On Planning Through LLMs

Anonymous submission

Abstract

In recent years, various studies have been carried out to assess whether Large Language Models (LLMs) possess different reasoning capabilities, including those required in automated planning. Typically, these studies provide the LLM with a planning domain and a problem, specified by an initial state and a goal, and require the LLM model to generate a plan solving the problem. Despite this common configuration, such studies significantly differ in the used models, the information provided to the model, the possible involvement of symbolic planners, and the experimental approaches used for the evaluation. Motivated by the growing interest in LLMs and in the understanding of their reasoning abilities, in this work we offer a concise review of recent studies on using LLMs for planning. We outline the main research trends and discuss their most notable findings. Furthermore, we identify key challenges and highlight critical aspects to consider when evaluating a LLM in terms of learning to plan and generating solution plans.

Introduction

The remarkable success obtained by pre-trained Large Language Models (LLMs) based on the Transformer architecture (Vaswani et al. 2017), such as the GPT models developed by OpenAI (Radford et al. 2018), have been opening new research lines which aim to understand the capabilities of such models. Due to the vast amount of textual data used in their training, they possess huge knowledge about real-world entities in subjects like geography or history (Jiang et al. 2020). Additionally, they can perform simple lexical operations (Madasu and Srivastava 2022), some common sense reasoning (Geva et al. 2021), and solve mathematical problems (Wei et al. 2022b). However, the claim that LLMs possess genuine reasoning abilities remains a subject of intense debate within the scientific community. Some studies, such as (Bender et al. 2021; Zhang et al. 2023; Hicks, Humphries, and Slater 2024), argue that LLMs often simulate reasoning in constrained domains by leveraging statistical patterns in the data, allowing them to solve reasoning tasks without truly "understanding" them.

As a result, there has been a growing interest in studying LLMs using more rigorous and challenging tasks, such as those found in automated planning. Solving automated planning tasks involves understanding complex relationships between entities and objects, determining when actions can be executed, comprehending their consequences, and organizing actions to achieve a specific goal. These capabilities are interesting not only from a theoretical point of view. For instance, consider the problem of finding optimal plans. Reasoning-based planners are typically slow, whereas LLMs could generate plans within seconds. Moreover, although LLMs (like all machine learning models) do not offer formal guarantees of always providing a correct solution, they hold the potential for integration with traditional planners to enhance performance. This approach mirrors existing efforts where machine learning systems calculate heuristics that are subsequently used by planners (Karia and Srivastava 2021; Shen, Trevizan, and Thiébaux 2020), or where deep neural networks are applied for generalized planning (Toyer et al. 2020) and goal recognition (Chiari et al. 2023).

In recent years, several studies have explored the application of LLMs in automated planning (Pallagani et al. 2023b; Silver et al. 2022, 2024; Valmeekam et al. 2022). Despite the common goal, these studies vary in several key aspects, such as the LLM they consider and how that model is used: whether it is exploited through zero-shot or few-shot prompting (Wei et al. 2022a), with a Chain-of-Thought (Wei et al. 2022b), through a more complex fine-tuning process (Pallagani et al. 2023b), or even training a GPT model specifically for planning (Rossetti et al. 2024b). Most notably, the results obtained, how they were evaluated, and the conclusions that can be drawn from them differ significantly. While the authors of (Valmeekam et al. 2022) claim that "LLMs still can't plan", more promising results were obtained by (Pallagani et al. 2023b; Hazra, Martires, and Raedt 2024; Rossetti et al. 2024b).

This paper aims to review and survey these studies, offering guidelines to contextualize their findings and highlighting their similarities and differences. We also discuss the methodologies used for evaluation and the achieved results. Finally, we examine the primary challenges and future directions in this emerging field, considering perspectives from both automated planning and deep learning.

Background

In this section, we provide a brief overview and the background on Large Language Models and Automated Planning.

Transformer-based models and LLMs

In 2017, Vaswani et al. propose a deep learning architecture called **Transformer** (Vaswani et al. 2017). Although it was originally conceived for machine translation, in the following years Transformer-based architecture (such as GPT) became the state-of-the-art in most Natural Language Processing (NLP) tasks.

A transformer is made by two main parts: the Encoder, which converts a text sequence into an embedded representation, and the Decoder, which exploits the embedded representation in an auto-regressive procedure to generate the translated sentence iteratively word by word. Both these parts are made by a stack of several layers (for instance, 12 in the original Transformer, 96 for GPT-3).

A transformer processes text sequences first by separating them into smaller units called tokens (words or parts of words). Different NLP models use different tokenization methods and algorithms, such as Byte-Pair Encoding and WordPiece. After tokenization, an embedding layer converts each token into a corresponding real-valued vector. Therefore, an input text is transformed into a sequence of vectors.

The most important component in the Transformer architecture is the self-attention mechanism introduced in (Vaswani et al. 2017). Intuitively, the self-attention allows to "pay attention" to different parts of the sentence and to incorporate this information into the embedded representation of each token. More formally, first the model projects the embedded representation of each word E into three new representations called *key* (K), *query* (Q) and *value* (V) by multiplying it with three weight matrices W_k, W_q and W_v. The new representation Z, is then calculated as $Z = softmax(\frac{QK^T}{\sqrt{d_k}})V$, where d_k is the size of the embedded representation. Transformer architectures combine several parallel self-attention mechanisms (in the so called Multi-Head Attention) with feed-forward neural layers and residual connections across all the encoding and decoding layers.

The works in (Devlin et al. 2019) and (Radford et al. 2018) derived two models based on the two main parts of the Transformer architecture: **Bidirectional Encoder Representations from Transformers** (BERT), which is based on the Encoder, (Devlin et al. 2019), and **Generative Pretrained Transformer** (GPT) (Radford et al. 2018), which is based on the Decoder. BERT, as other Encoder-based models, is typically trained with *Masked Language Modeling*, i.e. by masking several words in input and having the model predict them. In contrast, GPT (which is the most famous Decoder-based model) is trained to generate text starting from an initial textual prefix. This procedure is also called *Causal Language Modeling*.

