to generate partial PDDL such as $\mathbb{P}\mathbb{F}$ only, while we generate the entire PDDL including $\mathbb{P}\mathbb{F}$ and $\mathbb{D}\mathbb{F}$. Note the $\mathbb{D}\mathbb{D}$ and $\mathbb{P}\mathbb{D}$ are always provided and the $\mathbb{D}\mathbb{F}$ and $\mathbb{P}\mathbb{F}$ are always generated by the LLM.

$\mathbb{P}\mathbb{F}$, before using a deterministic planner to arrive at the plan (Figure 2). Our work will focus on the second while using the first as a baseline. LLM-as-formalizer is in essence neuro-symbolic, where LLMs help define the structured representation that is otherwise prohibitively costly to annotate and brittle to generalize. Existing works in this line demonstrated success but only generate a partial PDDL representation, while assuming the rest. For example, some works generate only the $\mathbb{P}\mathbb{F}$ goals while the $\mathbb{D}\mathbb{F}$ and remaining $\mathbb{P}\mathbb{F}$ are given (Lyu et al., 2023; Xie et al., 2023). Others generate the entire $\mathbb{P}\mathbb{F}$ while the $\mathbb{D}\mathbb{F}$ is given (Liu et al., 2023a; Zhang et al., 2024a; Zuo et al., 2024). There have also been works that generate the action semantics in the $\mathbb{D}\mathbb{F}$ while the remaining parts of the $\mathbb{D}\mathbb{F}$ and entire $\mathbb{P}\mathbb{F}$ are provided (Zhang et al., 2024c; Zhu et al., 2024), and work that generates the $\mathbb{D}\mathbb{F}$ while the $\mathbb{P}\mathbb{F}$ is given (Wong et al., 2023; Guan et al., 2023). While these assumptions of given PDDL components simplify the task and evaluation, they are often unrealistic. Therefore, we bridge this gap by asking the LLM to predict the entire PDDL – both the $\mathbb{D}\mathbb{F}$ and $\mathbb{P}\mathbb{F}$.²

²It is however minimally necessary to provide the action space, the identifiers and parameters of the actions in $\mathbb{D}\mathbb{F}$, so the agent knows *what* actions are possible.

4 Evaluation: Datasets, Metrics, Models

To evaluate both approaches above, we work with *fully-observed* textual environments. Here, the provided \mathbb{D} and \mathbb{P} contain all necessary information for the model to make a complete plan.

4.1 Datasets

We consider four simulated planning domains, BlocksWorld, Logistics, Barman from the International Planning Competition (IPC, 1998), and Mystery BlocksWorld (Valmeekam et al., 2024a).

BlocksWorld, also used in Liu et al. (2023a), is a domain to rearrange stacks of blocks on a table using a robotic arm. The domain includes 1 type of entity, 5 predicates, and 4 actions.

Mystery BlocksWorld obfuscates the original BlocksWorld domain by replacing all the names of the types, predicates, actions, and objects with nonsensical words, akin to a *wug test* (Berko, 1958). This dataset as a control group can effectively test whether models create plans via lexical pattern-matching and memorization.

Logistics, also used in Guan et al. (2023), is a domain to transport packages across different locations using both trucks and airplanes. In this domain, there are 6 types of entities, 3 predicates, and 6 actions.

Barman, also used in Zhu et al. (2024), is a domain to create cocktails from ingredients using different containers and two robotic arms. In this domain, there are 7 types of entities, 13 predicates, and 12 actions.

Each dataset comes with ground-truth PDDL files describing the domain (\mathbb{D}) and problem (\mathbb{P}). The input to the model is a natural language description of the domain (\mathbb{D}) and problem (\mathbb{P}). The output of the model is a plan, namely a sequence of instantiated actions defined in the \mathbb{D} . For each of these datasets, the natural language description (\mathbb{D} s and \mathbb{P} s) were created in 3 different levels of naturalness.

Heavily Templated. For BlocksWorld, Logistics and Barman, the Heavily Templated \mathbb{D} and \mathbb{P} are generated using the same template as Mystery BlocksWorld (Valmeekam et al., 2024a). This description is almost a word-by-word translation of PDDL. For example, for the ‘pick-up’ action in BlocksWorld, the ground-truth PDDL \mathbb{D} would be the following:

```
(:action pick-up
:parameters (?b - block)
:precondition (and (clear ?b) (on-table ?b)
(arm-empty))
:effect (and (not (on-table ?b)) (not
(clear ?b)) (not (arm-empty)) (holding ?b))
)
```

while the Heavily Templated \mathbb{D} is:

```
To perform Pickup action, the following
facts need to be true: clear block, block
on table, arm-empty.
Once Pickup action is performed the
following facts will be true: holding
block.
Once Pickup action is performed the
following facts will be false: clear block,
block on table, arm-empty.
```

From an application point of view, spelling out all preconditions and effects in terms of the predicates is paradoxical, as it assumes the user already has the algorithmic awareness of PDDL.

Moderately Templated. The \mathbb{D} and \mathbb{P} are generated using the same template as the original BlocksWorld dataset, following Valmeekam et al. (2024a). For example, the Moderately Templated description of the ‘pick-up’ action is:

```
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
clear. A block is clear if the block has no
other blocks on top of it and if the block
is not picked up.
```

While more natural-sounding than the Heavily Templated version, the description still explicitly discusses the preconditions and effects as well as predicates like ‘clear’.

Natural. A realistic pair of \mathbb{D} and \mathbb{P} should emulate how real-life users would describe the planning domain and problem, such that a human problem solver would understand and have just enough information to find a plan. To create such descriptions, we use a human-in-the-loop, model-assisted data generation approach.

To generate \mathbb{D} s, we ask gpt-4o with high temperature to generate and paraphrase a seed annotated \mathbb{D} , and then manually verify the correctness by making sure it lists the correct predicates, preconditions and effects, which are not unique. We next verify the naturalness of the generated text by making sure there were variations in language throughout all descriptions generated, but unnecessary information was not given out.

To generate \mathbb{P} s, we provide the model with a symbolic configuration containing the initial and goal states of the problem. The model then ‘human-

izes’ the problem by making it sound natural, given a couple of seed exemplars. We manually verify the correctness of the datasets of the non-templated problems by hand by comparing them against the problem configurations. We then verify the naturalness of the $\mathbb{P}\mathbb{D}$ by making sure there is variation but no ambiguity in its language.

In the following example of a Natural description of BlocksWorld:

The robot arm can pick up and move one block at a time from one position to another. It is only able to move the top block from any stack or table, and have only one block held by the robot arm at a time. The main actions available are ‘pick up’, ...

The example no longer discusses the preconditions and effects of each action one by one, but rather focuses on the general rules to the domain. These rules apply to not only ‘pick-up’ but also other actions. Therefore, the $\mathbb{D}\mathbb{D}$ can be much more concise, requires less algorithmic awareness, and is more realistic.

In total, we construct 100 problems varying in complexity for all domains. For each of the two Templated descriptions, there is one $\mathbb{D}\mathbb{D}$ paired with each of 100 $\mathbb{P}\mathbb{D}$ s. For the Natural descriptions, there are 100 different pairs of $\mathbb{D}\mathbb{D}$ s and $\mathbb{P}\mathbb{D}$ s. We refer to these datasets as BlocksWorld-100, MysteryBlocksWorld-100, Logistics-100 and Barman-100. Moderately Templated and Natural datasets are available only for BlocksWorld-100 and Logistics-100 due to the size of Barman-100. Data examples can be found in Appendix A.

4.2 Metrics

Following past work (Guan et al., 2023; Zhu et al., 2024), the plan produced from LLM-as-Planner is validated using VAL (Howey et al., 2004) against the ground-truth $\mathbb{D}\mathbb{F}$ s and $\mathbb{P}\mathbb{F}$ s provided above. This is done instead of being compared against “ground-truth” plans like some work (Lyu et al., 2023; Liu et al., 2023b; Pan et al., 2023) since there could be multiple correct plans. Similarly, the predicted $\mathbb{D}\mathbb{F}$ and $\mathbb{P}\mathbb{F}$ for the LLM-as-formalizer approach are not compared against the ground-truth, as only the eventual plan is validated because there might be more than one way to formalize the planning domain and problem in PDDL.

We evaluate the predicted plans following metrics from Zuo et al. (2024): using *solvability* and *correctness*. Solvability only applies to LLM-as-formalizer and indicates the percentage of solv-

able *predicted* $\mathbb{D}\mathbb{F}$ s and $\mathbb{P}\mathbb{F}$ s, regardless of whether the resulting plan can be executed based on the *ground-truth* $\mathbb{D}\mathbb{F}$ and $\mathbb{P}\mathbb{F}$. Correctness indicates the percentage of actually correct plans. However, different from Zuo et al. (2024), Solvability was determined using the planner `dual-bfws-ffparser` implemented by Muise (2016) and Correctness was evaluated using VAL³. More information about the solver and validator can be found in the Appendix, Section C.

