

# SayCanPay: Heuristic Planning with Large Language Models Using Learnable Domain Knowledge

Rishi Hazra<sup>1</sup>, Pedro Zuidberg Dos Martires<sup>1</sup>, Luc De Raedt<sup>1,2</sup>

<sup>1</sup>Centre for Applied Autonomous Sensor Systems (AASS), Örebro University, Sweden

<sup>2</sup>KU Leuven, Belgium

{rishi.hazra, pedro.zuidberg-dos-martires, luc.de-raedt}@oru.se

## Abstract

Large Language Models (LLMs) have demonstrated impressive planning abilities due to their vast “world knowledge”. Yet, obtaining plans that are both feasible (grounded in affordances) and cost-effective (in plan length), remains a challenge, despite recent progress. This contrasts with heuristic planning methods that employ domain knowledge (formalized in action models such as PDDL) and heuristic search to generate feasible, optimal plans. Inspired by this, we propose to combine the power of LLMs and heuristic planning by leveraging the world knowledge of LLMs and the principles of heuristic search. Our approach, SayCanPay, employs LLMs to generate actions (*Say*) guided by learnable domain knowledge, that evaluates actions’ feasibility (*Can*) and long-term reward/payoff (*Pay*), and heuristic search to select the best sequence of actions. Our contributions are (1) a novel framing of the LLM planning problem in the context of heuristic planning, (2) integrating grounding and cost-effective elements into the generated plans, and (3) using heuristic search over actions. Our extensive evaluations show that our model surpasses other LLM planning approaches.

## Introduction

With the rise of Large Language Models (LLMs), there has been a growing interest in leveraging their generative capabilities for planning tasks (Huang et al. 2022a; Valmeekam et al. 2022; Silver et al. 2022; Liu et al. 2023). These models have the ability to generate long-horizon plans, capitalizing on their extensive “world knowledge” gained from training on vast amounts of data (e.g. eggs are typically stored in the refrigerator, and placing an apple in the fridge will cool it). Such expansive knowledge can be exploited to plan in an open-world context (Ding et al. 2023). Moreover, planning in the natural language space offers significant flexibility especially, with the advent of multimodal foundation models (Lakhotia et al. 2021; Du et al. 2022; Brohan et al. 2023). Such models have made it easier to represent various modalities such as vision, speech, and even actions in the form of natural language, thus bypassing the need to have domain-specific knowledge (e.g. PDDL) that traditional planning approaches require. However, LLM-based planning often faces challenges, particularly in generating feasible plans. It can

fail to model action affordances (or pre-conditions)<sup>1</sup> due to difficulty in modeling the state of the world (e.g. *grab milk from the fridge* even if the door is closed) or having a pre-trained world model that is not aligned with the current environment (e.g. *using a controller to regulate the heater* where only a knob exists), leading to infeasible plans. Moreover, such models focus greedily on the next actionable step without considering its relevance to the ultimate goal, resulting in longer, cost-inefficient plans (Valmeekam et al. 2023). Recent works like SayCan (Ahn et al. 2022) have sought to address the affordance problem by using pretrained skills to evaluate the action’s executability – *Can the action be executed in the current state?* However, the plan cost remains a concern.

In contrast, traditional planning provides an established approach to developing a sequence of actions to transition from an initial state to a goal state. It uses a domain file (with action models defined in PDDL specifying pre- and post-conditions) and heuristic search planners like Fast Downward (Helmert 2006) to ensure feasibility through grounding in preconditions, and generating cost-effective plans by employing search trees to select the best (or shortest) sequence of actions. However, obtaining a domain file for complex real-world environments is difficult, and its use restricts planning to a closed-world setting. These methods also struggle to handle partial observations, although approximate planning (Kaelbling, Littman, and Cassandra 1998) can alleviate it.

Integrating LLMs with classical planning offers a promising research path, merging the generative abilities and (open) world knowledge of LLMs with the methodological rigor of planning algorithms. To this end, we extend the following contributions. **(1)** We propose to frame language model planning in the context of heuristic planning, which to our knowledge, is the first of its kind. **(2)** We incorporate feasibility and cost-effective elements into the generated plans using a joint scoring named **SayCanPay**<sup>2</sup>. As shown in Figure 1, it guides the planning through three key steps: (i) Say: Given a goal and an initial observation, the LLM

<sup>1</sup>In robotics, affordances refer to possible actions that *can* be executed, which is conceptually similar to inferring preconditions in planning – what actions are feasible in a certain situation.

<sup>2</sup>**Code link:** <https://rishihazra.github.io/SayCanPay/>

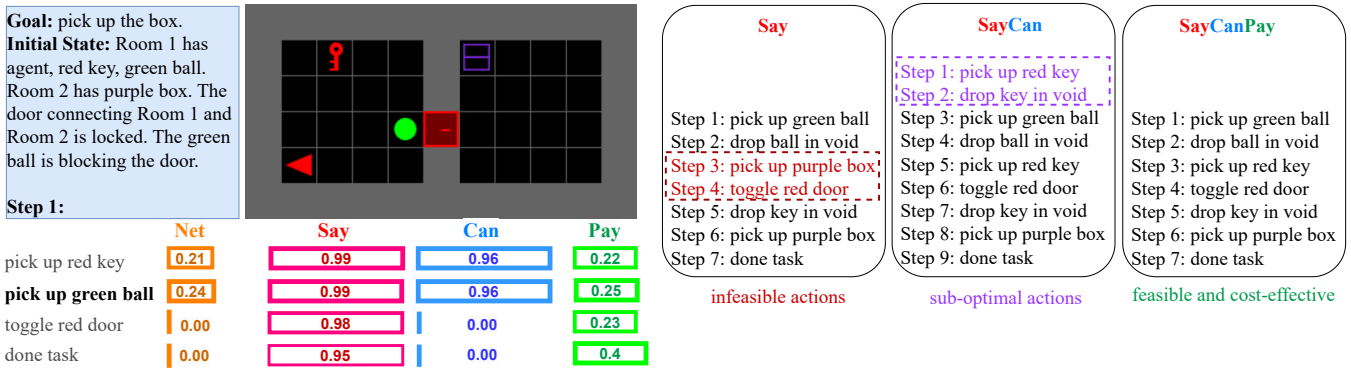


Figure 1: Figure illustrates how SayCanPay scores the next action in BabyAI environment (Chevalier-Boisvert et al. 2019). Given inputs: goal  $g$  and initial observation  $o_0$ , the Say model generates candidate actions with associated probabilities. These are then scored for feasibility by the Can model and for payoff by the Pay model. Here, the Can model deems both *pick up red key* and *pick up green ball* equally probable (i.e. both preconditions are satisfied). However, the Pay model ensures a better payoff for *pick up green ball*. We compare plans generated by Say, SayCan, and SayCanPay scoring. Say scoring can lead to infeasible plans and SayCan to feasible but longer plans. The displayed grid is purely illustrative, with no visual inputs used.

generates *likely* candidate actions at each step; (ii) Can: An affordance model scores these actions’ *feasibility*, mirroring the evaluation of preconditions; (iii) Pay: Another model scores the actions according to their *estimated payoff*, akin to heuristic estimators. The Can and Pay models undergo domain-specific training to align the plans with the current environment. (3) Using this combined score as a heuristic, we search for the most feasible and cost-effective plan. We demonstrate both quantitatively and qualitatively, how our proposed joint scoring and heuristic search improve over the current LLM planning frameworks.