Training, Fine-tuning and Prompting Transformer-based models are typically pre-trained on large corpora of unlabeled text to understand and generate language, enabling them to acquire vast knowledge. These models are referred to as pre-trained language models. Due to their scale, often comprising billions of parameters, they are also known as large language models (LLMs).

Pre-trained LLM's capabilities can be further enhanced by refining them on a smaller, task-specific dataset through a process called *fine-tuning*. Fine-tuning enables the model to specialize its general knowledge and better manage the complexities of a specific task. This is typically performed by attaching a simple feed-forward neural network to the model. This neural network is devoted to solve the given task and during its training the LLM weights are updated accordingly. Moreover, the performance of LLMs can be improved through the instruction tuning process, which also exploits reinforcement learning and human feedback(Ouyang et al. 2022). Alternatively, it is possible to train a language model from scratch in order to perform a specific task (Chen et al. 2021). However, this comes with the limitation that the model is restricted to perform only the specific task it was trained for, lacking the wide-ranging knowledge that pre-trained LLMs possess.

For Decoder-based models, and in particular GPT-3 and GPT-4 models, users can interact with LLMs by asking something in natural language and receiving the answer provided by the model. The request made to the model is typically called *prompt* and the overall process of interacting with the LLM in different ways is called *prompt engineering* (Liu et al. 2023b). More specifically, there are three main prompting approaches:

- **Zero-Shot**, in which the LLM is asked something without any examples;
- **Few-Shot**, in which the LLM is provided with some examples which can be used for understanding a more general strategy to address the user's request;
- **Chain-of-Thought** (CoT), in which there are several interactions between the LLM and the user, which may correct mistakes and provide useful information progressively.

Classical Planning

We assume that the reader is familiar with the standard planning language PDDL (Ghallab et al. 1998) for representing deterministic, fully observable planning problems.

A classical planning problem is a pair $P = (D, I)$ where D is a planning domain and I is a problem instance. The planning domain D contains a set of predicate symbols p and a set of action schemas with preconditions and effects given by atoms $p(x_1, ..., x_k)$ where each x_i is an argument of the schema. The problem instance is a tuple $I = (O, Init, Goal)$ where O is a (finite) set of objects names c_i, and $Init$ and $Goal$ are sets of ground atoms $p(c_1, ..., c_k)$ representing the initial state and the goal of the problem. A classical problem $P = (D, I)$ encodes a state model $S(P) = (S, s_0, S_G, Act, A, f)$ where each state $s \in S$ is a set of ground atoms from P, s_0 is the initial state $Init$, S_G is the set of goal states $s \in S$ such that $Goal \subseteq s$, Act is the set of ground actions in P, $A(s)$ is the set of ground actions whose preconditions are true in s, and f is the transition function so that $f(a, s)$ for $a \in A(s)$ represents the state resulting from applying action a to state s. An action sequence $a_0, ..., a_n$ is applicable in P if $a_i \in A(s_i)$ and $s_{i+1} = f(a_i, s_i)$, for $i = 0, ..., n$, and it is a plan if

$s_{n+1} \in S_G$. The cost of a plan is assumed to be given by its length, and a plan is optimal if there is no shorter plan.

Literature selection

We conducted a comprehensive survey of the existing literature to explore the recent trend of using language models as tools for automated plan generation. This research identified 18 published papers that present different approaches to leverage language models for plan generation within automated planning. Our search spanned multiple academic databases, conferences, and journals, using keywords such as *LLM in planning*, *Reasoning with LLM*, *Heuristics with LM*, and *Transformer*, among others. From the initial pool of results, we carefully selected papers formally published in journals, conferences, or workshops or gained substantial citations on platforms like arXiv. The selection process was rigorous and focused on studies where neural language models are applied to generate plans in automated planning tasks. Differently from (Pallagani et al. 2024), which proposed an overview on LLMs in many different contexts (such as translating natural language to PDDL and creating domain models (Oswald et al. 2024), robotics, or coordinating different planning tools), the present study focuses only on how Language Models can be exploited for plan generation. Therefore, we excluded works where LLMs are used for performing some kind of activities outside the symbolic domains of automated planning, such as reinforcement learning (Chen et al. 2021) or robotics applications (Singh et al. 2023). Over the past three years, most of the papers we reviewed have been presented at prominent conferences in artificial intelligence (IJCAI, AAAI), deep learning (NIPS, ICML, EMNLP), automated planning (ICAPS), and robotics (ICRA) highlighting the growing interest in LLMs and their potential reasoning capabilities in this field.

Overview of LLM-based methods for Generating Plans

First, we present a general description of the claims and contributions from the articles we reviewed.

A significant line of work in the examined literature concerns the prompting capabilities of pre-trained LLMs in terms of reasoning and planning without external validation. (Valmeekam et al. 2022) analyzed GPT-3, and proposed a benchmark to address these experiments in a thorough way (Valmeekam et al. 2023a). In these works, planning problems are provided to an LLM (in the form of natural language or PDDL), and the study focuses on evaluating their behaviour under different conditions and prompts. As summed up in (Valmeekam et al. 2022), the takeaway of these works is that LLMs by themselves cannot solve planning tasks.

Nonetheless, several studies have attempted to leverage the basic reasoning capabilities of LLMs and to refine them for planning tasks in more sophisticated ways. These approaches include using chain-of-thought prompting combined with LLM-based validators (Stechly, Valmeekam, and Kambhampati 2024) and symbolic validators (Guan et al. 2023; Kambhampati et al. 2024; Zhou et al. 2024) or integrating search algorithms and planners (Valmeekam et al. 2023b; Hao et al. 2023; Silver et al. 2022; Liu et al. 2023a). Another interesting approach is presented in (Silver et al. 2024), in which GPT-4 is used to generate Python programs that can solve various planning problems in the same domain. The results across these studies vary significantly, leading to differing conclusions. Some works express optimism about LLMs' potential in planning, such as (Hao et al. 2023; Silver et al. 2024; Zhou et al. 2024), while others, like (Kambhampati et al. 2024), offer a more cautious and pessimistic outlook on their effectiveness in these tasks.