4.3 Models

For both of the LLM-as-planner and LLM-as-formalizer approaches, we consider a number of models, including open-source and closed-source LLMs varying in size, including gemma-2-9b127b-it (Team et al., 2024), llama-3.1-8B170B1405B-Instruct (Dubey et al., 2024), DeepSeek-R1-Distill-Llama-8B170B1671B (Guo et al., 2025), gpt-4o-mini-2024-07-18, gpt-4o-2024-08-06, and o3-mini-2025-01-31⁴. We query these models using KANI (Zhu et al., 2023) with default hyper-parameters. The open-source models are run using 4 RTX A6000 GPUs, averaging about 1062 input and output tokens for the LLM-as-formalizer approach in BlocksWorld-100.

To emulate real-life application with minimal user interference, we use zero-shot prompts for all naturalness levels across all datasets (see prompts in Appendix B). We would also like to note that since the LLM has to generate the entire $\mathbb{D}\mathbb{F}$ and $\mathbb{P}\mathbb{F}$, using few-shot prompting would not be meaningful. Each dataset corresponds to one unique domain and we cannot provide the LLM the $\mathbb{D}\mathbb{F}$ in the prompt as the model is supposed to predict it. As for the $\mathbb{P}\mathbb{F}$, we also cannot provide it to the model because the typing, predicates, and usage in the $\mathbb{P}\mathbb{F}$ would also give away critical information from the file. Nevertheless, we performed a small set of experiments using different prompting methods, such as few-shot and chain-of-thought prompting on Heavily Templated and Natural BlocksWorld-100. Results and discussion of these experiments can be found in the Appendix Sections D.5 and D.6.

5 Results and Observations

In this section, we display our results as well as perform an in-depth analysis of the strengths and weaknesses of LLMs in formal planning, to under-

³nms.kcl.ac.uk/planning/software/val.html

⁴platform.openai.com/docs/models




Figure 3: Performance of both usages of LLMs on Heavily Templated BlocksWorld-100, Logistics-100, and MysteryBlocksWorld-100. Detailed results are shown in Appendix D. Due to the cost and zero solvability on our easiest dataset, results of Llama-405B and DeepSeek-8B | 70B are omitted.

stand the impact of the model choice, naturalness of the description, content of the task, and difficulty of the problem.

5.1 Can LLMs formalize?

We seek to understand the extent to which LLMs can act as a formalizer to generate *entire* PDDL, rather than partial components seen in previous work. Figure 3 displays the results of our experiments on Heavily Templated BlocksWorld-100, Logistics-100 and MysteryBlocksWorld-100. Results on the most complex domain Barman-100 is omitted due to close-to-zero performance for all models.

These results demonstrate that **GPT-family LLMs and DeepSeek-R1 can decently generate PDDL**, while open-source models like Gemma and Llama (even up to 405B parameters) struggle. As a formalizer, gpt-4o-mini, gpt-4o and o3-mini demonstrate non-trivial and increasing performance on BlocksWorld-100. On the more complex Logistics-100, gpt-4o-mini succumbs to zero performance, whereas the other two show decreased performance. The solvability of gpt-4o-mini is often much higher than its correctness, suggesting a good grasp of the PDDL syntax but a lack of semantic understanding. In contrast, the solvability of gpt-4o and o3-mini is often 80% to 100% of their correctness. On the other hand, **open-sourced models can rarely generate PDDL**, a low-resource language, despite them being reportedly strong at generating high-resource languages like Python (Cassano et al., 2022). All Llama models up to 405B cannot generate any solvable PDDL across all three datasets, while Gemma models show poor, though

non-zero, performance on BlocksWorld-100 and Logistics-100, and strong performance on MysteryBlocksWorld-100.

5.2 Should LLMs formalize?

Between LLM-as-planner and LLM-as-formalizer, which is the preferred methodology? Figure 3 shows that on BlocksWorld-100, gpt-4o is able to generate solvable PDDL 64 times, out of 100, and of those 64 plans, 60 of them are correct. This far surpasses the LLM-as-planner baseline, which only found correct plans 33/100 times. This trend holds for Logistics-100 as well as the Moderately Templated and Natural BlocksWorld-100 data (Figure 4). On MysteryBlocksWorld-100, we can see that LLM-as-formalizer can generate 70/100 correct plans, which far surpassed LLM-as-planner which did not find a single correct plan as the description becomes unorthodox. The superiority of LLM-as-formalizer also extends to gpt-4o-mini but not o3-mini, who shows strong performance as a planner. These results demonstrate that **LLM-as-formalizer greatly outperforms LLM-as-planner** in most cases, whenever these LLMs can formalize PDDL *at all*. However, these results also show that models that cannot formalize (e.g., Llama models) can still plan, though with close-to-zero performance.

5.3 The more natural, the harder?

We now examine whether using humanized descriptions makes the problem more difficult. Results from Figure 4 show that on BlocksWorld-100 as the problem sounds more similar to PDDL and less natural, the performance of all the models improves. Similar results hold for the Logistics do-

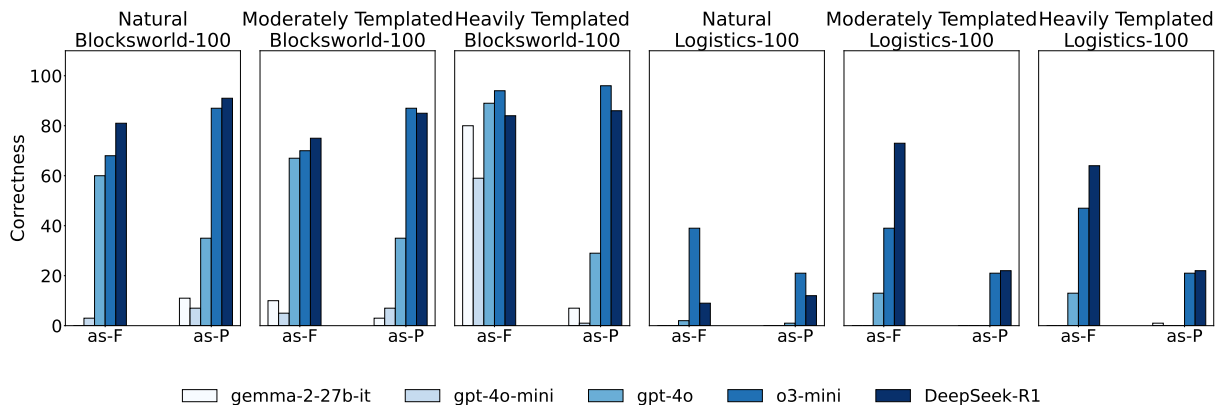


Figure 4: Performance of LLM-as-planner (as-P) and LLM-as-formalizer (as-F) across different naturalness level of description on BlocksWorld-100 and Logistics-100. Detailed results are shown in Appendix D.

main (see results in Appendix D). This suggests that **a more natural-sounding domain and problem description is much more challenging** than templated, less natural sounding descriptions. One potential explanation is that pattern matching a template back to PDDL is much easier than having to first parse all the predicates and objects from a passage. Another reason is a more natural sounding description may leave out implicit common-sense. For example, the Natural BlocksWorld-100 data does not explicitly specify that a block is ‘clear’, because any human who reads that a block is “on top of a stack” can understand that there is no block on top of it and hence ‘clear’ to be moved. However, models often fail to invoke this knowledge and will leave out the ‘clear’ predicate, leading to unsolvable PDDL or incorrect plans.

5.4 Do LLMs memorize pretraining?

Do LLMs generate plans or formalize PDDL based on what they have memorized in their training data? We determine this by looking at the results on MysteryBlocksWorld-100, a derivative of BlocksWorld where all names are perturbed and nonsensical. From Figure 3, we can see that LLM-as-planner was not able to find a single correct plan using either gpt-4o-mini or gpt-4o. However, gpt-4o-as-formalizer surpassed this baseline with a Correctness score of 70/100. Surprisingly, Gemma models performed very well on the MysteryBlocksWorld-100 dataset with near perfect solvability and correctness for both models. We conclude that the results are mixed across models, which contradicts previous works cited involving PDDL generation which seem to claim that LLM-as-formalizer is clearly better than

other methodologies. Despite the inconclusiveness among models, there are more clear-cut results between two groups of models: those that can formalize (eg. gpt-4o and o3-mini) and those that cannot formalize (eg. all Llama models and gpt-4o-mini). Among the models that can formalize, we find that the results are robust to lexical perturbations.