## Related Work on Planning with LLMs

Table 1 categorizes LLM planning works into two broad categories based on whether the inputs (goals, states) and output actions (I/O) are natural language (NL) or symbolic (PDDL, scripting language). The approaches in the first category (Huang et al. 2022a; Valmeekam et al. 2022) often fail to model action affordances and the state of the world, leading to the generation of infeasible plans (Valmeekam et al. 2022). To improve the groundedness, recent works have explored planning guided by learnable domain-specific models that score the actions’ *feasibility* akin to preconditions (Huang et al. 2023; Lin et al. 2023). Notably, SayCan (Ahn et al. 2022) uses pretrained low-level skills to ground the LM-generated actions. Others have used online planning with environmental and human feedback (Huang et al. 2022b). A limitation of such models, however, is their short-sighted nature, as they focus greedily on the next feasible action without considering its long-term relevance to the goal. Moreover, the plans are generated in an online fashion, interleaving action generation and execution, thus simplifying state tracking. In contrast, SayCanPay performs offline planning (i.e. complete plan generation while maintaining an internal world state) with both precondition and heuristic estimators, improving plan feasibility and cost-efficiency.

Another line of work employs LLMs to create offline

symbolic plans, leveraging LLMs’ training on open-source codebases, where actions appear as function calls (Singh et al. 2023; Liang et al. 2023). The feasibility of plans is ensured through assertion checks (`assert (preconditions)`), that may trigger recovery actions. However, it relies solely on the LLM’s domain knowledge which is limited to its training data and may not be aligned with the agent’s current environment (e.g. espresso machine operations vary widely). Conversely, SayCanPay uses additional models trained with domain-specific knowledge collected from the current environment. There are also efforts to fine-tune LLMs like CodeT5 (Wang et al. 2021) to generate plans in PDDL (Pallagani et al. 2022). This requires a significant amount of training data (given LLMs’ minimal PDDL exposure) which is not entirely justified by their performance.

Yet another exciting line of work explores hybrid I/O systems like LLM+P (Liu et al. 2023) wherein, given a PDDL domain file (with a predefined action model), the LLM maps the NL inputs (task description, input observation) to a PDDL problem file. A symbolic planner then generates the plan. However, its effectiveness is limited by the closed-world constraint of the domain file, the necessity for fully observable states, and the LLM’s restricted capability in translating NL to PDDL (Xie et al. 2023).

## Preliminaries

**Planning Framework.** We formulate our planning problem, based on approximate planning (Golowich, Moitra, and Rohatgi 2022), as a finite-horizon Partially Observable Markov Decision Process (POMDP) given by the tuple  $\langle \mathcal{S}, \mathcal{S}_G, b_0, \mathcal{A}, \mathcal{O}, R, \mathbb{T} \rangle$ . Here,  $\mathcal{S}$  is state space,  $\mathcal{S}_G \subseteq \mathcal{S}$  is a set of goal states,  $b_0$  is the initial belief state,  $\mathcal{A}$  is the set of actions,  $\mathcal{O}$  is a set of observations retrieved from states via an observation function  $O$ ,  $R : \mathcal{O} \rightarrow \mathbb{R}$  is a known reward function,  $\mathbb{T} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta^{\mathcal{S}}$  is a known stochastic transition function and  $\Delta^{\mathcal{S}}$  is a distribution over states. Belief states represent the agent’s knowledge of the environment at any

Model	I/O	Planner	Domain Knowledge		Search	Planning
			Affordances	Heuristics		
HSP (Bonet and Geffner 2001)	Symbolic	Symbolic	✓	✓	Heuristic	Offline
LLM+P (Liu et al. 2023)	Hybrid	Symbolic	✓	✓	Heuristic	Offline
Planning LM (Huang et al. 2022a)	NL	LLM	✗	✗	Greedy*	Offline
SayCan (Ahn et al. 2022)	NL	LLM	✓	✗	Greedy*	Online
Grounded Decoding (Huang et al. 2023)	NL	LLM	✓	✗	Greedy*	Online
Text2Motion (Lin et al. 2023)	NL	LLM	✓	✗	Greedy*	Online
ProgPrompt (Singh et al. 2023)	Symbolic	LLM	✓	✗	Greedy*	Offline
Plansformer (Pallagani et al. 2022)	Symbolic	LLM	✓	✗	Greedy*	Offline
<b>SayCanPay</b> (Beam-Action)	NL	LLM	✓	✓	Heuristic	Offline

Table 1: Table contrasts SayCanPay with existing works. I/O: input (goal/task, observation/state) / output (actions), NL: natural language. Here, Greedy\* suggests the algorithm greedily selects actions while (possibly) searching over tokens.

point, given as  $b \in \Delta^{\mathcal{S}}$ . Additionally, let  $\mathcal{H}_t := (\mathcal{A} \times \mathcal{O})_{t-1}$  denote the set of *histories* at step  $t$ , namely the set of action/observation sequences  $(o_0, a_1, o_1, \dots, a_{t-1}, o_{t-1})$  or  $(a_{1:t-1}, o_{0:t-1})$  the agent has access to before selecting action  $a_t$ . It is assumed that the goal states are fully observable.

Unlike MDPs, the optimal policy in a POMDP typically takes actions depending on not just the most recent observation but the entire history. The objective of the planning algorithm is to find the optimal sequence of actions  $a_{1:T}$  (i.e. an optimal plan) from an initial belief state  $b_0$  to a given goal state  $g \in \mathcal{S}_G$ . Here,  $T$  is the length of the horizon.

**Heuristic Search Planning.** In real-world scenarios where the state space can be exponentially large to explore exhaustively, heuristic search planning (HSP) becomes useful (Bonet and Geffner 2001). Essentially, it uses heuristic functions  $f_{\text{heur}} : \mathcal{H}_t \times \mathcal{S}_G \rightarrow \mathbb{R}$  to guide the search process in the planning problem, by computing a cost estimate from a given history of actions and observations. An example is the Best-First Search algorithms that select the most promising (next) action(s) using a linear combination of previously accumulated **cost**  $f_{\text{acc}}$  for history  $h_{t-1}$ , and the estimated **cost**  $f_{\text{heur}}$  from updated history  $h_t = (h_{t-1}, a_t)$  and goal  $g$ .

$$f(h_t) = z_1 \cdot f_{\text{acc}}(h_{t-1}) + z_2 \cdot f_{\text{heur}}(h_t, g) \quad (1)$$

Here  $z_1, z_2 \in \{0, 1\}$ . The next action  $a_t = \arg \min_{h_t} f(h_t)$ . Special cases are the  $A^*$  algorithm algorithm ( $z_1 = 1$  and  $z_2 = 1$ ) and Greedy Best-First Search ( $z_1 = 0$  and  $z_2 = 1$ ).

## Language Model Planning Framework

We keep the same POMDP formulation while updating our interpretations of the tuple. Previous works have shown that language models (LMs) trained on extensive data would internalize rich world knowledge that can be queried for downstream tasks like planning (Hao et al. 2023). This is akin to an internal transition function  $\mathbb{T}^{int}$ . Similarly, LMs also maintain and update an internal belief state  $b_t^{int}$  over tokens (or actions). An observation function maps states to NL observations,  $\mathcal{O} : \mathcal{S} \rightarrow \mathcal{O}$ . The updated POMDP is now given as  $\langle \mathcal{S}, \mathcal{S}_G, b_0^{int}, \mathcal{A}, \mathcal{O}, R, \mathbb{T}^{int} \rangle$ . In our offline planning experiments, we assume the following: (i)  $\mathcal{O} = \{o_0\}$  inducing belief state  $b_0^{int} = \mathbb{1}_{s_0}$ , while  $o_t = \emptyset \forall t > 0$ , due to lack

of environmental feedback; (ii) sparse rewards = 1 for plan success, else 0. While our LM does not utilize the reward function, one could use it for alignment (Ziegler et al. 2020).