Most of the studies mentioned above focus on closed-source, pre-trained models like GPT-3.5 and GPT-4, which have gained significant attention due to their high relevance and commercial success. However, a major limitation of using such models is the challenge of adapting them to specific tasks like automated planning. To address this, several studies have fine-tuned or even trained from scratch smaller Transformer-based models, aiming to create specialized language models tailored to solving planning tasks. One of the first specialized models is Plansformer (Pallagani et al. 2023b), which utilizes a fine-tuned T5 model (Raffel et al. 2020) and generates a corresponding solution plan given a formalized planning problem in PDDL. PlanGPT (Rossetti et al. 2024b) achieved better results by training a GPT-2 model from scratch, and integrating a validator (Rossetti et al. 2024a). A more complex configuration has been studied in (Hazra, Martires, and Raedt 2024) by using three different Language Models to generate and evaluate actions. Another interesting approach is presented in (Lehnert et al. 2024), in which the authors train a Transformer to emulate the A* algorithm. Finally, the works in (Hirsch, Uziel, and Anaby-Tavor 2024) focuses on generating valid actions and heuristic values with LLMs.

In the following sections, we propose additional categorizations of these works, examining key aspects such as their input and output, how they utilize LLMs, what LLMs they used, and whether they incorporate validators or planners into their analyses. For clarity of the pictures, in Figures 1,2 and 3 we named the considered papers with the first three letters of the first author's surname, followed by the last two digits of the publication year. For instance, (Guan et al. 2023) is reported as [GUA23].

LLM Procedures to Solve Planning Problems

In this section, we provide a more detailed description on how LLMs are exploited to solve planning problems. First, we discuss their input and output. Next, we analyze the procedures used for interacting with the LLM and the results.

Input and Output

As illustrated on the x-axis of Figure 1, there are three main ways to provide an automated planning problem as input to a LLM: the first is using **Natural Language** (**NL**), typically English without an explicit logical formalism; the second is using **PDDL**; the third (**Hybrid**) is through a combination of PDDL and prompts based on natural language.

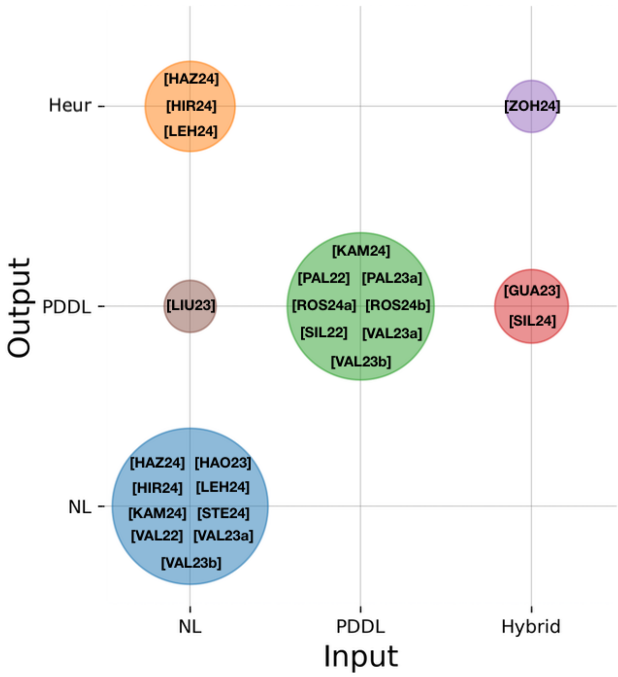

Figure 1: Visualization of the input categories (on the x-axis) and output categories (on the y-axis) of the considered LLM approaches for automated planning. Each work is identified by the first three letters of the first author's surname, followed by the last two digits of the publication year. NL stands for Natural Language, Heur stands for Heuristics.

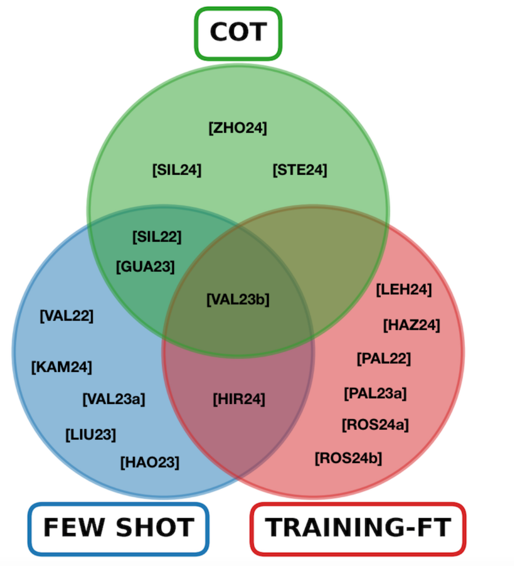

Figure 2: Venn Diagram of the different ways to exploit LLMs for generating plans: Zero and Few-Shot Prompting, in blue, Training and Fine-Tuning, in red and Chain-of-Thought, in green.

More specifically, the approaches based on Natural Language (Valmeekam et al. 2022, 2023b) translate a PDDL domain and a specific problem into a series of simple sentences understandable by a LLM using a custom domain specific translator. Other approaches, such as (Pallagani et al. 2023b; Rossetti et al. 2024b), work directly with PDDL, focusing on processing formalized planning problems based on their logical structure. Hybrid approaches typically use the PDDL formalism combined with natural language to interact with the model via external verifiers and assist the generation (Silver et al. 2024; Guan et al. 2023). Notably, the work by Silver et al. employs a hybrid input strategy, where the system receives not only PDDL but also code fragments and error traces.