5.5 What kind of errors?

In this section, we discuss the kinds of errors made in PDDL generation. We perform an error analysis on a random 20 sample subset of problems where a plan was not found, or the found plan was not correct. From there, we categorize the errors by syntax errors in either file, semantic errors in the \mathbb{DF} , and semantic errors in the \mathbb{PF} . Of the errors in the \mathbb{DF} , we determine finer-grained errors such as incorrect or missing actions, preconditions, effects, predicates, and parameters. The error analysis for both Natural BlocksWorld-100 and Logistics-100 can be found in Figure 5.

For the open-source models, the most common error is syntax errors on BlocksWorld-100 and Logistics-100. For example, models repeatedly use the keyword ‘preconditions’ instead of ‘precondition’ which might suggest a lack of grasp of the PDDL syntax. Gemma models like gemma-2-27b make significantly fewer syntax errors (3 out of 20 on BlocksWorld-100) than Llama models like Llama-3.1-70B (20 out of 20), despite being smaller. Despite the syntax errors, **open-sourced models still make many semantic errors in the \mathbb{DF} and \mathbb{PF}** , which include missing predicates and incorrect effects in the actions. We also observe a significant gap between the number of

plans that were found, and the number of found plans that were correct. We find that the most common error made was swapping the parameters in the preconditions of the ‘stack’ action, leading to incorrect plans. While making much fewer syntax errors, **GPT models frequently suffer from semantic errors**. Interestingly, the most common error made for gpt-4o come from the $\mathbb{P}\mathbb{F}$, which is intuitively easier to generate than the $\mathbb{D}\mathbb{F}$. A common error in the $\mathbb{P}\mathbb{F}$ is incorrect predicates in the initial state and goal state. A common error in the $\mathbb{D}\mathbb{F}$ is incorrect effects in the actions. For example, in the ‘unstack’ action, the model does not make the next block ‘clear’ when the top block has been placed in the hand. For MysteryBlocksWorld-100, there are barely any syntax errors or semantic errors in the $\mathbb{P}\mathbb{F}$ but rather the most common errors come from the $\mathbb{D}\mathbb{F}$. Since this domain is a result of lexical perturbation, formalizing in PDDL is akin to symbolic information extraction and translation, devoid of much use of commonsense knowledge. Due to the Heavily Templated descriptions, all the predicates would be listed out in the $\mathbb{P}\mathbb{D}$ and the model would just need to match them to PDDL syntax in the $\mathbb{P}\mathbb{F}$. While a similar essence, this is more challenging for the $\mathbb{D}\mathbb{F}$ since the clauses of preconditions and effects are more involved. From Table 9, a similar trend between BlocksWorld-100 and MysteryBlocksWorld-100 also suggests that the LLM-as-formalizer methodology is robust to such perturbation.

5.6 The curious case of DeepSeek-R1

The only exception to our main findings is DeepSeek-R1 for all datasets. The model shows strong performance as both a formalizer and planner for BlocksWorld-100, while showing mixed results for Logistics-100, and following a similar trend as GPT models for MysteryBlocksWorld-100. We can conclude that for simple datasets, using DeepSeek-R1 as a formalizer is preferred for some datasets, and when formalizing, the model is robust to lexical perturbations. However, for more complex datasets, using the model as a planner is preferred. This result may be due to how DeepSeek-R1 generates its response. The model outputs a reasoning string that thinks through every step of the plan and attempts to check its work by stepping through the plan. After examining several reasoning strings, we notice that when we use the model as a planner, it will think through the plan and when it gets stuck, it will at-

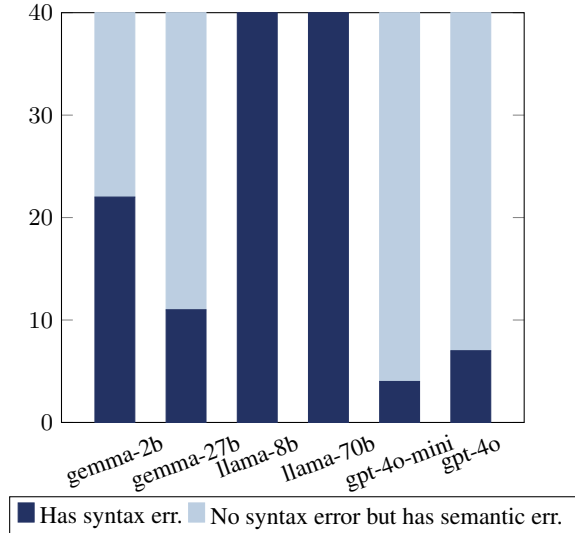


Figure 5: Number of syntax errors (solver cannot parse PDDL) and semantic errors (solver outputs either no plan or a wrong plan) in 40 manually annotated, randomly sampled errors in Natural BlocksWorld-100 and Logistics-100. See full results in Table 8.

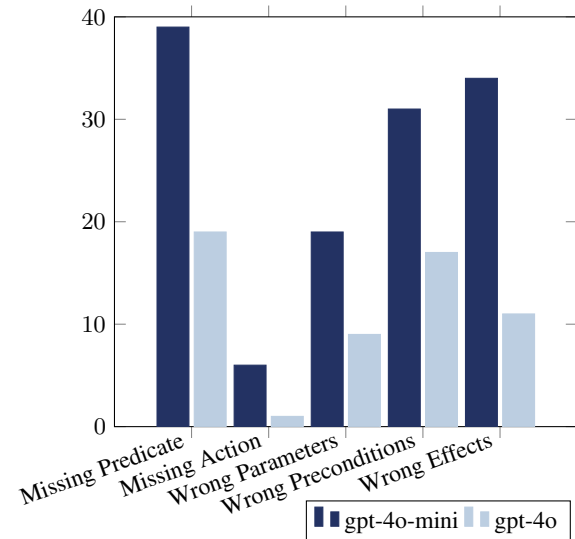


Figure 6: Number of fine-grained $\mathbb{D}\mathbb{F}$ errors in 40 manually annotated, randomly sampled errors in Natural BlocksWorld-100 and Logistics-100. See full results in Table 9.

tempt to pattern-match the example output given in the prompt. While the example output for each domain is not found in any of the datasets, we believe this may have caused the higher performance on more complex domains.

6 Related Work

Planning with LLMs There has been a large amount of research using LLMs for planning tasks. Some use LLMs for informal planning, also known

as script or procedure learning (Zhang et al., 2020; Lyu et al., 2021; Lal et al., 2024). While modern LLMs can make coherent and plausible informal plans, they are ungrounded and so lack executability and verifiability. Work that uses LLMs for formal planning in grounded environments generally concludes with the inability of such LLMs-as-planners (Silver et al., 2024; Valmeekam et al., 2024a; Stechly et al., 2024; Valmeekam et al., 2024b). Follow-up work tackles this shortcoming by using the LLM as a heuristic, not just a planner, such as by proposing candidate plans that are iteratively verified (Valmeekam et al., 2023; Kambhampati et al., 2024). While we consider the standard LLM-as-planner as a baseline, our focus is on LLM-as-formalizer, an alternative methodology for the same problem.

LLMs as PDDL formalizer Here, LLMs do not provide plans but rather generate a PDDL representation of the domain and problem, which is then run through a solver to find the plan. This methodology has proven successful in a number of recent works, where the LLM generates different parts but not all of the PDDL for simplified evaluation. Zuo et al. (2024); Zhang et al. (2024a); Liu et al. (2023a) use the LLM to predict the entire $\mathbb{P}\mathbb{F}$, while Xie et al. (2023); Lyu et al. (2023) predict just the goal for the $\mathbb{P}\mathbb{F}$. Some predict parts of the $\mathbb{D}\mathbb{F}$, such as Zhang et al. (2024c); Zhu et al. (2024) which generate the action semantics of the $\mathbb{D}\mathbb{F}$ and Wong et al. (2023) who also predicts the predicates from a candidate list. Closest to our work is Guan et al. (2023) which predicts the $\mathbb{D}\mathbb{F}$ as well as the $\mathbb{P}\mathbb{F}$ goal. However, our work of holistically generating PDDL shows that coming up with the initial state in the $\mathbb{P}\mathbb{F}$ is non-trivial (Section 5.5). Moreover, we vary the level of naturalness of descriptions in addition to the templated ones, which prove to be more challenging and insightful (Section 5.3).

While the above discussions pertain to LLMs generating PDDL, many work on embodied agents outside the NLP community tackle similar problems with different focus (Li et al., 2024).