**Problem Statement:** Given a NL goal  $g$ , history  $h_0 = (o_0)$ , and a LM generating actions  $a_t$  with probability  $p(a_t|h_{t-1}, g)$ , generate the most likely plan  $(a_{1:T})$  to go from  $b_0^{int}$  to  $g$ , i.e.,  $\arg \max_{a_{1:T}} P(a_{1:T}|h_0, g)$ .

We aim to maximize the plan’s probability, reframing LM planning as a classical search problem, where we repeatedly expand the current plan  $a_{1:t-1}$  by adding action  $a_t$ . Rewriting the probability  $P(a_{1:T}|h_0, g)$  recursively as:

$$\begin{aligned} &= P(a_{1:t-1}, a_t, a_{t+1:T}|h_0, g) \\ &= p(a_{1:t-1}|h_0, g)p(a_t|h_0, a_{1:t-1}, g)p(a_{t+1:T}|h_0, a_{1:t}, g) \\ &= p(a_{1:t-1}|h_0, g) \cdot p(a_t|h_{t-1}, g) \cdot p(a_{t+1:T}|h_t, g) \end{aligned}$$

To align with Eq 1 of the planning problem, we take log on both sides and **maximize** rather than minimize. We get accumulated **value**  $f_{\text{acc}}(h_{t-1}) = \log p(a_{1:t-1}|h_0, g)$ , heuristic **payoff**  $f_{\text{heur}}(h_t, g) = p(a_{t+1:T}|h_t, g)$ , and  $f(h_t) = \log P(a_{1:T}|h_0, g)$ . Rewriting the above equation:

$$f(h_t) = f_{\text{acc}}(h_{t-1}) + \log (p(a_t|h_{t-1}, g) \cdot f_{\text{heur}}(h_t, g)) \quad (2)$$

The additional  $p(a_t|h_{t-1}, g)$  reflects that, unlike classical planning which evaluates only feasible actions based on pre-conditions, LMs assign probabilities to each action. Here, next action  $a_t = \arg \max_{h_t} f(h_t)$ .

Technically, the LM generates actions wherein each action is a sequence of tokens until the end-of-sequence token,  $\langle \text{EOS} \rangle$ . For each action step  $a = (w_1, \dots, w_n)$  composed of tokens  $w_i$ , the LM computes the action probability as  $p(a) = p(w_1) \prod_{i=2}^n p(w_i|w_{1:i-1})$ . Planning LM (Huang et al. 2022a) proposed a greedy decoding strategy wherein the LM greedily picks the next token, henceforth referred to as **Greedy-Token** baseline (Figure 2 Left). The generated action is then appended to the history  $h_t = (h_{t-1}, a_t)$ , and the generation process repeats until a “done task” action is generated. Subsequent works (Lin et al. 2023) have investigated beam search over tokens. However, we are mainly interested in searching on the level of actions and not tokens.

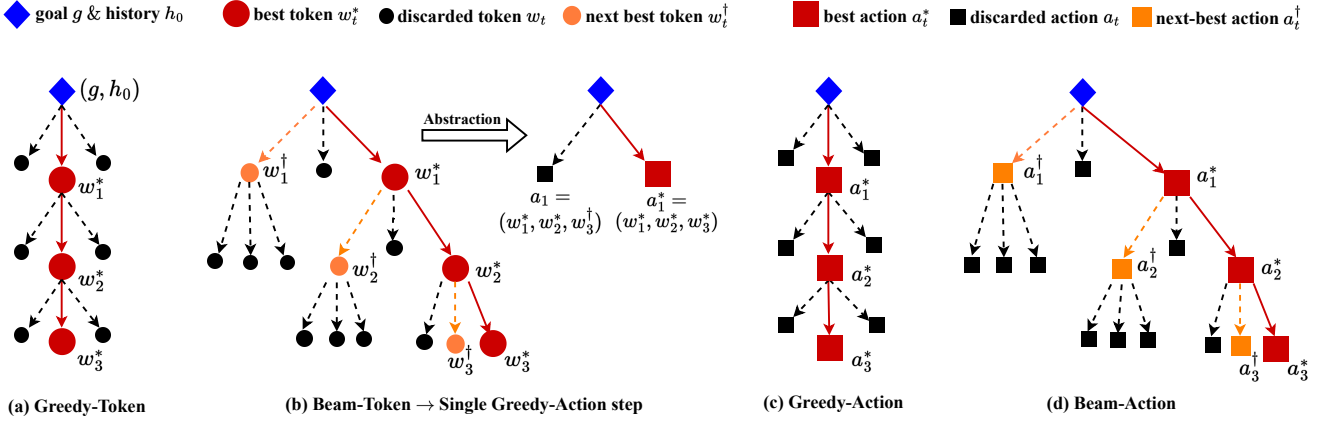


Figure 2: The figure outlines decoding strategies – Greedy-Token, Greedy-Action, and Beam-Action. Greedy-Token greedily selects the next best token by its probability. Greedy-Action (which is a beam search over tokens) greedily selects the next best action based on a specific decoding score. Beam-Action uses a beam search over actions, maintaining  $k$  beams and selecting the best sequence as the plan. Here, nodes represent either tokens  $w_t$  or actions  $a_t$ . The best plan is given by  $(a_1^*, a_2^*, a_3^*)$  and represented in red. The second-best node is in orange, discarded ones in black. Here, for Beam-Action,  $m = 3$  and  $k = 2$ .

## SayCanPay Inference

The core concept of SayCanPay is to guide LMs in generating feasible and cost-effective plans. The process unfolds in three key steps: (1) **Say**: At each step  $t$ , the LM generates the *top-m* candidate actions with associated probabilities  $\{p(a_t^i|h_{t-1}, g)\}_{i=1}^m$ . This generation employs a beam search over tokens. (2) **Can**: Next, a trained domain-specific model weighs these candidate actions on their feasibility, mirroring precondition evaluation. (3) **Pay**: Finally, a trained domain-specific estimator weighs the candidate actions according to their estimated payoff. The probabilities from these three components are then combined to select the next action. An overview of SayCanPay is provided in Figure 1.

In what follows, we instantiate the LM planning problem with two decoding strategies (or search algorithms that select the next action(s)): **Greedy Action** and **Beam Action**. Each strategy is explored using three distinct decoding scores (i.e. score used by the search algorithm to select the next action) – Say, SayCan, SayCanPay. We then elaborate on the training of Can and Pay models.

### Greedy-Action

In this decoding strategy, we maintain a single action sequence and at each step, greedily choose the next *best* action based on a specific decoding score. This is akin to performing Greedy Best-First Search with  $z_1 = 0$  and  $z_2 = 1$ . The decoding score for each candidate action  $a_t^i$  is given as:

$$f(h_t^i) = \log(p(a_t^i|h_{t-1}, g) \cdot f_{\text{heur}}(h_t^i, g))$$

Here, the best action  $a_t^* = \arg \max_{h_t^i} f(h_t^i)$ , where  $h_t^i = (h_{t-1}, a_t^i)$  denotes the current history with  $i^{\text{th}}$  candidate action. As shown in Figure 2, this approach can be viewed as being “greedy” with respect to actions while using “beams” over the tokens. Now, we explore three variations of the strategy based on how the decoding score is computed.