In terms of the output generated by the LLM, as shown on the y-axis of Figure 1, the works we considered can be divided into three main categories. The first is composed by works (such as (Valmeekam et al. 2023b; Kambhampati et al. 2024)) in which the LLM generates an answer written in natural language (NL) from which a PDDL plan is derived. Similarly, the second category is composed by works in which the LLM directly generates a valid sequence of actions to solve problem in PDDl, such as (Guan et al. 2023; Liu et al. 2023a). A notable alternative approach is presented in (Silver et al. 2024), where the LLM generates a Python program that includes a policy designed to solve the planning problem writing the plan in PDDL. In the third category, called Heuristics in Figure 1, LLMs are used to compute heuristic values, which can then be combined with search algorithms to generate valid plans. For instance, in (Hirsch, Uziel, and Anaby-Tavor 2024) the authors train the decoder component of a T5 model to predict heuristic values, while Hao et al. (Hao et al. 2023) calculates heuristic values by combining the action's likelihood with the LLM's confidence. A hybrid approach is adopted by Hazra, Martires, and Raedt (Hazra, Martires, and Raedt 2024), where knowledge learned by two LLMs, focused on action applicability and best action selection, guides a beam search carried out by a third pre-trained LLM, integrating multiple models to enhance the planning process.

Procedures and Takeaways

Another critical aspect of the considered works is how LLMs are exploited to solve planning problems, and the results they obtained. As we show in Figure 2, we identified three main strategies attempted: **Zero and Few-Shot** prompting, **Chain-of-thought** (CoT) prompting, and **Fine-tuning and Training**.

In the studies belonging to the first category, the authors analyze the reasoning capabilities of pre-trained LLMs using a prompting setting without external validation. They show that LLMs possess poor planning abilities, even when handling simple problems, as they struggle to reason about action applicability and its effects, generating invalid plans. These capabilities do not increase by providing examples (Valmeekam et al. 2022).

For the second category, some researches have proposed enhancing LLM capabilities by making them process the problem step-by-step, as in the Chain-of-Thought, or even

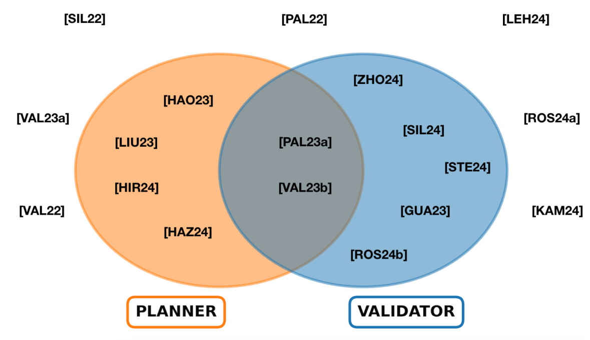

Figure 3: Venn Diagram of the integration with Validators (reported in orange) and Planners (reported in blue). The approaches that do not integrate with either a planner or a validator are reported outside the two sets.

combining them with reasoning tools, for instance the VAL validator (Howey, Long, and Fox 2004) or a Python interpreter (Silver et al. 2024). Instead, the work in (Stechly, Valmeekam, and Kambhampati 2024) utilizes another LLM in a sort of self-verification tool to detect and correct errors. The use of external validators has yielded mixed results. The approach has produced excellent outcomes in (Silver et al. 2024; Zhou et al. 2024). However, in other studies, like (Stechly, Valmeekam, and Kambhampati 2024; Silver et al. 2022; Guan et al. 2023), the integration of validation tools led to only modest improvements.

Finally, considering the studies of the third category, some authors fine-tune and train smaller Transformer-based models on planning datasets to instruct these models how to generate a plan. This marks an important difference between the other two categories. In fact these works, instead of relying solely on the knowledge previously gained by a pre-trained LLM, try to inject some domain-specific planning knowledge into them. Among these works, the creators of Plansformer (Pallagani et al. 2023b, 2022, 2023a) fine-tuned a Transformer-base model (called CodeT5 (Raffel et al. 2020)) trained on code of several programming language. The results show that the fine-tuning significantly improves the planning capabilities of the model, up to a coverage higher than 80% across 6 planning domains.

Instead of fine-tuning an LLM, in (Rossetti et al. 2024b,a), a new GPT model, called PlanGPT, is trained from scratch to learn a general policy to solve many planning instances in a given domain. Exploiting a large amount of training data (about 63000 planning problems per domain), PlanGPT obtains impressive results in terms of coverage (more than 90% on IPC problems in Blocksworld, Driverlog, Floortile, Visitall and Zenotravel).

Integration with Validators and Planners

This section provides an overview of how the considered studies integrate LLMs with symbolic tools, in particular with reasoning-based validators, such as VAL (Howey, Long, and Fox 2004) and planners. These three forms (No Integration, Validator and Planner) are shown in Figure 3.

Analyzing the approaches which exploit validators, the work in (Guan et al. 2023) integrates the VAL validator (Howey, Long, and Fox 2004) during the LLM-based plan generation, reporting poor results. Conversely, the study in (Zhou et al. 2024) achieves improved coverage over standalone LLM performance on various benchmarks using the same approach. Additionally, (Silver et al. 2024) combines VAL and Python validators within the LLM workflow to enhance program synthesis and address domain-specific planning tasks.