LLM code generation Our work hinges on modern LLMs’ ability to generate code (Chen et al., 2021). In addition to writing or debugging programs (Jiang et al., 2024), LLMs are also used to generate formal, interim representations that are not necessarily PDDL for problem solving. For example, Gao et al. (2023); Lyu et al. (2023); Tang et al. (2024) use the LLM to generate executable Python code for solving symbolic problems. In other work, the gen-

erated code may not be executable and is provided to another LLMs to facilitate reasoning (Madaan et al., 2022; Zhang et al., 2023).

A table comparing a couple of these works can be seen in Table 1 in the Appendix (Section F).

7 Conclusion and Future Work

Amidst the trend of using LLMs as formalizers, especially to generate PDDL for planning, our comprehensive analysis can be distilled into the following key findings and suggestions for future work.

State-of-the-art LLMs can plan and formalize at a high level on simple domains like BlocksWorld, and to some degree on more complex domains like Logistics. Moreover, reasoning models trained to produce reasoning tokens before the answer (such as o3-mini and DeepSeek-R1) are superior planners than those who are not (such as gpt-4o). We suggest future work closely study these models (such as Valmeekam et al. (2024b)) from various angles such as the soundness of their reasoning chain, time and monetary cost, etc.

Most small open LLMs cannot plan nor formalize even on simple domains, succumbing to syntax errors. To improve their performance on formalizing, we suggest future efforts on methods such as finetuning, modular prompting, constrained decoding, retrieval by language documentation, wrapper of a high-resource language, and so on to improve the performance of generating not just PDDL, but also low-resource languages in general.

Formalizing is often superior to planning, but not always, with positive evidence for models like gpt-4o but negative evidence for models like DeepSeek-R1. Given the rising attention towards the LLM-as-formalizer methodology, we suggest future work validate this claim and extend to other environments such as more complex ones, partially observable ones, etc.

Natural descriptions are challenging in planning tasks, and **so is predicting complete PDDL**. We acknowledge that translating a simple natural language instruction to a specific PDDL component might be sufficient in some confined domains and a fruitful endeavor. However, we argue that generating a planning representation as complete as possible given a natural language query as complex and realistic as possible is a prerequisite to open-domain planning.

8 Limitation

A common and valid criticism for using these simulations or text problems for evaluation is that their settings may be too contrived and removed from reality. Nevertheless, it is likely that LLMs' satisfactory performance on these datasets is a necessary condition to success in real life.

While we only consider zero-shot prompting, and a few experiments with few-shot and chain-of-thought prompts, without any attempt for prompt tuning, it is possible that the models' performance significantly increases otherwise. Therefore, experimental results in all settings may be underestimated. Moreover, advanced prompting techniques such as self-refine, and voting can all potentially improve model performance. However, the study of those is out of the scope of this work.

While we advocate for the LLM-as-formalizer methodology over LLM-as-planner, the former's success may be dependent on the task. Highly symbolic tasks which can be relatively easily described, like BlocksWorld, are likely to favor LLM-as-formalizer. However, LLM-as-planner might shine in tasks with a more complex action space requiring common-sense knowledge that is easily accessed by pretraining. Furthermore, while we only consider the most straightforward LLM-as-planner prompting method, more involved methods, like Kambhampati et al. (2024) that combines LLM-as-planner with symbolic validation, will likely lead to a stronger baseline.

Since this work uses only the BlocksWorld, Mystery BlocksWorld, Logistics and Barman domains, it is a small toy example to the usage of LLMs as formalizers and are not representative to problems in the real world, which would be much more challenging. This may pose a risk to users using this code on real world problems.

The datasets we use and we propose are all under the MIT License.

Acknowledgment

We thank Peter Clark for providing invaluable input throughout this work and Andrew Zhu for providing technical support on LLM inference.

References

Jean Berko. 1958. The child's learning of english morphology. *Word*, 14(2-3):150–177.

Federico Cassano, John Gouwar, Daniel Nguyen, Sydney Nguyen, Luna Phipps-Costin, Donald Pinckney, Ming-Ho Yee, Yangtian Zi, Carolyn Jane Anderson, Molly Q Feldman, et al. 2022. Multipl-e: A scalable and extensible approach to benchmarking neural code generation. *arXiv preprint arXiv:2208.08227*.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Pal: Program-aided language models. In *International Conference on Machine Learning*, pages 10764–10799. PMLR.

Lin Guan, Karthik Valmeekam, Sarath Sreedharan, and Subbarao Kambhampati. 2023. Leveraging pre-trained large language models to construct and utilize world models for model-based task planning. *Advances in Neural Information Processing Systems*, 36:79081–79094.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

R. Howey, D. Long, and M. Fox. 2004. 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.

IPC. 1998. International planning competition. <https://www.icaps-conference.org/competitions>.

IPC. 2000. International planning competition. <https://www.icaps-conference.org/competitions>.

Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. 2024. A survey on large language models for code generation. *arXiv preprint arXiv:2406.00515*.

Subbarao Kambhampati, Karthik Valmeekam, Lin Guan, Mudit Verma, Kaya Stechly, Siddhant Bhambri, Lucas Saldyt, and Anil Murthy. 2024. Llms can't plan, but can help planning in llm-modulo frameworks. *arXiv preprint arXiv:2402.01817*.

Yash Kumar Lal, Li Zhang, Faeze Brahman, Bodhisattwa Prasad Majumder, Peter Clark, and Niket Tandon. 2024. Tailoring with targeted precision:

- Edit-based agents for open-domain procedure customization. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 15597–15611, Bangkok, Thailand. Association for Computational Linguistics.
- Manling Li, Shiyu Zhao, Qineng Wang, Kangrui Wang, Yu Zhou, Sanjana Srivastava, Cem Gokmen, Tony Lee, Li Erran Li, Ruohan Zhang, et al. 2024. Embodied agent interface: Benchmarking llms for embodied decision making. *arXiv preprint arXiv:2410.07166*.
- Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and Peter Stone. 2023a. Llm+ p: Empowering large language models with optimal planning proficiency. *arXiv preprint arXiv:2304.11477*.
- Hanmeng Liu, Ruoxi Ning, Zhiyang Teng, Jian Liu, Qiji Zhou, and Yue Zhang. 2023b. Evaluating the logical reasoning ability of chatgpt and gpt-4. *arXiv preprint arXiv:2304.03439*.
- Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. Faithful chain-of-thought reasoning. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 305–329, Nusa Dua, Bali. Association for Computational Linguistics.
- Qing Lyu, Li Zhang, and Chris Callison-Burch. 2021. Goal-oriented script construction. In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 184–200, Aberdeen, Scotland, UK. Association for Computational Linguistics.
- Aman Madaan, Shuyan Zhou, Uri Alon, Yiming Yang, and Graham Neubig. 2022. Language models of code are few-shot commonsense learners. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1384–1403, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Christian Muise. 2016. Planning.Domains. In *The 26th International Conference on Automated Planning and Scheduling - Demonstrations*.
- Liangming Pan, Alon Albalak, Xinyi Wang, and William Wang. 2023. Logic-LM: Empowering large language models with symbolic solvers for faithful logical reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3806–3824, Singapore. Association for Computational Linguistics.
- Jendrik Seipp, Álvaro Torralba, and Jörg Hoffmann. 2022. PDDL generators. <https://doi.org/10.5281/zenodo.6382173>.
- Tom Silver, Soham Dan, Kavitha Srinivas, Joshua B Tenenbaum, Leslie Kaelbling, and Michael Katz. 2024. Generalized planning in pddl domains with pretrained large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 20256–20264.
- Kaya Stechly, Karthik Valmeekam, and Subbarao Kambhampati. 2024. Chain of thoughtlessness: An analysis of cot in planning. *arXiv preprint arXiv:2405.04776*.
- Hao Tang, Darren Key, and Kevin Ellis. 2024. World-coder, a model-based llm agent: Building world models by writing code and interacting with the environment. *Preprint*, arXiv:2402.12275.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2024a. Planbench: An extensible benchmark for evaluating large language models on planning and reasoning about change. *Advances in Neural Information Processing Systems*, 36.
- Karthik Valmeekam, Sarath Sreedharan, Matthew Marquez, Alberto Olmo, and Subbarao Kambhampati. 2023. On the planning abilities of large language models (a critical investigation with a proposed benchmark). *Preprint*, arXiv:2302.06706.
- Karthik Valmeekam, Kaya Stechly, and Subbarao Kambhampati. 2024b. Llms still can't plan; can lrms? a preliminary evaluation of openai's o1 on planbench. *Preprint*, arXiv:2409.13373.
- Daniel S Weld. 1999. Recent advances in ai planning. *AI magazine*, 20(2):93–93.
- Lionel Wong, Jiayuan Mao, Pratyusha Sharma, Zachary S Siegel, Jiahai Feng, Noa Korneev, Joshua B Tenenbaum, and Jacob Andreas. 2023. Learning adaptive planning representations with natural language guidance. *arXiv preprint arXiv:2312.08566*.
- Yaqi Xie, Chen Yu, Tongyao Zhu, Jinbin Bai, Ze Gong, and Harold Soh. 2023. Translating natural language to planning goals with large-language models. *arXiv preprint arXiv:2302.05128*.
- Hongming Zhang, Muhao Chen, Haoyu Wang, Yangqiu Song, and Dan Roth. 2020. Analogous process structure induction for sub-event sequence prediction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1541–1550, Online. Association for Computational Linguistics.