- **Say**: In this decoding score, we set the estimated payoff  $f_{\text{heur}}(h_t^i, g) = 1 \forall i \in \{1, \dots, m\}$ . Hence, the action is selected solely based on the LM generation probability, without considering feasibility or payoff.

$$f(h_t^i) = \log \left( \underbrace{p(a_t^i|h_{t-1}, g)}_{=:p_{a_t^i}^{\text{say}}} \right) \quad (3)$$

- **SayCan**: Here, the action feasibility is also considered. Let,  $\sigma_t = (a_t, pre(a_t))$  where  $pre(a_t)$  denotes the preconditions of  $a_t$ . The “can” probability<sup>3</sup>, is denoted by  $p(pre(a_t)|h_{t-1}, g)$ . Again,  $f_{\text{heur}}(h_t^i, g) = 1 \forall i$ .

$$\begin{aligned} f(h_t^i) &= \log(p(\sigma_t^i|h_{t-1}, g)) \\ &= \log \left( \underbrace{p(a_t^i|h_{t-1}, g)}_{=:p_{a_t^i}^{\text{say}}} \cdot \underbrace{p(pre(a_t^i)|h_{t-1}, g)}_{=:p_{a_t^i}^{\text{can}}} \right) \end{aligned} \quad (4)$$

- **SayCanPay**: This decoding score accounts for the estimated payoff in addition to the abovementioned scores. Hence, the best action is selected based on a combined score of Say, Can, and Pay scores.

$$\log \left( \underbrace{p(a_t^i|h_{t-1}, g)}_{=:p_{a_t^i}^{\text{say}}} \cdot \underbrace{p(pre(a_t^i)|h_{t-1}, g)}_{=:p_{a_t^i}^{\text{can}}} \cdot \underbrace{f_{\text{heur}}(h_t^i, g)}_{=:p_{a_t^i}^{\text{pay}}} \right) \quad (5)$$

### Beam-Action

In heuristic planning, multiple potential plans (i.e. action sequences) are simultaneously maintained and iteratively expanded until the goal is achieved. To simulate this behavior, we propose to manage  $k$  action sequences. It works as

<sup>3</sup>The goal  $g$  is used to evaluate the preconditions of “done task”.

follows – each sequence is expanded with  $m$  candidate actions (where  $m \geq k$ ) from the LM, resulting in a total of  $k \times m$  sequences. Then, top- $k$  sequences are retained using a specific decoding score accumulated over the sequence, as shown below. Once all  $k$ -beams have terminated, we select the sequence with the highest (length-normalized)<sup>4</sup> accumulated score. To avoid repetition, we only show the SayCan-Pay version. The rest can be similarly formulated.

$$\text{top-}k \left[ \frac{1}{|h_t^{ij}|} \left( f_{\text{acc}}(h_{t-1}^i) + \log p(\sigma_t^j | h_{t-1}^i, g) \cdot f_{\text{heur}}(h_t^{ij}, g) \right) \right]$$

Here,  $i \in \{1, \dots, k\}$ ,  $j \in \{1, \dots, m\}$ ,  $k \leq m$ . The updated history  $h_t^{ij} = (h_{t-1}^i, a_t^j)$  is obtained by adding the action  $a_t^j$  to the  $i^{\text{th}}$  beam history  $h_{t-1}^i$ . The outcome becomes the value for  $f_{\text{acc}}(h_t)$  for the next iteration. Note, that setting  $k = 1$  results in Greedy-Action decoding.

Our proposed decoding has similarities with Tree-of-Thoughts inference (Yao et al. 2023) which also maintains multiple reasoning paths to decide the next step. However, our method is specifically tailored for planning problems. It uses search and evaluation techniques akin to planning methods, making it more suited for such challenges. Now, we discuss training details of the Can and Pay models.

## Learning the Can and Pay Models

To train our domain-specific Can and Pay models, we collect  $N$ -expert trajectories  $\mathcal{E} = \{\tau\}_{n=1}^N$  for each environment using an oracle planner, where  $\tau_i = (o_0, g, a_1, a_2, \dots, a_T, r)$ . Note,  $r = 1$  for all expert trajectories.

### Can Model

We model it as a classification problem, where the positive action (i.e., the action whose preconditions are satisfied) is assigned the highest probability from a set of one positive and a few negative actions. Specifically, we sample a batch of actions  $[h_{t-1}, g, a_t, a_{\bar{t} \neq t}, \tilde{a}]^{1:B}$  from expert trajectories  $\mathcal{E}$ . We then train a model  $\mathcal{M}^{\text{can}}$  with the aim of minimizing the InfoNCE loss (van den Oord, Li, and Vinyals 2019):

$$-\frac{1}{B} \sum_{i=1}^B \log \frac{\mathcal{M}^{\text{can}}(h_{t-1}^i, g^i, a_t^i)}{\sum_{a \in \{a_t^i, a_{\bar{t} \neq t}^i, \tilde{a}^i\}} \mathcal{M}^{\text{can}}(h_{t-1}^i, g^i, a)}$$

Here,  $B$  is the batch size,  $a_t$  is the positive action from trajectory  $\tau_i$  executed in the context of history  $h_{t-1}$  with goal  $g$ ,  $a_{\bar{t} \neq t}$  is a negative action sampled from the same trajectory  $\tau_i$ , but at a different time-step  $\bar{t}$ , and  $\tilde{a}$  is a negative action sampled from a different trajectory  $\tau_{j \neq i}$  with a different initial observation  $o_0$  and goal  $g$ .  $\mathcal{M}^{\text{can}}$  consists of an uncased Bert model (Devlin et al. 2019) with a probe layer and is trained end-to-end to correctly identify the positive action. The input to  $\mathcal{M}^{\text{can}}$  is of the format ‘⟨Goal⟩{ $g$ } ⟨History⟩{ $h_{t-1}$ } ⟨NXT⟩{ $a_t$ }’. Here, ‘⟨\*⟩’ serves as special tokens. The output is the Can probability  $p_{a_t}^{\text{can}} := \mathcal{M}^{\text{can}}(h_{t-1}, g, a_t)$ . The model is trained across multiple batches for F1-score convergence on the validation set.

<sup>4</sup>Since different beams can have different sequence lengths.

Our approach is different from SayCan (Ahn et al. 2022) which trains multiple affordance functions (corresponding to different skills), through temporal-difference-based reinforcement learning to predict the likelihood of a particular skill succeeding (i.e., executing) in the current state. Here, we show two training I/O examples, one with positive action and another one with negative action.

**Input** ⟨Goal⟩ pick up the purple box. ⟨Initial State⟩ Room 1 has yellow key, agent. Room 2 has purple box. The door connecting Room 1 and Room 2 is locked. ⟨Step 1⟩ pick up yellow key. ⟨NXT⟩ toggle yellow door.

**Output** 1.0 // feasible

**Input** ⟨Goal⟩ pick up the purple box. ⟨Initial State⟩ Room 1 has yellow key, agent. Room 2 has purple box. The door connecting Room 1 and Room 2 is locked. ⟨Step 1⟩ pick up yellow key. ⟨NXT⟩ pick up purple box.