Regarding the approaches integrated into planners, we can select two papers where a planner is exploited for modifying a candidate plan (which may be valid or invalid) generated by a LLM (Valmeekam et al. 2023b; Silver et al. 2022). The core idea behind these works is that even an imperfect, LLM-generated plan can provide helpful information that a planner can leverage to enhance performance. More specifically, in (Valmeekam et al. 2023b), a pre-trained GPT model on a general text corpus is combined with the LPG planner (Gerevini and Serina 2002). Differently, the work in (Silver et al. 2022) employs the plan resulting from the LLM to initialize the queue of expanded nodes in Greedy Best First Search (GBFS), using the LLM-generated candidates as a starting point for the search process. These two studies demonstrate that combining neural-based LLMs with symbolic planning approaches improves overall performance over LLMs alone, and reduces the search space compared to planners alone. However, these approaches are slower than using satisficing planners such as LPG (Gerevini and Serina 2002) and LAMA (Richter and Westphal 2010).

Considering the approaches that focus on computing heuristics with LLMs rather than generating plans directly, we can see also a more in-depth integration among LLMs and search-based techniques. The most notable papers are (Hirsch, Uziel, and Anaby-Tavor 2024) and (Hao et al. 2023), which combine LLM-derived heuristics with Greedy Best First Search (GBFS) and Monte Carlo Tree Search (MCTS), respectively. In these works, starting from the current state s of a planning problem, the LLM assigns a heuristic value to the next states derived from $A(s)$. This process guides the search by prioritizing the expansion of states with the highest heuristic values until all goals are satisfied. Although both these works share a similar idea, it is important to note that the scoring function in (Hirsch, Uziel, and Anaby-Tavor 2024) derives directly from the actions which the LLM associates to a higher probability. A more complex approach is presented in (Hao et al. 2023) which combines the probabilities of actions and states calculated by the LLM with a task-specific heuristic to obtain the final heuristic value. However, computing such heuristics is computationally very expensive, as an LLM must be prompted for each new applicable state. Moreover, there is no in-depth study of how these approaches compare to heuristics based on reasoning and planners.

Planning Capabilities Evaluation

Among the works we considered in this study, there is a great variety in how the LLMs are evaluated in terms of their planning capabilities. While establishing a standardized evalua-

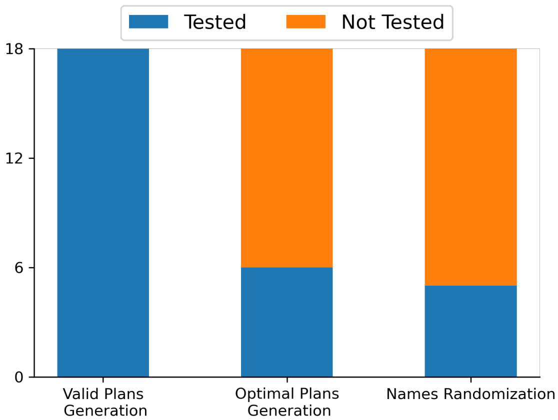

Figure 4: Types of evaluation conducted by the considered studies, in terms of Valid Plans Generations, Optimal Plans Generation, and whether an experiment with randomized names is present (Names Randomization). For each column, the blue bar represents the number of papers which perform the actual test, while the orange bar represents the number of papers which does not perform the test.

tion method across these systems, which vary significantly in architectures and setups (as explained in the previous sections), is not the main scope of this paper, we can identify several common evaluation points shared across the studies.

Figure 4 reports an overview of the main features commonly evaluated among the considered works. The most important aspect is undoubtedly whether they are able to generate valid plans. This is evaluated in terms of coverage, i.e. the percentage of planning problems correctly solved by the LLM-based technique proposed by the paper. In fact, as expected, all the papers assess whether the proposed architectures can generate a valid plan (i.e. that does not violate any action precondition and reaches all the problem goals). However, several studies also verify whether the solution provided is optimal (Valmeekam et al. 2022, 2023a; Pallagani et al. 2023a; Hazra, Martires, and Raedt 2024) or perform a more general evaluation of the plan quality, by comparing the quality of the generated plan with the optimal one (Valmeekam et al. 2022, 2023a). Interestingly, to improve the quality of the generated plans, the works by (Pallagani et al. 2023a) and (Hazra, Martires, and Raedt 2024) proposed neural language models specifically fine-tuned with optimal plans.

Another interesting experiment performed in the considered works is evaluating the impact of object and action names in PDDL domains and problems on the models' performance (Valmeekam et al. 2023b; Silver et al. 2022; Pallagani et al. 2023a). In most benchmark domains, the names used for objects and actions are derived from English words, which the LLMs likely encountered during their training. This evaluation aims to strip away the models' reliance on their linguistic knowledge to assess purely their reasoning ability to generate a plan. In (Silver et al. 2022), the objects and actions names are replaced by English words with different meanings. The work in (Valmeekam et al. 2023b,a) tests the effects of randomization by using both random strings and different words. The results indicate that all tested LLMs have a substantial drop in performance, producing almost no valid plan.

Considering the works that fine-tuned or even completely trained LLMs, those in (Rossetti et al. 2024b; Hazra, Martires, and Raedt 2024) assessed only the plan generation capability on problems matching the complexity level of the training set. Moreover, they do not test their system on more challenging problems with respect to those used in the training set, and they do not test the effects of randomization and word substitution. In contrast, Plansformer (Pallagani et al. 2023a) was also tested on randomized problems and on problems with more objects than seen during training. Both tests lead to a performance decline, with Plansformer solving only a small portion of the tested problems.

Discussion and Conclusions

We have analyzed the main characteristics of the works concerning the relationship among LLMs and automated planning, identifying various approaches for testing and developing new models. Our brief survey and categorization of these works highlight several differences in terms of input and output (typically, text versus PDDL), exploitation of pre-trained models (prompting, fine-tuning, or even training from scratch), integration of a planner or a plan validator. In the following, we offer some conclusions and a discussion to place the works we analyzed within a broader context.

Are LLMs Capable of Planning?

In order to understand whether LLMs can be used effectively as planners, it is important to understand the most crucial planners' properties and how they align with LLMs. Three fundamental properties in planning are completeness, soundness and domain independence. A planner is complete if it always finds a solution when a solution exists, and it is sound if the generated plans are guaranteed to be valid solutions.