Li Zhang, Peter Jansen, Tianyi Zhang, Peter Clark, Chris Callison-Burch, and Niket Tandon. 2024a. PDDLEGO: Iterative planning in textual environments. In *Proceedings of the 13th Joint Conference on Lexical and Computational Semantics (*SEM 2024)*, pages 212–221, Mexico City, Mexico. Association for Computational Linguistics.

Li Zhang, Hainiu Xu, Abhinav Kommula, Chris Callison-Burch, and Niket Tandon. 2024b. OpenPI2.0: An improved dataset for entity tracking in texts. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 166–178, St. Julian’s, Malta. Association for Computational Linguistics.

Li Zhang, Hainiu Xu, Yue Yang, Shuyan Zhou, Weiqiu You, Manni Arora, and Chris Callison-Burch. 2023. Causal reasoning of entities and events in procedural texts. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 415–431, Dubrovnik, Croatia. Association for Computational Linguistics.

Tianyi Zhang, Li Zhang, Zhaoyi Hou, Ziyu Wang, Yuling Gu, Peter Clark, Chris Callison-Burch, and Niket Tandon. 2024c. PROC2PDDL: Open-domain planning representations from texts. In *Proceedings of the 2nd Workshop on Natural Language Reasoning and Structured Explanations (@ACL 2024)*, pages 13–24, Bangkok, Thailand. Association for Computational Linguistics.

Andrew Zhu, Liam Dugan, Alyssa Hwang, and Chris Callison-Burch. 2023. Kani: A lightweight and highly hackable framework for building language model applications. In *Proceedings of the 3rd Workshop for Natural Language Processing Open Source Software (NLP-OSS 2023)*, pages 65–77, Singapore. Association for Computational Linguistics.

Wang Zhu, Ishika Singh, Robin Jia, and Jesse Thomason. 2024. Language models can infer action semantics for classical planners from environment feedback. *arXiv preprint arXiv:2406.02791*.

Max Zuo, Francisco Piedrahita Velez, Xiaochen Li, Michael L Littman, and Stephen H Bach. 2024. Planetarium: A rigorous benchmark for translating text to structured planning languages. *arXiv preprint arXiv:2407.03321*.

A Data Examples

As discussed above, each dataset comes with ground-truth PDDL describing domains ($\mathbb{D}\mathbb{F}$) and problems ($\mathbb{P}\mathbb{F}$). To maximize flexibility when performing analysis, we construct problem instances ourselves for some datasets, so that we can measure complexity with metrics like the number of blocks in BlocksWorld, for which we ensure a uniform distribution to avoid biases. Instances

for BlocksWorld were randomly generated by varying the number of blocks and number of stacks in the initial and goal states from 2 to 15. Instances for Mystery BlocksWorld were randomly sampled from (Valmeekam et al., 2024a). Instances of Logistics were taken directly from (IPC, 1998) and (IPC, 2000). Instances of Barman were generated by using (Seipp et al., 2022) and varying the number of shot-glasses, ingredients and cocktails from 1 to 9. Below are the example PDDL and descriptions for BlocksWorld-100 and MysteryBlocksWorld-100. Due to the length of the examples, the PDDL and natural language descriptions for Logistics-100 and Barman-100 can be found at <https://github.com/CassieHuang22/llm-as-pddl-formalizer/tree/main/examples>.

A.1 BlocksWorld-100 PDDL

Listings 1 and 2 display examples of the ground-truth $\mathbb{D}\mathbb{F}$ and $\mathbb{P}\mathbb{F}$ for BlocksWorld-100.

The $\mathbb{D}\mathbb{F}$ contains all four actions (pickup, put-down, stack and unstack) and their preconditions and post-conditions, as well as predicates needed for the domain.

The $\mathbb{P}\mathbb{F}$ contains the objects, initial state and goal state for the problem.

A.2 BlocksWorld-100 $\mathbb{D}\mathbb{D}$ and $\mathbb{P}\mathbb{D}$

Listings 3 and 4 display example $\mathbb{D}\mathbb{D}$ and $\mathbb{P}\mathbb{D}$ for the Heavily Templated setting in the BlocksWorld-100 dataset. In this setting, all preconditions and post-conditions are written out explicitly and sound similar to PDDL. The $\mathbb{P}\mathbb{D}$ is similar, in that it lists all the predicates needed for to solve the task.

Listings 5 and 6 display example $\mathbb{D}\mathbb{D}$ and $\mathbb{P}\mathbb{D}$ for Moderately Templated data in the BlocksWorld-100 dataset. In this setting, the $\mathbb{D}\mathbb{D}$ and $\mathbb{P}\mathbb{D}$ are much more natural than the Heavily Templated data, but all predicates are still listed.

Finally Listings 7 and 8 display example $\mathbb{D}\mathbb{D}$ and $\mathbb{P}\mathbb{D}$ for Natural BlocksWorld-100 data. For the Natural data, we can see that the $\mathbb{D}\mathbb{D}$ and $\mathbb{P}\mathbb{D}$ give all necessary information to complete the task, but does not sound like PDDL, and does not describe all predicates needed to perform the task.

We can see that the descriptions have all the same components as the $\mathbb{D}\mathbb{F}$ and $\mathbb{P}\mathbb{F}$ in PDDL, but written in different levels of naturalness.

A.3 MysteryBlocksWorld-100 PDDL

Listings 9 and 10 display examples of the ground-truth \mathbb{DF} and \mathbb{PF} for MysteryBlocksWorld-100.

A.4 MysteryBlocksWorld-100 \mathbb{DD} and \mathbb{PD}

Listings 11 and 12 are example \mathbb{DD} and \mathbb{PD} of the Heavily Templated MysteryBlocksWorld-100. Text written in green demonstrates natural sounding text while text written in red demonstrates text that sounds the most like PDDL.

B Prompts

Listings 13, 14, 15 and 16 displays the prompts for all domains given to all models. Whenever possible, we asked the model to return the output in a JSON object for easier parsing.

C Experimental Setup Details

C.1 Planner

We use a dual-bfws-ffparser planner to find plans from the generated PDDL. This is a best first width search planner used with a fast forward planner.

C.2 VAL

The VAL library takes in a ground-truth PDDL \mathbb{DF} , \mathbb{PF} and a plan and tries to execute the plan in the environment by checking whether each action in the found plan can be executed based on the preconditions written in the ground-truth files. If the plan is executable, VAL then checks whether the final state after executing the plan matches the goal state found in the ground-truth \mathbb{PF} . If either the plan is not executable or the final state does not match the goal state, VAL will return an error, therefore the plan found by the planner is not correct.

C.3 Evaluation Metrics

We use solvability and correctness to evaluate the results from the LLMs. Solvability was calculated by running the planner and counting how many plans were returned. If a plan was not returned due to not finding a plan or an error occurred, it did not count towards solvability. Correctness was calculated by counting how many plans returned with no errors after running VAL, the plan was counted if VAL was able to execute the plan and there were no errors, and the final state is equal to the goal state in ground-truth problem file.

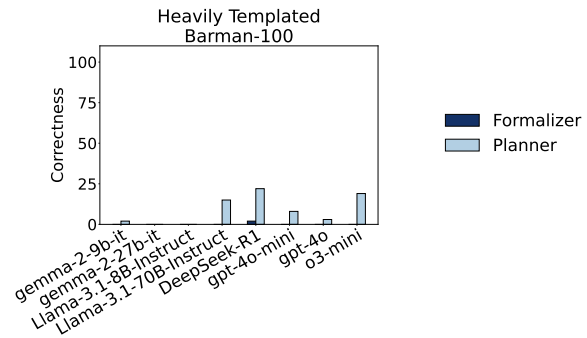


Figure 7: Performance for Barman-100.

D Detailed Results

Beyond the visualizations above, we show the detailed results of all models on all simulations of all naturalness levels.

D.1 Results for BlocksWorld-100

Table 2 displays results for all results for BlocksWorld-100.

D.2 Results for MysteryBlocksWorld-100

Table 3 displays the results for Heavily Templated MysteryBlocksWorld-100.