**Output** 0.0 // infeasible

### Pay Model

We model it as a regression problem to estimate action payoffs. Using expert trajectories  $\mathcal{E}$ , we create a dataset with each batch as  $[g, h_{t-1}, a_t, r]^{1:B}$ . Given sparse rewards (i.e.  $r_T = 1$ ), we use temporal discounting  $\delta \in (0, 1)$  to assign discounted rewards to previous actions in the trajectory<sup>5</sup>. This ensures that actions closer to the end receive higher rewards and vice versa. Specifically,  $r_{T-1} = \delta, r_{T-2} = \delta^2$ , and so on. We also sample negative actions from other paths (akin to the Can model) with a reward of 0. The model is trained to align the discounted reward of the action and the predicted reward from  $\mathcal{M}^{\text{pay}}$  by minimizing the mean squared error (MSE) loss  $\frac{1}{B} \sum_{i=1}^B (r^i - \mathcal{M}^{\text{pay}}(g^i, h_{t-1}^i, a_t^i))^2$ . The model uses an uncased Bert plus a regression layer whose output is bounded in  $[0, 1]$  via a sigmoid activation. The input format is the same as the Can model. The output is the estimated payoff,  $f_{\text{heur}}(h_t, g) = \mathcal{M}^{\text{pay}}(g, h_{t-1}, a_t)$ . I/O examples:

**Input** ⟨Goal⟩ pick up the purple box. ⟨Initial State⟩ Room 1 has yellow key, agent. Room 2 has purple box. The door connecting Room 1 and Room 2 is locked. ⟨Step 1⟩ pick up yellow key. ⟨Step 2⟩ toggle yellow door. ⟨Step 3⟩ drop key in void. ⟨Step 4⟩ pick up blue box. ⟨NXT⟩ done picking up.

**Output** 1.0 // end of plan

**Input** ⟨Goal⟩ pick up the purple box. ⟨Initial State⟩ Room 1 has yellow key, agent. Room 2 has purple box. The door connecting Room 1 and Room 2 is locked. ⟨Step 1⟩ pick up yellow key. ⟨Step 2⟩ toggle yellow door. ⟨Step 3⟩ drop key in void. ⟨NXT⟩ pick up blue box.

**Output** 0.6 //  $\delta \cdot r$

**Input** ⟨Goal⟩ pick up the purple box. ⟨Initial State⟩ Room 1 has yellow key, agent. Room 2 has purple box. The door connecting Room 1 and Room 2 is locked. ⟨Step 1⟩ pick up yellow key. ⟨Step 2⟩ toggle yellow door. ⟨Step 3⟩ drop key in void. ⟨NXT⟩ pick up green box.

**Output** 0 // very low payoff

<sup>5</sup> $\delta$  for the Pay model training is unrelated to the POMDP.

Environment	Example Goal	Example Initial Observation	Plan Length	$ \mathcal{A} $
Ravens (Tower of Hanoi seq)	Move the gray disk in rod 2	Blue disk on top of gray disk. Gray disk on top of green disk. Green disk in rod 1. The disks can be moved in rod 1, rod 2, rod 3.	3.3	7.5
Ravens (Put Blocks in Bowls)	Put the yellow blocks in gray bowls	There is a gray bowl 1, gray bowl 2, gray bowl 3, yellow block 1, yellow block 2, yellow block 3, blue bowl 1, red block 1, green bowl 1, orange block 1.	6.1	25
BabyAI (Pickup)	Pick up the ball	Room 1 has purple ball. Room 2 has yellow key, agent. Room 3 has red key. The door connecting Room 1 and Room 2 is locked. The door connecting Room 2 and Room 3 is locked.	6.7	7.7
VirtualHome	Read book		5.9	150

Table 2: Table displays tasks from each environment, average plan length, and average action space size  $|\mathcal{A}|$ . For VirtualHome, we do not specify an initial observation since it is hard to describe a room environment. Here, the action space varies with episodes, depending for instance on the number of objects.

Setup	Say Model	Greedy-Token	Greedy-Action			Beam-Action		
			Say	SayCan	SayCanPay	Say	SayCan	SayCanPay
Ravens (tower of hanoi)	Vicuna	45	48	48	50	54	68	<b>70</b>
	Flan-T5	30	30	39	42	38	<b>50</b>	<b>50</b>
Ravens (put blocks in bowls)	Vicuna	30	51	52	54	52	52	<b>56</b>
	Flan-T5	96	96	96	96	<b>98</b>	<b>98</b>	<b>98</b>
BabyAI (pickup)	Vicuna	59	62	81	88	72	<b>94</b>	<b>94</b>
	Flan-T5	0	0	30	<b>36</b>	1	<b>36</b>	30
VirtualHome	Vicuna	0	32	49	52	48	52	<b>53</b>
	Flan-T5	0	0	30	48	30	41	<b>50</b>

Table 3: Table shows the *planning success* (i.e. the number of plans out of 100 that reached the goal within limited steps) on the test split across different environments using Vicuna, Flan-T5 models. It can be observed that the best decoding strategy is Beam-Action and the best decoding score is SayCanPay.

## Experimental Setup

### Say Model

The Say model does not undergo any fine-tuning and is used only for inference. We experimented with two types of transformer architectures. (i) **Decoder type**: 13b-parameter Vicuna model (Chiang et al. 2023) trained by fine-tuning LLaMA (Touvron et al. 2023). (ii) **Encoder-decoder type**: Flan-T5-11b (Chung et al. 2022) which is the instruction fine-tuned version of the T5 transformer (Raffel et al. 2020). Existing works have demonstrated the planning abilities of both the decoder type (Pallagani et al. 2022) and the encoder-decoder type architectures (Valmeekam et al. 2023, 2022). Since the generated plan is in free-form language and may contain unrecognizable (for the environment) words or incorrect syntax, it cannot be directly translated into actionable steps in the environment. Following Huang et al. (2022a), we use an exhaustive list of admissible actions (feasible and otherwise), and at the end of each action step, map the generated action to the closest admissible action using minimum edit distance. Interleaving action generation and mapping ensures that all subsequent steps are conditioned on admissible actions, thus mitigating compounding errors. We provide 3 examples (input goal and observation, output plan) to the model via few-shot prompting.

### Environments

We tested in three environments, detailed in Table 2.

- **Ravens** (Zeng et al. 2021) is a PyBullet simulated task set focusing on “pick and place”. It includes 10 tabletop tasks, of which we use two: (i) Tower of Hanoi (sequence), a variation of the classic puzzle focusing on specific intermediate goals, like moving a particular disk to a designated rod while keeping the traditional constraints. This creates more goal diversity; (ii) Put blocks in bowls requires placing blocks into bowls based on rules like *put yellow block in green bowls*. We adapt the environment for language tasks, observations, and actions.
- **BabyAI** (Chevalier-Boisvert et al. 2019) is a 2D-grid world environment where a bot is provided a language task sampled from a predefined grammar. We focus on *pickup* tasks where the agent navigates to collect an object, often unlocking doors or moving obstacles. Task difficulty varies with rooms, obstacles, and distractor objects. The agent’s actions include high-level commands like *pickup* and *drop* which are composed of atomic actions: “left”, “right”, “forward”, “pick”, and “drop” (see Figure 1)
- **VirtualHome** (Puig et al. 2018) is an interactive platform to simulate complex household activities via interactions with the environment, such as picking up objects, and switching on/off appliances. We utilize the VirtualHome-Env dataset (Liao et al. 2019), comprising daily household activities from 7 scenes gathered via crowdsourcing. We only use the goal as the input (see Table 2).