The models presented in the reviewed works meet at most one of these properties. For instance, works that exploit zero-shot and few-shot prompting on pre-trained LLMs, such as GPT-3, are neither sound nor complete. Although, in theory, they possess the expressiveness to generate any possible solution across any different domain, they have a limited coverage, and, even when a solution is provided, there is no formal guarantee of it being valid, as shown in (Valmeekam et al. 2022, 2023b; Silver et al. 2022).

Considering works that train or fine-tune an LLM (such as (Pallagani et al. 2022; Hazra, Martires, and Raedt 2024)), they can be considered sound for a limited set of instances if their solution is authenticated by a validator. However, these approaches are not complete. Since they are trained on a finite vocabulary that includes a specific number of objects, the models are constrained by this vocabulary and extending the model's capability to handle a greater number of objects would require a complete retraining of the system. Clearly, this limits the flexibility and generalization capabilities of

LLMs beyond the specific instances they were trained on. Moreover, all these approaches have been realised for specific domains, and therefore they are not domain independent. Finally, the preliminary results from these fine-tuned or trained LLMs suggest they struggle with solving more complex problems, especially those that exceed the complexity of the tasks used during training (Pallagani et al. 2023a). Consequently, while these models can be helpful within their trained scope, they cannot fully generalize to a broader range of planning tasks without significant adjustments.

The highlighted limitations of LLMs in planning can be considered as part of a bigger problem that is currently being addressed by the research community, called the "stochastic parrot" problem (Bender et al. 2021; Hicks, Humphries, and Slater 2024). This problem posits that LLMs only mimic human language without possessing relevant reasoning capabilities and a real "understanding" of semantics. In planning, the stochastic parrot problem would justify the identified limitations and the scarce generalization performance that the LLMs (in particular, the trained and fine-tuned ones) achieve. Consider for instance the generated plans reported in (Valmeekam et al. 2022) for the Blocksworld domain. We can see that LLMs can generate plausible-sound solutions, in which the generated plan consists of stacking and unstacking blocks. However, they result in invalid plans. Similar results can be seen in (Rossetti et al. 2024b; Silver et al. 2024). In other words, these results could suggest that LLMs may parrot back learnt patterns, without a proper understanding of complex reasoning tasks, such as planning. However, these aspects are widely debated in the scientific community, and today there is no definitive conclusion regarding the true nature of LLMs capabilities or the best way to address these challenges.

How can we have a fair evaluation?

In the evaluation process, typically the ability of LLMs to generate valid plans is assessed. Additionally, some works test generation robustness in terms of names and actions randomization, and problem complexity. However, a contention point among the considered articles concerns the datasets and benchmarks employed.

A first benchmark, PlanBench, was proposed in (Valmeekam et al. 2023a), where the authors introduce 600 hand-crafted problems in natural language for the Blocksworld and Logistics domains in order to test the planning capabilities of LLMs, even with respect to several forms of randomizations of names and fluents. However, this benchmark has not been adopted in other works yet. Furthermore, it covers only two domains, which limits its ability to represent the full complexity of automated planning tasks. Other planning benchmarks can be obtained from the 2023 International Planning Competitions (IPC) (Taitler et al. 2024), which includes a broader range of domains but offers a limited number of problems per domain. While these problems are sufficient to test symbolic planners, they are often too few to evaluate deep learning systems effectively, which typically require large and diverse test sets (such as the 20% of the overall dataset, which is typically built of thousands of instances) with problems of different types and sizes. This is particularly relevant given the importance of ensuring that a deep learning model (such as a LLM) has genuinely learned to solve planning tasks, rather than just memorizing solutions from the training data. One of the first attempts to solve this problem appears in (Rossetti et al. 2024b), where the authors published a dataset containing 5000 problems for eight planning domains with a complexity similar to the IPC setup.

Another factor that needs to be evaluated is LLMs' capability to generate correct plans following order permutations of the fluents in the initial state and the goals or name randomization (of objects, fluents, and actions). This capability is fundamental because the reasoning process should be independent of the order of its fluents or name permutation, meaning that the same logical plan should be generated regardless of these variations. For example, considering the Blocksworld domain, the LLM should produce the same plan whether an object named $BlockA$ is renamed to $BlockY$ or if the fluents describing the initial state or goal are presented in a different order.

However, it is important to note that the knowledge of LLMs trained on text has been acquired from processing documents that reflect and describe the real world. Therefore, their ability to solve, for instance, a logistic problem stems not only from the logical structure of the problem, but also from understanding its main components, such as what is a *plane* object. Randomizing or replacing these meaningful object names with gibberish, while keeping the formal structure of the problem, could significantly impair the LLM's capabilities. The model may rely on some associations between real-world entities, learned during its pretraining, and their roles in planning contexts. By removing these associations, the LLM could struggle to generate a correct plan, as it would no longer recognize the objects or actions in a way that reflects their real-world functions. This behaviour contrasts sharply with the reasoning process of traditional planners. Unlike LLMs, planners are designed to be unaffected by word randomization because they treat words as mere placeholders without any intrinsic meaning. Planners rely solely on the logical relationships and formal structures defined by the problem specification, not on any semantic understanding of the objects or actions *names*.

Finally, an important aspect to assess is how LLMs handle complex planning problems, particularly those involving an increased numbers of objects, potential dead ends, and a high branching factor of fluents and actions within the domain. For training and fine-tuning approaches, where models are specifically trained on problems of a certain complexity, it is crucial to determine whether these models can generalize to more complex problems (in terms of number of objects and a higher branching factor) than those used during training. Suppose an LLM can solve problems beyond the complexity of its training dataset, in that case, we have clear evidence that the model has learned generalizable reasoning rather than merely memorizing solutions from the training data. However, if the model's performance deteriorates when faced with more complex problems, this indicates that the training process has limited generalization properties and that the LLM may only be effective within a narrow range of specific problems (Pallagani et al. 2023a).