D.3 Results for Logistics-100

Table 4 displays results for all naturalness settings for Logistics-100.

D.4 Results for Barman-100

Table 5 displays results for Barman-100.

D.5 Few-Shot Prompting Results

As mentioned in Section 4.3, we cannot provide the entire \mathbb{DF} and \mathbb{PF} in the prompt for few-shot prompting as we are asking the LLM to predict the entire \mathbb{DF} and \mathbb{PF} . Therefore, we performed experiments where we provide a template for the \mathbb{DF} and \mathbb{PF} where portions of the PDDL are stubbed out. We performed experiments on Heavily Templated and Natural BlocksWorld-100 on a small subset of models. From Tables 2 and 6, we can see that results are mixed, with performance improving for some models like Llama-3.1-8B-Instruct and worse performance for others like gpt-4o on the Natural BlocksWorld-100 dataset. This suggests few-shot prompting does not improve performance for our experiments.

	Environment	LLM predicts?	Natural Descriptions?
Zuo et al. (2024)	fully-observed	PF	Y
Zhang et al. (2024a)	partially-observed	PF	N
Liu et al. (2023a)	fully-observed	PF	N
Xie et al. (2023)	fully-observed & partially-observed	PF goal	N
Lyu et al. (2023)	fully-observed	PF goal	N
Zhang et al. (2024c)	procedural texts	DF action semantics	N
Wong et al. (2023)	partially-observed	DF	N
Guan et al. (2023)	fully-observed	DF & PF goal	N
Zhu et al. (2024)	fully-observed	DF action semantics	N
Tang et al. (2024)	partially-observed	Python	N/A
Silver et al. (2024)	fully-observed	plan	N
Valmeekam et al. (2024a)	fully-observed	plan	N
Stechly et al. (2024)	fully-observed	plan	N
This work	fully-observed	DF & PF	Y

Table 1: Comparison with related work.

D.6 Chain-of-Thought Prompting Results

For chain-of-thought prompting, we ask the model to first write a domain description, then write the DF, then write a problem description then write the PF. These experiments were performed on Heavily Templated and Natural BlocksWorld-100 on a small subset of models. From Tables 2 and 7, we can see that results are mixed, with performance improvement on some models like gpt-3.5-turbo and performance decline on other models like gpt-4o on the Natural BlocksWorld-100 dataset. This suggests that chain-of-thought prompting does not improve performance for our experiments.

E Detailed Error Analysis

E.1 Syntax and Semantic Errors

Table 8 displays syntax and semantic errors of models on Natural BlocksWorld-100 and Logistics-100. We do not analyze Barman-100 which is too challenging for models.

Table 9 displays fine-grained DF errors of models on Natural BlocksWorld-100 and Logistics-100.

E.2 Sample Model Output

We display an example of the DF and PF that Llama-3.1-8B-Instruct generated. We can see that there are syntax errors, as well as semantic errors in the DF and PF. We give the model the DD and PD in Listing 17 as input, which returned the DF and PF shown in Listings 18 and 19. Text in red displays errors outputted from the model. We

can see that in the DF there are syntax errors (incorrect keyword “preconditions”) as well as semantic errors (incorrect predicates in preconditions and effects). For the output PF there are semantic errors, such as incorrect preconditions (a block cannot be clear and have another block on top of it) in the init section.

F Related Works Comparison

Table 1 compares works related to this paper. We can see that other works as the LLM to predict either the plan, parts of PDDL files and other languages. We can also see that other works have mostly templated natural language descriptions, while this work uses both templated and natural descriptions.

G Methodology Comparison

Both methodologies used in this paper incorporate LLMs into planning. There are pros and cons to each methodology. When using LLM-as-planner, we have a lightweight solution that returns results very quickly. However, due to the lack of reasoning skills in LLMs, they often struggle to create formal plans. Meanwhile using LLM-as-formalizer provides better executability and interpretability, though it uses a solver, which may result in getting results slower. We believe that for performance reasons that LLM-as-formalizer is the superior methodology.

Listing 1: Domain File for BlocksWorld-100

```
(define (domain blocksworld)
  (:predicates (clear ?x)
               (on-table ?x)
               (arm-empty)
               (holding ?x)
               (on ?x ?y))

  (: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))))

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

  (:action stack
   :parameters (?ob ?underob)
   :precondition (and (clear ?underob) (holding ?ob))
   :effect (and (arm-empty) (clear ?ob) (on ?ob ?underob)
               (not (clear ?underob)) (not (holding ?ob))))

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

Listing 2: Problem File for BlocksWorld-100

```
(define (problem blocksworld-p99)
  (:domain blocksworld)
  (:objects red blue green yellow )
  (:init
   (on-table red)
   (on blue red)
   (clear blue)
   (on-table green)
   (on yellow green)
   (clear yellow)
   (arm-empty)
  )
  (:goal (and
   (on-table red)
   (on green red)
   (on yellow green)
   (on blue yellow)
  ))
)
```

Listing 3: Domain Description for Heavily Templated BlocksWorld-100

I am playing with a set of blocks. Here are the actions I can do

Pickup block
 Unstack block from another block
 Putdown block
 Stack block on another block

I have the following restrictions on my actions:

To perform Pickup action, the following facts need to be true: clear block, block on table, arm-empty.

Once Pickup action is performed the following facts will be true: holding block.

Once Pickup action is performed the following facts will be false: clear block, block on table, arm-empty.
 To perform Putdown action, the following facts need to be true: holding block.
 Once Putdown action is performed the following facts will be true: clear block, block on table, arm-empty.
 Once Putdown action is performed the following facts will be false: holding block.
 To perform Stack action, the following needs to be true: clear block2, holding block1.
 Once Stack action is performed the following will be true: arm-empty, clear block1, block1 on block2.
 Once Stack action is performed the following will be false: clear block2, holding block1.
 To perform Unstack action, the following needs to be true: block1 on block2, clear block1, arm-empty.
 Once Unstack action is performed the following will be true: holding block1, clear block2.
 Once Unstack action is performed the following will be false:, block1 on block2, clear block1, arm-empty.

Listing 4: Problem Description for Heavily Templated BlocksWorld-100

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

Listing 5: Domain Description for Moderately Templated BlocksWorld-100

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.
 Once you stack a block on top of a second block, the second block is no longer clear.

Listing 6: Problem Description for Moderately Templated BlocksWorld-100

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

Listing 7: Domain Description for Natural BlocksWorld-100

The Blocksworld game involves a set of blocks of different colors, which can be stacked on top of each other or placed on the table. The objective is to move the blocks from an initial configuration to a goal configuration using a series of legal moves. Legal moves in Blocksworld include: picking up a block from the table or from the top of another block, stacking a block onto the table, or stacking a block onto another block.

Listing 8: Problem Description for Natural BlocksWorld-100

In this particular game, there are 4 blocks: a red block, a blue block, a green block, and a yellow block. At the start, the red block is on the table, the blue block is on top of the red block, the green block is on the table, and the yellow block is on top of the green block. The goal is to have the red block on the table, the green block on top of the red block, the yellow block on top of the green block, and the blue block on top of the yellow block.

Listing 9: Domain File for Mystery_BlocksWorld-100

```
(define (domain mystery_blocksworld)
  (:predicates (province ?x)
               (planet ?x)
               (harmony)
               (pain ?x)
               (craves ?x ?y))

  (:action attack
   :parameters (?ob)
   :precondition (and (province ?ob) (planet ?ob) (harmony))
   :effect (and (pain ?ob) (not (province ?ob)) (not (planet ?ob))
                (not (harmony))))

  (:action succumb
   :parameters (?ob)
   :precondition (pain ?ob)
   :effect (and (province ?ob) (harmony) (planet ?ob)
                (not (pain ?ob))))

  (:action overcome
   :parameters (?ob ?underob)
   :precondition (and (province ?underob) (pain ?ob))
   :effect (and (harmony) (province ?ob) (craves ?ob ?underob)
                (not (province ?underob)) (not (pain ?ob))))

  (:action feast
   :parameters (?ob ?underob)
   :precondition (and (craves ?ob ?underob) (province ?ob) (harmony))
   :effect (and (pain ?ob) (province ?underob)
                (not (craves ?ob ?underob)) (not (province ?ob)) (not (harmony))))))
```

Listing 10: Problem File for Mystery_BlocksWorld-100