Setup	Say Model	Greedy-Token	Greedy-Action			Beam-Action		
			Say	SayCan	SayCanPay	Say	SayCan	SayCanPay
Ravens (tower of hanoi)	Vicuna	12	24	55	<b>58</b>	20	47	52
	Flan-T5	34	34	46	47	38	54	<b>56</b>
Ravens (put blocks in bowls)	Vicuna	16	36	40	48	38	42	<b>56</b>
	Flan-T5	63	65	71	<b>74</b>	67	<b>74</b>	<b>74</b>
BabyAI (pickup)	Vicuna	48	50	53	54	56	56	<b>62</b>
	Flan-T5	0	0	26	28	1	30	<b>34</b>
VirtualHome	Vicuna	0	14	23	29	20	26	<b>30</b>
	Flan-T5	0	0	6	15	4	19	<b>26</b>

Table 4: Table shows the *cost-effectiveness* (i.e. the number of plans out of 100 that reached the goal within limited steps and also had the same plan length as the expert plan) on the **test** split across different environments using Vicuna, Flan-T5 models. It can be observed that the best decoding strategy is Beam-Action and the best decoding score is SayCanPay.

Setup	Say Model	Greedy-Token	Greedy-Action			Beam-Action		
			Say	SayCan	SayCanPay	Say	SayCan	SayCanPay
Ravens (tower of hanoi)	Vicuna	32	30	18	18	27	<b>34</b>	<b>34</b>
	Flan-T5	24	22	18	16	<b>26</b>	<b>26</b>	<b>26</b>
Ravens (put blocks in bowls)	Vicuna	8	<b>30</b>	10	6	<b>30</b>	10	6
	Flan-T5	94	94	26	18	<b>96</b>	22	24
BabyAI (pickup)	Vicuna	0	1	4	<b>12</b>	9	<b>12</b>	10
	Flan-T5	0	1	<b>28</b>	<b>28</b>	1	15	<b>28</b>
VirtualHome	Vicuna	0/20	2/20	3/20	3/20	<b>5/20</b>	<b>5/20</b>	<b>5/20</b>
	Flan-T5	0/20	0/20	0/20	3/20	1/20	3/20	<b>5/20</b>

Table 5: Table shows the generalization results (i.e. the number of plans out of 100 that reached the goal) on *test-generalize* split across different environments using Vicuna and Flan-T5 models. It can be observed that Beam-Action outperforms other decoding strategies.

**Data Splits and Evaluation.** We aim to assess the success, cost-effectiveness, and out-of-distribution (OOD) generalization of the generated plans. We created three data splits for each environment using expert trajectories. (i) **train** split for Can, Pay model training and few-shot prompting of the Say Model; (ii) **test** split assesses the LM planners’ ability to generate successful plans (i.e. reach the goal within limited steps), and also the planners’ ability to generate cost-effective plans (i.e. plans that succeed and also have the same plan length as the expert plan<sup>6</sup>). (iii) **test-generalize** split focuses on the generalization capabilities like handling novel initial observations (e.g., unseen colors of blocks and bowls, distractors in BabyAI), longer sequence lengths (e.g., more blocks or disks in Ravens, more rooms in BabyAI), and unseen tasks in VirtualHome. All test splits have # total episodes = 100 unless specified otherwise. Moreover, all splits are disjoint (i.e. no overlap).

**Baselines.** At the action level, we evaluate our decoding scores (Say, SayCan, SayCanPay) using various decoding strategies (Greedy and Beam-Action). Therefore, our baselines employ a mix of these strategies and scores. For tokens, we use the Greedy-Token decoding strategy as a reference. Notably, Greedy-Action SayCan is the offline planning version of the original SayCan paper (Ahn et al. 2022).

<sup>6</sup>We split test into two parts of 100 samples to evaluate success, cost-effectiveness. For VirtualHome, we use the annotated plans from its dataset.

**Training and Inference Details.** We use 800 expert train trajectories for each Ravens task and 400 for BabyAI. For VirtualHome, we retained  $\approx 800$  compatible trajectories for the current simulator. An additional 100 expert trajectories were collected for each test split (20 for VirtualHome test-generalize). The Can and Pay models were trained on 7 NVIDIA-DGX V-100 GPUs using the Distributed Data-Parallel framework across 20 epochs. Training parameters included a  $1e-4$  learning rate, AdamW optimizer with  $1e-5$  weight decay, a batch size of 50, a train-validation split of 80-20. For inference, the Say model was loaded using Model Parallel on the same GPUs. Inference hyperparameters are listed in Table 6. Parameters like beam groups and diversity penalty encourage diversity among the beams, mitigating issues like returning multiple similar sequences. We used 8-bit precision for GPU-efficient model loading (Dettmers et al. 2022).

## Results

We analyze the results along the following axes: decoding strategies, decoding scores, and transformer architectures. We assessed planning success and generalization by executing the generated plans in simulators such as Ravens and BabyAI, which have built-in validation checks to determine goal achievement. For the more open-ended VirtualHome environment, we manually reviewed fully executed plans to ensure they met the intended task objectives. For cost-effectiveness, we acquired expert trajectories for each test sample using an oracle planner.




Figure 3: From left to right: Planning success, cost-effectiveness, generalization for different beam sizes. Note, that generalization on the test-generalize split for VirtualHome is reported as a percentage.

**Comparing decoding scores.** From Tables 3, 4, the performance across various decoding scores can be summarized as  $\text{Say} < \text{SayCan} \leq \text{SayCanPay}$ . **(i) planning success:** The SayCanPay and SayCan scores lead to comparable performances, often outperforming Say. The Pay model’s minor performance edge could be due to its focus on selecting actions based on long-term relevance, potentially avoiding irreversible (*breaking an egg*) or even absorbing states (*discharged cellphone*) from where it is impossible to reach the goal (i.e. planning is non-ergodic). **(ii) cost-effectiveness:** SayCanPay leads to a significant improvement over both Say ( $\approx 11 - 97\%$  for Beam-Action) and SayCan ( $\approx 0 - 33\%$  for Beam-Action and  $\approx 1 - 150\%$  for Greedy-Action). **(iii) generalization:** From Table 5, while the overall performance for SayCan and SayCanPay improves over Say, a noticeable drop in performance was observed for Ravens. This led to the hypothesis that the learned domain models (Can, Pay) are not generalizing to OOD data in certain environments.

**Comparing decoding strategies.** From Tables 3, 4, 5, the overall performance across decoding strategies follows the pattern: Greedy-Token < Greedy-Action < Beam-Action across all splits. The Beam-Action Say, SayCan, and SayCanPay versions show improvement over their corresponding Greedy-Action counterparts. **(i) planning success:** Beam-Action SayCanPay beats Greedy-Action SayCanPay by  $\approx 1 - 40\%$ . Similar gains are also observed with other decoding scores. **(ii) cost-effectiveness:** Beam-Action SayCanPay improves over Greedy-Action SayCanPay by  $\approx 0 - 73\%$ . **(iii) generalization:** Beam-Action SayCanPay beats Greedy-Action SayCanPay by  $\approx 0 - 89\%$ .

**Comparing Transformer Architectures.** We did not observe a consistent performance gain for any particular architecture, suggesting that either is apt for planning. We lack a definitive explanation, and further research is required to understand how each LM component impacts reasoning.

**Ablation Details.** We performed the following ablations for a deeper understanding of the individual modules.

- Effect of beam-size  $k$ : As seen in Figure 3, in general, both plan success and cost-effectiveness increases with increase in beam size with  $k = 1$  (Greedy-Action), 2, 3

(Beam-Action). All experiments used the SayCanPay decoding score. However, no clear patterns were observed for generalization results.