Limitations

It is important to emphasize that all fine-tuning and training models are domain-specific. Even the current best-performing architecture, PlanGPT (Rossetti et al. 2024b), was implemented and tested by training several domain-specific models, each with its vocabulary and dataset. These models can learn a general policy and, therefore, solve many planning problems in a specific domain, but need different training for different domains. Therefore, their results should be analyzed and understood carefully to detect whether they demonstrate some forms of reasoning or, instead, learn some specific statistical patterns from the training data. In contrast, general-purpose LLMs trained on text do not precisely know any particular domain and should be more robust. However, to date, they are not able to plan even if they are trained with massive datasets and refined with human feedback (Valmeekam et al. 2023a).

Moreover, the training and evaluation of these architectures require huge computational resources and, as shown in (Silver et al. 2022), they can be slower than classical planners such as LAMA. These issues, along with the reliance on world knowledge (often expressed in PDDL) for approaches that use chain-of-thought reasoning, external validators, or human experts, leave open the question of whether in practice it is convenient to use LLMs to planning tasks, given that current domain independent (classical) planners are already highly efficient and optimized. An interesting direction for future work is combining these two technologies, taking the best from both worlds (Hao et al. 2023; Hirsch, Uziel, and Anaby-Tavor 2024).

References

Bender, E. M.; Gebru, T.; McMillan-Major, A.; and Shmitchell, S. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *FAccT*, 610–623. ACM.

Chen, L.; Lu, K.; Rajeswaran, A.; Lee, K.; Grover, A.; Laskin, M.; Abbeel, P.; Srinivas, A.; and Mordatch, I. 2021. Decision Transformer: Reinforcement Learning via Sequence Modeling. In *NeurIPS*, 15084–15097.

Chiari, M.; Gerevini, A. E.; Percassi, F.; Putelli, L.; Serina, I.; and Olivato, M. 2023. Goal Recognition as a Deep Learning Task: The GRNet Approach. In *ICAPS*, 560–568. AAAI Press.

Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT (1)*, 4171–4186. Association for Computational Linguistics.

Gerevini, A.; and Serina, I. 2002. LPG: A Planner Based on Local Search for Planning Graphs with Action Costs. In *AIPS*, 13–22. AAAI.

Geva, M.; Khashabi, D.; Segal, E.; Khot, T.; Roth, D.; and Berant, J. 2021. Did Aristotle Use a Laptop? A Question Answering Benchmark with Implicit Reasoning Strategies. *Trans. Assoc. Comput. Linguistics*, 9: 346–361.

Ghallab, M.; Knoblock, C.; Wilkins, D.; Barrett, A.; Christianson, D.; Friedman, M.; Kwok, C.; Golden, K.; Penberthy, S.; Smith, D.; Sun, Y.; and Weld, D. 1998. PDDL - The Planning Domain Definition Language.

Guan, L.; Valmeekam, K.; Sreedharan, S.; and Kambhampati, S. 2023. Leveraging Pre-trained Large Language Models to Construct and Utilize World Models for Model-based Task Planning. In *NeurIPS*.

Hao, S.; Gu, Y.; Ma, H.; Hong, J. J.; Wang, Z.; Wang, D. Z.; and Hu, Z. 2023. Reasoning with Language Model is Planning with World Model. In *EMNLP*, 8154–8173. Association for Computational Linguistics.

Hazra, R.; Martires, P. Z. D.; and Raedt, L. D. 2024. SayCanPay: Heuristic Planning with Large Language Models Using Learnable Domain Knowledge. In *AAAI*, 20123–20133. AAAI Press.

Hicks, M. T.; Humphries, J.; and Slater, J. 2024. ChatGPT is bullshit. *Ethics Inf. Technol.*, 26(2): 38.

Hirsch, E.; Uziel, G.; and Anaby-Tavor, A. 2024. What's the Plan? Evaluating and Developing Planning-Aware Techniques for LLMs. *CoRR*, abs/2402.11489.

Howey, R.; Long, D.; and Fox, M. 2004. VAL: Automatic Plan Validation, Continuous Effects and Mixed Initiative Planning Using PDDL. In *ICTAI*, 294–301. IEEE Computer Society.

Jiang, Z.; Xu, F. F.; Araki, J.; and Neubig, G. 2020. How Can We Know What Language Models Know. *Trans. Assoc. Comput. Linguistics*, 8: 423–438.

Kambhampati, S.; Valmeekam, K.; Guan, L.; Verma, M.; Stechly, K.; Bhambri, S.; Saldyt, L.; and Murthy, A. 2024. Position: LLMs Can't Plan, But Can Help Planning in LLM-Modulo Frameworks. In *ICML*. OpenReview.net.

Karia, R.; and Srivastava, S. 2021. Learning Generalized Relational Heuristic Networks for Model-Agnostic Planning. In *AAAI*, 8064–8073. AAAI Press.

Lehnert, L.; Sukhbaatar, S.; McVay, P.; Rabbat, M.; and Tian, Y. 2024. Beyond A*: Better Planning with Transformers via Search Dynamics Bootstrapping. *CoRR*, abs/2402.14083.

Liu, B.; Jiang, Y.; Zhang, X.; Liu, Q.; Zhang, S.; Biswas, J.; and Stone, P. 2023a. LLM+P: Empowering Large Language Models with Optimal Planning Proficiency. *CoRR*, abs/2304.11477.

Liu, P.; Yuan, W.; Fu, J.; Jiang, Z.; Hayashi, H.; and Neubig, G. 2023b. Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. *ACM Comput. Surv.*, 55(9).

Madasu, A.; and Srivastava, S. 2022. What do Large Language Models Learn beyond Language? In *EMNLP (Findings)*, 6940–6953. Association for Computational Linguistics.