```
(define (problem mystery_blocksworld-p01)
  (:domain mystery_blocksworld)
  (:objects a b c d)
  (:init
   (craves a b)
   (craves b c)
   (harmony)
   (planet c)
   (planet d)
   (province a)
   (province d)
  )
  (:goal (and
         (craves a d)
         (craves c a)
  ))
)
```

Listing 11: Domain Description for Mystery_BlocksWorld-100

I am playing with a set of green objects. Here are the actions I can do

Attack object
Feast object from another object
Succumb object
Overcome object from another object

I have the following restrictions on my actions:

To perform Attack action, the following facts need to be true: Province object, Planet object, Harmony.

Once Attack action is performed the following facts will be true: Pain object.

Once Attack action is performed the following facts will be false: Province object, Planet object, Harmony.

To perform Succumb action, the following facts need to be true: Pain object.

Once Succumb action is performed the following facts will be true: Province object, Planet object, Harmony.

Once Succumb action is performed the following facts will be false: Pain object.

To perform Overcome action, the following needs to be true: Province other object, Pain object.

Once Overcome action is performed the following will be true: Harmony, Province object, Object Craves other object.

Once Overcome action is performed the following will be false: Province other object, Pain object.

To perform Feast action, the following needs to be true: Object Craves other object, Province object, Harmony.

Once Feast action is performed the following will be true: Pain object, Province other object.

Once Feast action is performed the following will be false:, Object Craves other object, Province object, Harmony.

Listing 12: Problem Description for Mystery_BlocksWorld-100

As initial conditions I have that, object a craves object b, object b craves object c, harmony, planet object c, planet object d, province object a and province object d.

My goal is to have that object a craves object d and object c craves object a.

Listing 13: Prompt for LLM-as-Planner on BlocksWorld-100

Here is a game we are playing.

{domain_description}

{problem_description}

Write the plan that would solve this problem.

These are the available actions:

(PICK-UP block): pick up a block from the table

(PUT-DOWN block): put down a block on the table

(STACK block1 block2): stack block1 onto block2

(UNSTACK block1 block2): unstack block1 from block2

Here is what the output should look like:

(PICK-UP A)

(STACK A B)

(UNSTACK A B)

(PUT-DOWN A)

Listing 14: Prompt for LLM-as-Planner on Mystery_BlocksWorld-100

```
Here is a game we are playing.

{domain_description}

{problem_description}

Write the plan that would solve this
problem.
These are the available actions:
(ATTACK object): attack object
(SUCCUMB object): succumb
(OVERCOME object1 object2): overcome object1 from object2
(FEAST object1 object2): feast object1 from object2

Here is what the output should look like:
(ATTACK A)
(OVERCOME A B)
(FEAST A B)
(SUCCUMB A)
```

Listing 15: Prompt for LLM-as-Planner on Logistics-100

```
Here is a game we are playing.

{domain_description}

{problem_description}

Write the plan that would solve this
problem.
These are the available actions:
(load-truck package truck location): load a package onto a truck at a location
(load-airplane object airplane location): load a package onto an airplane at a location
(unload-truck package truck location): unload a package from a truck at a location
(unload-airplane package airplane location): unload a package from an airplan at a location
(drive-truck truck location1 location2 city): drive a truck from location1 to location2 in a city
(fly-airplane airplane location1 location2): fly an airplane from location1 to location2