- Impact of Say Model: Planning failures may arise because the Say model fails to propose a right action amongst the candidates, making Can and Pay ineffective. We studied the Say model’s impact on overall performance using a *Perfect Say* that always recommends the correct action along with random distractors. From Table 7, we observed 16-84% improvements in SayCan and SayCanPay performance across various environments, indicating the potential of an improved Say model. Thus, using a larger model trained on more data could potentially enhance performance.
- Plan length comparison: We compute a *relative length* =  $\text{oracle plan length} / \text{generated plan length}$ , which compares the generated and oracle plan lengths. A value = 1 indicates equal lengths, and a value = 0, that the plan length is infinity (i.e. an unsuccessful plan). As shown in Figure 4, Beam-Action slightly improves over Greedy-Action. Furthermore, SayCanPay scoring achieves the best relative length ( $\approx 1$ ) for both Greedy and Beam-Action strategies signifying its cost-efficiency.
- Impact of problem size on planning time. Effect of action space: Planning time remains unaffected since the Say model generates rather than discriminates between actions. Effect of plan length: Greedy-Token run time increases by  $\sim 2\text{s}$  for each action step. Effect of decoding strategy:  $\sim 9\text{s}$  for Greedy-Token,  $\sim 17\text{s}$  for Greedy-Action,  $\sim 35\text{s}$  for Beam-Action. Effect of decoding score: Planning time is unaffected since the Can and Pay are small LMs with negligible overheads. Quantization techniques and advanced hardware can further reduce run time and is an active research area (Dettmers et al. 2023; Frantar et al. 2023).
- Qualitative Analysis: The Can model effectively selects feasible actions (Figure 1). The Pay model prioritizes actions that lead to quicker goal achievement. While Pay gives a high probability to the “done task” action linking it to plan completion, the Can score negates it due to unsatisfied “done task” preconditions.




Figure 4: [Best viewed in color] The error plot represents the variance in relative length over models Vicuna and Flan-T5. Due to the open-ended nature of VirtualHome, the crowdsourced trajectories are not optimal, which explains why certain cases have a relative length  $> 1.0$ . Note that Greedy-Token decoding in VirtualHome has a relative length = 0 since no generated plans were executed successfully for both Vicuna and Flan-T5.

Parameter	Value	Exceptions
max new tokens	10	11 Vicuna (Ravens-Blocks), 3 (VirtualHome)
beam groups	3	4 for Flan-T5 (BabyAI)
diversity penalty	2.0	
candidates ( $m$ )	6	8 for Flan-T5 (Baby-AI)
beam-size ( $k$ )	3	

Table 6: Inference hyperparameters. Here the candidates ( $m$ ) and the beam-size ( $k$ ) parameter are over actions. The rest of the beam search parameters are over tokens.

## Limitations and Future Work

The main limitations are (i) the need for expert trajectories to train domain models, and (ii) the domain models’ limited adaptability to OOD data. These challenges are inherent to deep learning models. However, recent advances in LLMs offer promising solutions. For example, newer studies have leveraged LLMs for reward design due to their ability to infer intentions from minimal prompts (Kwon et al. 2023). Notably, LLMs outperform smaller counterparts like Bert in generalization. Since both Can and Pay scores resemble rewards, future studies could use LLMs to mitigate training and improve generalization. Another potential direction could be to experiment with symbolic methods and non-parameterized heuristics like comparing the current generated plan with the successful/expert trajectories in the buffer.

## Conclusion

We proposed to combine the world knowledge and generative capabilities of LLMs with the systematicity of classical planning by formulating a heuristic search-based planning framework for LLMs. We demonstrated how to generate plans that are both feasible and cost-effective. While LLMs still cannot generate long-horizon plans on par with classical planners, our method overcomes issues inherent to LLM-based planning and extends traditional planning with the advantages of language models, marking significant progress

	Score	LM	Perfect
Ravens-Hanoi	SayCan	48	88
	SayCanPay	50	92
Ravens-Blocks	SayCan	52	70
	SayCanPay	54	75
BabyAI	SayCan	81	90
	SayCanPay	88	92
VirtualHome	SayCan	49	60
	SayCanPay	52	64

Table 7: The table depicts the impact of the Say model on planning success performance. In this context, both “LM” and “Perfect” represent Say models. “LM” corresponds to the Vicuna model, while “Perfect Say” is an oracle Say model that consistently proposes the correct action along with two other distractor actions as next candidates. For all experiments, we used the Greedy-Action decoding strategy.

for planning research with LLMs.

## Acknowledgements

This work was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation, and is also part of the EU H2020 ICT48 project “TAILOR” under contract 952215, and the KU Leuven Research Fund (C14/18/062).

## References

Ahn, M.; Brohan, A.; Brown, N.; Chebotar, Y.; Cortes, O.; David, B.; Finn, C.; Fu, C.; Gopalakrishnan, K.; Hausman, K.; Herzog, A.; Ho, D.; Hsu, J.; Ibarz, J.; Ichter, B.; Irpan, A.; Jang, E.; Ruano, R. J.; Jeffrey, K.; Jesmonth, S.; Joshi, N. J.; Julian, R.; Kalashnikov, D.; Kuang, Y.; Lee, K.-H.; Levine, S.; Lu, Y.; Luu, L.; Parada, C.; Pastor, P.; Quiambao, J.; Rao, K.; Rettinghouse, J.; Reyes, D.; Sermanet, P.; Sievers, N.; Tan, C.; Toshev, A.; Vanhoucke, V.; Xia, F.; Xiao, T.; Xu, P.; Xu, S.; Yan, M.; and Zeng, A. 2022. Do As I Can,