Oswald, J. T.; Srinivas, K.; Kokel, H.; Lee, J.; Katz, M.; and Sohrabi, S. 2024. Large Language Models as Planning Domain Generators. In *ICAPS*, 423–431. AAAI Press.

Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P. F.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*.

Pallagani, V.; Muppasani, B.; Murugesan, K.; Rossi, F.; Horesh, L.; Srivastava, B.; Fabiano, F.; and Loreggia, A. 2022. Plansformer: Generating Symbolic Plans using Transformers. *CoRR*, abs/2212.08681.

Pallagani, V.; Muppasani, B.; Murugesan, K.; Rossi, F.; Srivastava, B.; Horesh, L.; Fabiano, F.; and Loreggia, A. 2023a. Understanding the Capabilities of Large Language Models for Automated Planning. *CoRR*, abs/2305.16151.

Pallagani, V.; Muppasani, B.; Srivastava, B.; Rossi, F.; Horesh, L.; Murugesan, K.; Loreggia, A.; Fabiano, F.; Joseph, R.; and Kethepalli, Y. 2023b. Plansformer Tool: Demonstrating Generation of Symbolic Plans Using Transformers. In *IJCAI*, 7158–7162. ijcai.org.

Pallagani, V.; Muppasani, B. C.; Roy, K.; Fabiano, F.; Loreggia, A.; Murugesan, K.; Srivastava, B.; Rossi, F.; Horesh, L.; and Sheth, A. P. 2024. On the Prospects of Incorporating Large Language Models (LLMs) in Automated Planning and Scheduling (APS). In *ICAPS*, 432–444. AAAI Press.

Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training.

Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *J. Mach. Learn. Res.*, 21: 140:1–140:67.

Richter, S.; and Westphal, M. 2010. The LAMA Planner: Guiding Cost-Based Anytime Planning with Landmarks. *J. Artif. Intell. Res.*, 39: 127–177.

Rossetti, N.; Tummolo, M.; Gerevini, A. E.; Olivato, M.; Putelli, L.; and Serina, I. 2024a. Enhancing GPT-Based Planning Policies by Model-Based Plan Validation. In *NeSy (2)*, volume 14980 of *Lecture Notes in Computer Science*, 328–337. Springer.

Rossetti, N.; Tummolo, M.; Gerevini, A. E.; Putelli, L.; Serina, I.; Chiari, M.; and Olivato, M. 2024b. Learning General Policies for Planning through GPT Models. In *ICAPS*, 500–508. AAAI Press.

Shen, W.; Trevizan, F. W.; and Thiébaux, S. 2020. Learning Domain-Independent Planning Heuristics with Hypergraph Networks. In *ICAPS*, 574–584. AAAI Press.

Silver, T.; Dan, S.; Srinivas, K.; Tenenbaum, J. B.; Kaelbling, L. P.; and Katz, M. 2024. Generalized Planning in PDDL Domains with Pretrained Large Language Models. In *AAAI*, 20256–20264. AAAI Press.

Silver, T.; Hariprasad, V.; Shuttleworth, R. S.; Kumar, N.; Lozano-Pérez, T.; and Kaelbling, L. P. 2022. PDDL planning with pretrained large language models. In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.

Singh, I.; Blukis, V.; Mousavian, A.; Goyal, A.; Xu, D.; Tremblay, J.; Fox, D.; Thomason, J.; and Garg, A. 2023. ProgPrompt: Generating Situated Robot Task Plans using Large Language Models. In *ICRA*, 11523–11530. IEEE.

Stechly, K.; Valmeekam, K.; and Kambhampati, S. 2024. On the Self-Verification Limitations of Large Language Models on Reasoning and Planning Tasks. *CoRR*, abs/2402.08115.

Taitler, A.; Alford, R.; Espasa, J.; Behnke, G.; Fiser, D.; Gimelfarb, M.; Pommerening, F.; Sanner, S.; Scala, E.; Schreiber, D.; Segovia-Aguas, J.; and Seipp, J. 2024. The 2023 International Planning Competition. *AI Mag.*, 45(2): 280–296.

Toyer, S.; Thiébaux, S.; Trevizan, F. W.; and Xie, L. 2020. ASNets: Deep Learning for Generalised Planning. *J. Artif. Intell. Res.*, 68: 1–68.

Valmeekam, K.; Hernandez, A. O.; Sreedharan, S.; and Kambhampati, S. 2022. Large Language Models Still Can't Plan (A Benchmark for LLMs on Planning and Reasoning about Change). *CoRR*, abs/2206.10498.

Valmeekam, K.; Marquez, M.; Hernandez, A. O.; Sreedharan, S.; and Kambhampati, S. 2023a. PlanBench: An Extensible Benchmark for Evaluating Large Language Models on Planning and Reasoning about Change. In *NeurIPS*.

Valmeekam, K.; Marquez, M.; Sreedharan, S.; and Kambhampati, S. 2023b. On the Planning Abilities of Large Language Models - A Critical Investigation. In *NeurIPS*.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *NIPS*, 5998–6008.

Wei, J.; Bosma, M.; Zhao, V. Y.; Guu, K.; Yu, A. W.; Lester, B.; Du, N.; Dai, A. M.; and Le, Q. V. 2022a. Finetuned Language Models are Zero-Shot Learners. In *ICLR*. OpenReview.net.

Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E. H.; Le, Q. V.; and Zhou, D. 2022b. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *NeurIPS*.

Zhang, H.; Li, L. H.; Meng, T.; Chang, K.; and den Broeck, G. V. 2023. On the Paradox of Learning to Reason from Data. In *IJCAI*, 3365–3373. ijcai.org.

Zhou, Z.; Song, J.; Yao, K.; Shu, Z.; and Ma, L. 2024. ISR-LLM: Iterative Self-Refined Large Language Model for Long-Horizon Sequential Task Planning. In *ICRA*, 2081–2088. IEEE.