Here is what the output should look like:
(load-truck package truck city1-1)
(drive-truck truck city1-1 city1-2 city1)
(unload-truck package truck city1-2)
(load-airplane package plane city1-2)
(fly-airplane plane city1-2 city2-1)
(unload-airplane package plane city2-1)
```

Listing 16: Prompt for LLM-as-Planner on Barman-100

Here is a game we are playing.

{domain_description}

{problem_description}

Write the plan that would solve this problem.

These are the available actions:

(GRASP hand container): grasp container with hand

(LEAVE hand container): leave container with hand

(FILL-SHOT shot ingredient hand1 hand2 dispenser): fill shot with ingredient from dispenser using hand1 and hand2

(REFILL-SHOT shot ingredient hand1 hand2 dispenser): re-fill shot with ingredient from dispenser using hand1 and hand2

(EMPTY-SHOT hand shot beverage): empty shot containing beverage with hand

(CLEAN-SHOT shot beverage hand1 hand2): clean shot containing beverage using hand1 and hand2

(POUR-SHOT-TO-CLEAN-SHAKER shot ingredient shaker hand level level1): pour shot containing ingredient into clean shaker using hand so that the level in the shaker changes from level to level1

(POUR-SHOT-TO-USED-SHAKER shot ingredient shaker hand level level1): pour shot containing ingredient into used shaker using hand so that the level in the shaker changes from level to level1

(EMPTY-SHAKER hand shaker cocktail level level1): empty shaker containing cocktail with hand so that the level in the shaker changes from level to level1

(CLEAN-SHAKER hand1 hand2 shaker): clean shaker using hand1 and hand2

(SHAKE cocktail ingredient1 ingredient2 shaker hand1 hand2): shake shaker containing cocktail of ingredient1 and ingredient2 using hand1 and hand2

(POUR-SHAKER-TO-SHOT beverage shot hand shaker level level1): pour shaker containing beverage into shot using hand so that the level in the shaker changes from level to level1

Here is what the output should look like:

(GRASP right shot1)

(FILL-SHOT shot1 ingredient10 right left dispenser10)

(POUR-SHOT-TO-CLEAN-SHAKER shot1 ingredient10 shaker1 right l0 l1)

(CLEAN-SHOT shot1 ingredient10 right left)

(FILL-SHOT shot1 ingredient5 right left dispenser5)

(GRASP left shaker1)

(POUR-SHOT-TO-USED-SHAKER shot1 ingredient5 shaker1 right l1 l2)

(LEAVE right shot1)

(SHAKE cocktail1 ingredient5 ingredient10 shaker1 left right)

(LEAVE left shaker1)

(GRASP left shot1)

(CLEAN-SHOT shot1 ingredient5 left right)

(GRASP right shaker1)

(POUR-SHAKER-TO-SHOT cocktail1 shot1 right shaker1 l2 l1)

Models	Metrics					
	Natural		Moderately Templated		Heavily Templated	
	Solvability	Correctness	Solvability	Correctness	Solvability	Correctness
gemma-2-9b-it	3/100	3/100	0/100	0/100	61/100	10/100
gemma-2-9b-it ^P	-	9/100	-	5/100	-	5/100
gemma-2-27b-it	0/100	0/100	17/100	10/100	81/100	80/100
gemma-2-27b-it ^P	-	11/100	-	3/100	-	7/100
Llama-3.1-8B	0/100	0/100	0/100	0/100	0/100	0/100
Llama-3.1-8B ^P	-	1/100	-	1/100	-	0/100
Llama-3.1-70B	0/100	0/100	0/100	0/100	0/100	0/100
Llama-3.1-70B ^P	-	13/100	-	10/100	-	10/100
DeepSeek-R1-Distill-Llama-8B	0/100	0/100	0/100	0/100	0/100	0/100
DeepSeek-R1-Distill-Llama-8B ^P	-	2/100	-	2/100	-	2/100
DeepSeek-R1-Distill-Llama-70B	0/100	0/100	0/100	0/100	0/100	0/100
DeepSeek-R1-Distill-Llama-70B ^P	-	18/100	-	25/100	-	22/100
DeepSeek-R1	83/100	81/100	76/100	75/100	84/100	84/100
DeepSeek-R1 ^P	-	91/100	-	85/100	-	86/100
gpt-3.5-turbo	2/100	1/100	14/100	4/100	39/100	29/100
gpt-4o-mini	19/100	3/100	9/100	5/100	66/100	59/100
gpt-4o-mini ^P	-	7/100	-	7/100	-	1/100
gpt-4o	64/100	60/100	77/100	67/100	89/100	89/100
gpt-4o ^P	-	33/100	-	35/100	-	29/100
o1-preview	91/100	91/100	-	-	-	-
o1-preview ^P	-	82/100	-	-	-	-
o3-mini	79/100	68/100	82/100	70/100	94/100	94/100
o3-mini ^P	-	87/100	-	87/100	-	96/100

Table 2: Performance of LLM-as-formalizer and LLM-as-planner (^P) on all naturalness settings of the BlocksWorld-100 dataset.

Models	Solvability	Correctness
Llama-3.1-8B	0/100	0/100
Llama-3.1-8B ^P	-	0/100
Llama-3.1-70B	0/100	0/100
Llama-3.1-70B ^P	-	0/100
gemma-2-9b-it	100/100	99/100
gemma-2-9b-it ^P	-	9/100
gemma-2-27b-it	99/100	98/100
gemma-2-27b-it ^P	-	0/100
DeepSeek-R1	56/100	54/100
DeepSeek-R1 ^P	-	21/100
gpt-3.5-turbo	4/100	0/100
gpt-4o-mini	36/100	5/100
gpt-4o-mini ^P	-	0/100
gpt-4o	74/100	70/100
gpt-4o ^P	-	0/100
o3-mini	95/100	95/100
o3-mini ^P	-	74/100

Table 3: Performance of LLM-as-formalizer and LLM-as-planner (^P) on the Heavily Templated MysteryBlocksWorld-100 dataset.

Metrics						
Models	Natural		Moderately Templated		Heavily Templated	
	Solvability	Correctness	Solvability	Correctness	Solvability	Correctness
gemma-2-9b-it	1/100	0/100	0/100	0/100	1/100	0/100
gemma-2-9b-it ^P	-	0/100	-	0/100	-	0/100
gemma-2-27b-it	1/100	0/100	3/100	0/100	0/100	0/100
gemma-2-27b-it ^P	-	0/100	-	0/100	-	1/100
Llama-3.1-8B	0/100	0/100	0/100	0/100	0/100	0/100
Llama-3.1-8B ^P	-	1/100	-	0/100	-	0/100
Llama-3.1-70B	0/100	0/100	0/100	0/100	0/100	0/100
Llama-3.1-70B ^P	-	0/100	-	0/100	-	0/100
DeepSeek-R1	51/100	9/100	73/100	73/100	64/100	64/100
DeepSeek-R1 ^P	-	12/100	-	22/100	-	22/100
gpt-3.5-turbo	1/100	0/100	0/100	0/100	1/100	1/100
gpt-3.5-turbo ^P	-	0/100	-	0/100	-	0/100
gpt-4o-mini	13/100	0/100	29/100	0/100	26/100	0/100
gpt-4o-mini ^P	-	0/100	-	0/100	-	0/100
gpt-4o	20/100	2/100	34/100	13/100	33/100	13/100
gpt-4o ^P	-	1/100	-	0/100	-	0/100
o3-mini	43/100	7/100	51/100	39/100	55/100	47/100
o3-mini ^P	-	14/100	-	21/100	-	21/100

Table 4: Performance of LLM-as-formalizer and LLM-as-planner (^P) on all Logistics-100 datasets.

Models	Solvability	Correctness
Llama-3.1-8B	0/100	0/100
Llama-3.1-8B ^P	-	0/100
Llama-3.1-70B	0/100	0/100
Llama-3.1-70B ^P	-	15/100
gemma-2-9b-it	0/100	0/100
gemma-2-9b-it ^P	-	2/100
gemma-2-27b-it	0/100	0/100
gemma-2-27b-it ^P	-	0/100
DeepSeek-R1	10/100	2/100
DeepSeek-R1 ^P	-	22/100
gpt-3.5-turbo	0/100	0/100
gpt-3.5-turbo ^P	-	4/100
gpt-4o-mini	0/100	0/100
gpt-4o-mini ^P	-	8/100
gpt-4o	1/100	0/100
gpt-4o ^P	-	3/100
o3-mini	1/100	0/100
o3-mini ^P	-	19/100

Table 5: Performance of LLM-as-formalizer and LLM-as-planner (^P) on the Heavily Templated Barman-100 dataset.

Metrics				
Models	Natural		Heavily Templated	
	Solvability	Correctness	Solvability	Correctness
Llama-3.1-8B	1/100	1/100	49/100	45/100
Llama-3.1-70B	0/100	0/100	0/100	0/100
gpt-3.5-turbo	8/100	1/100	35/100	24/100
gpt-4o-mini	13/100	3/100	75/100	61/100
gpt-4o	60/100	48/100	93/100	92/100

Table 6: Performance of LLM-as-formalizer all the Natural and Heavily Templated BlocksWorld-100 datasets using few-shot prompting.

Metrics				
Models	Natural		Heavily Templated	
	Solvability	Correctness	Solvability	Correctness
Llama-3.1-8B	0/100	0/100	0/100	0/100
Llama-3.1-70B	0/100	0/100	0/100	0/100
gpt-3.5-turbo	6/100	0/100	23/100	17/100
gpt-4o-mini	19/100	4/100	49/100	45/100
gpt-4o	61/100	48/100	91/100	90/100

Table 7: Performance of LLM-as-formalizer all the Natural and Heavily Templated BlocksWorld-100 datasets using chain-of-thought prompting.

Models	Syntax Error	DF Error	PF Error
Natural BlocksWorld-100			
gemma-2-9b-it	15/20	20/20	20/20
gemma-2-27b-it	3/20	20/20	14/20
Llama-3.1-8B	20/20	20/20	18/20
Llama-3.1-70B	20/20	20/20	17/20
gpt-4o-mini	2/20	20/20	19/20
gpt-4o	2/20	2/20	18/20
Natural Logistics-100			
gemma-2-9b-it	7/20	20/20	15/20
gemma-2-27b-it	8/20	20/20	20/20
Llama-3.1-8B	20/20	20/20	20/20
Llama-3.1-70B	20/20	20/20	10/20
gpt-4o-mini	2/20	20/20	19/20
gpt-4o	5/20	20/20	19/20
MysteryBlocksWorld-100			
gpt-4o-mini	6/20	20/20	1/20
gpt-4o	5/20	16/20	0/20

Table 8: Annotated types of 20 randomly sampled errors of LLM-as-formalizer for various models and datasets.

Models	Wrong Precondition	Wrong Effect	Missing Predicate	Missing Action	Missing Parameters
Natural BlocksWorld-100					
gpt-4o-mini	11/20	18/20	19/20	1/20	2/20
gpt-4o	0/20	2/20	0/20	0/20	0/20
Natural Logistics-100					
gpt-4o-mini	20/20	16/20	20/20	5/20	17/20
gpt-4o	17/20	9/20	19/20	1/20	9/20
Heavily Templated MysteryBlocksWorld-100					
gpt-4o-mini	14/20	17/20	17/20	0/20	5/20
gpt-4o	13/20	14/20	0/20	0/20	2/20

Table 9: Analysis of errors found in DF for Natural BlocksWorld-100 and Heavily Templated MysteryBlocksWorld-100 out of 20 randomly sampled instances.

Listing 17: Example Domain Description and Problem Description given to Llama-3.1-8B-Instruct

```
;Domain Description:
Blocksworld is a puzzle domain where the aim is to move a set of blocks to achieve a specific goal
configuration. The domain consists of a flat surface, often called the 'table', and an arbitrary
number of blocks. Every block can be stacked on top of another block or placed on the table.
Blocks can be moved one at a time, and only the block at the top of a stack can be moved. The
main actions available in Blocksworld are:

1. Pick Up (Pick-Up B): Picks up block B from the table.
2. Put Down (Put-Down B): Places block B onto the table.
3. Stack (Stack B1 B2): Stacks block B1 on top of block B2.
4. Unstack (Unstack B1 B2): Removes block B1 from the top of block B2.

The goal is to transform an initial configuration of blocks into a desired target configuration using
these actions.

;Problem Description:
In this Blocksworld problem, we will consider a simple scenario with 2 blocks: Block1 and Block2. The
initial state of the blocks is as follows:

Initial State:
- Block1 is on the table.
- Block2 is on Block1.

The goal is to achieve the following configuration:

Goal State:
- Block2 is on the table.
- Block1 is on the table.

To solve this problem, you will need to use the available actions to rearrange the blocks from the
initial state to the goal state.
```

Listing 18: Generated Domain File from Llama-3.1-8B-Instruct

```
(define (domain blocksworld)
  (:requirements :strips :typing)
  (:types block)
  (:predicates (on-table?b - block) (on?b1 - block?b2 - block) (clear?b - block) )
  (:action pick-up
    :parameters (?b - block)
    :preconditions (and (on-table?b) (clear?b))
    :effects (and (not (on-table?b)) (not (clear?b)) (on-table?b))
  )
  (:action put-down
    :parameters (?b - block)
    :preconditions (and (not (on-table?b)) (clear?b))
    :effects (and (on-table?b) (clear?b))
  )
  (:action stack
    :parameters (?b1 - block?b2 - block)
    :preconditions (and (on-table?b1) (clear?b2))
    :effects (and (not (on-table?b1)) (not (clear?b2)) (on?b1?b2))
  )
  (:action unstack
    :parameters (?b1 - block?b2 - block)
    :preconditions (and (on?b1?b2) (clear?b1))
    :effects (and (on-table?b1) (not (on?b1?b2)) (clear?b2))
  )
)
```


Listing 19: Generated Problem File from Llama-3.1-8B-Instruct

```
(define (problem blocksworld-problem)
  (:domain blocksworld)
  (:objects block1 block2 - block)
  (:init
    (on-table block1)
    (on block2 block1)
    (clear block1)
    (clear block2)
  )
  (:goal
    (and (on-table block2)
          (on-table block1)
          (clear block1)
          (clear block2)
    )
  )
)
```