- Not As I Say: Grounding Language in Robotic Affordances. arXiv:2204.01691.
- Bonet, B.; and Geffner, H. 2001. Planning as heuristic search. *Artificial Intelligence*, 129(1-2): 5–33.
- Brohan, A.; Brown, N.; Carbajal, J.; Chebotar, Y.; Chen, X.; Choromanski, K.; Ding, T.; Driess, D.; Dubey, A.; Finn, C.; Florence, P.; Fu, C.; Arenas, M. G.; Gopalakrishnan, K.; Han, K.; Hausman, K.; Herzog, A.; Hsu, J.; Ichter, B.; Irpan, A.; Joshi, N.; Julian, R.; Kalashnikov, D.; Kuang, Y.; Leal, I.; Lee, L.; Lee, T.-W. E.; Levine, S.; Lu, Y.; Michalewski, H.; Mordatch, I.; Pertsch, K.; Rao, K.; Reymann, K.; Ryoo, M.; Salazar, G.; Sanketi, P.; Sermanet, P.; Singh, J.; Singh, A.; Soricut, R.; Tran, H.; Vanhoucke, V.; Vuong, Q.; Wahid, A.; Welker, S.; Wohlhart, P.; Wu, J.; Xia, F.; Xiao, T.; Xu, P.; Xu, S.; Yu, T.; and Zitkovich, B. 2023. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. arXiv:2307.15818.
- Chevalier-Boisvert, M.; Bahdanau, D.; Lahlou, S.; Willems, L.; Saharia, C.; Nguyen, T. H.; and Bengio, Y. 2019. BabyAI: First Steps Towards Grounded Language Learning With a Human In the Loop. In *International Conference on Learning Representations*, volume 105.
- Chiang, W.-L.; Li, Z.; Lin, Z.; Sheng, Y.; Wu, Z.; Zhang, H.; Zheng, L.; Zhuang, S.; Zhuang, Y.; Gonzalez, J. E.; Stoica, I.; and Xing, E. P. 2023. Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90%\* ChatGPT Quality.
- Chung, H. W.; Hou, L.; Longpre, S.; Zoph, B.; Tay, Y.; Fedus, W.; Li, Y.; Wang, X.; Dehghani, M.; Brahma, S.; Webson, A.; Gu, S. S.; Dai, Z.; Suzgun, M.; Chen, X.; Chowdhery, A.; Castro-Ros, A.; Pellat, M.; Robinson, K.; Valter, D.; Narang, S.; Mishra, G.; Yu, A.; Zhao, V.; Huang, Y.; Dai, A.; Yu, H.; Petrov, S.; Chi, E. H.; Dean, J.; Devlin, J.; Roberts, A.; Zhou, D.; Le, Q. V.; and Wei, J. 2022. Scaling Instruction-Finetuned Language Models. arXiv:2210.11416.
- Dettmers, T.; Lewis, M.; Belkada, Y.; and Zettlemoyer, L. 2022. LLM.int8(): 8-bit Matrix Multiplication for Transformers at Scale. arXiv:2208.07339.
- Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2023. QLoRA: Efficient Finetuning of Quantized LLMs. arXiv:2305.14314.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Ding, Y.; Zhang, X.; Amiri, S.; Cao, N.; Yang, H.; Kaminiski, A.; Esselink, C.; and Zhang, S. 2023. Integrating action knowledge and LLMs for task planning and situation handling in open worlds. *Autonomous Robots*, 47(8): 981–997.
- Du, Y.; Liu, Z.; Li, J.; and Zhao, W. X. 2022. A Survey of Vision-Language Pre-Trained Models. arXiv:2202.10936.
- Frantar, E.; Ashkboos, S.; Hoefler, T.; and Alistarh, D. 2023. GPTQ: Accurate Post-Training Quantization for Generative Pre-trained Transformers. arXiv:2210.17323.
- Golowich, N.; Moitra, A.; and Rohatgi, D. 2022. Planning in Observable POMDPs in Quasipolynomial Time. arXiv:2201.04735.
- 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. arXiv:2305.14992.
- Helmert, M. 2006. The fast downward planning system. *Journal of Artificial Intelligence Research*, 26: 191–246.
- Huang, W.; Abbeel, P.; Pathak, D.; and Mordatch, I. 2022a. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International Conference on Machine Learning*, 9118–9147. PMLR.
- Huang, W.; Xia, F.; Shah, D.; Driess, D.; Zeng, A.; Lu, Y.; Florence, P.; Mordatch, I.; Levine, S.; Hausman, K.; and Ichter, B. 2023. Grounded Decoding: Guiding Text Generation with Grounded Models for Embodied Agents. arXiv:2303.00855.
- Huang, W.; Xia, F.; Xiao, T.; Chan, H.; Liang, J.; Florence, P.; Zeng, A.; Tompson, J.; Mordatch, I.; Chebotar, Y.; Sermanet, P.; Brown, N.; Jackson, T.; Luu, L.; Levine, S.; Hausman, K.; and Ichter, B. 2022b. Inner Monologue: Embodied Reasoning through Planning with Language Models. arXiv:2207.05608.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2): 99–134.
- Kwon, M.; Xie, S. M.; Bullard, K.; and Sadigh, D. 2023. Reward Design with Language Models. In *The Eleventh International Conference on Learning Representations*.
- Lakhotia, K.; Kharitonov, E.; Hsu, W.-N.; Adi, Y.; Polyak, A.; Bolte, B.; Nguyen, T.-A.; Copet, J.; Baevski, A.; Mohamed, A.; and Dupoux, E. 2021. On Generative Spoken Language Modeling from Raw Audio. *Transactions of the Association for Computational Linguistics*, 9: 1336–1354.
- Liang, J.; Huang, W.; Xia, F.; Xu, P.; Hausman, K.; Ichter, B.; Florence, P.; and Zeng, A. 2023. Code as Policies: Language Model Programs for Embodied Control. arXiv:2209.07753.
- Liao, Y.-H.; Puig, X.; Boben, M.; Torralba, A.; and Fidler, S. 2019. Synthesizing Environment-Aware Activities via Activity Sketches. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6284–6292.
- Lin, K.; Agia, C.; Migimatsu, T.; Pavone, M.; and Bohg, J. 2023. Text2Motion: from natural language instructions to feasible plans. *Autonomous Robots*, 47(8): 1345–1365.
- Liu, B.; Jiang, Y.; Zhang, X.; Liu, Q.; Zhang, S.; Biswas, J.; and Stone, P. 2023. LLM+P: Empowering Large Language Models with Optimal Planning Proficiency. arXiv:2304.11477.
- Pallagani, V.; Muppasani, B.; Murugesan, K.; Rossi, F.; Horesh, L.; Srivastava, B.; Fabiano, F.; and Loreggia, A. 2022. Plansformer: Generating Symbolic Plans using Transformers. arXiv:2212.08681.

- Puig, X.; Ra, K.; Boben, M.; Li, J.; Wang, T.; Fidler, S.; and Torralba, A. 2018. Virtualhome: Simulating household activities via programs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8494–8502.
- 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. *The Journal of Machine Learning Research*, 21(1): 5485–5551.
- 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 *International Conference on Robotics and Automation (ICRA)*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023. LLaMA: Open and Efficient Foundation Language Models. arXiv:2302.13971.
- Valmeekam, K.; Olmo, A.; Sreedharan, S.; and Kambhampati, S. 2022. Large Language Models Still Can’t Plan (A Benchmark for LLMs on Planning and Reasoning about Change). In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.
- Valmeekam, K.; Sreedharan, S.; Marquez, M.; Olmo, A.; and Kambhampati, S. 2023. On the Planning Abilities of Large Language Models (A Critical Investigation with a Proposed Benchmark). arXiv:2302.06706.
- van den Oord, A.; Li, Y.; and Vinyals, O. 2019. Representation Learning with Contrastive Predictive Coding. arXiv:1807.03748.
- Wang, Y.; Wang, W.; Joty, S.; and Hoi, S. C. 2021. CodeT5: Identifier-aware Unified Pre-trained Encoder-Decoder Models for Code Understanding and Generation. In Moens, M.-F.; Huang, X.; Specia, L.; and Yih, S. W.-t., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 8696–8708. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Xie, Y.; Yu, C.; Zhu, T.; Bai, J.; Gong, Z.; and Soh, H. 2023. Translating Natural Language to Planning Goals with Large-Language Models. arXiv:2302.05128.
- Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T. L.; Cao, Y.; and Narasimhan, K. 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. arXiv:2305.10601.
- Zeng, A.; Florence, P.; Tompson, J.; Welker, S.; Chien, J.; Attarian, M.; Armstrong, T.; Krasin, I.; Duong, D.; Sindhwani, V.; and Lee, J. 2021. Transporter Networks: Rearranging the Visual World for Robotic Manipulation. In *Proceedings of the 2020 Conference on Robot Learning*, volume 155 of *Proceedings of Machine Learning Research*, 726–747. PMLR.
- Ziegler, D. M.; Stiennon, N.; Wu, J.; Brown, T. B.; Radford, A.; Amodei, D.; Christiano, P.; and Irving, G. 2020. Fine-Tuning Language Models from Human Preferences. arXiv:1909.08593.