

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 NATURAL LANGUAGE PDDL (NL-PDDL) FOR OPEN-WORLD GOAL-ORIENTED COMMONSENSE REGRESSION PLANNING IN EMBODIED AI

Anonymous authors

Paper under double-blind review

ABSTRACT

Planning in open-world environments, where agents must act with partially observed states and incomplete knowledge, is a central challenge in embodied AI. Open-world planning involves not only sequencing actions but also determining what information the agent needs to sense to enable those actions. Existing approaches using Large Language Models (LLM) and Vision-Language Models (VLM) cannot reliably plan over long horizons and complex goals, where they often hallucinate and fail to reason causally over agent-environment interactions. Alternatively, classical PDDL planners offer correct and principled reasoning, but fail in open-world settings: they presuppose complete models and depend on exhaustive grounding over all objects, states, and actions; they cannot address misalignment between goal specifications (e.g., “heat the bread”) and action specifications (e.g., “toast the bread”); and they do not generalize across modalities (e.g., text, vision). To address these core challenges: (i) we extend symbolic PDDL into a flexible natural language representation that we term NL-PDDL, improving accessibility for non-expert users as well as generalization over modalities; (ii) we generalize regression-style planning to NL-PDDL with commonsense entailment reasoning to determine what needs to be observed for goal achievement in partially-observed environments with potential goal-action specification misalignment; and (iii) we leverage the lifted specification of NL-PDDL to facilitate open-world planning that avoids exhaustive grounding and yields a time and space complexity independent of the number of ground objects, states, and actions. Our experiments in three diverse domains — classical Blocksworld and the embodied ALFWorld environment with both textual and visual states — show that NL-PDDL substantially outperforms existing baselines, is more robust to longer horizons and more complex goals, and generalizes across modalities.

1 INTRODUCTION

Open-world planning¹, where the agent operates under partial observability and incomplete knowledge, is essential for embodied agents to perform real-world tasks. Such embodied environments are inherently open-world and involve reasoning over a multitude of objects, posing challenges for planning. For example, we consider Figure 1, where the agent is asked to “*toast the bread and leave it on a plate*”. To efficiently identify the sequence of actions that achieve this goal given only partial observations of the state and an incomplete environment model, the agent must find the relevant objects (e.g., *bread*, *toaster*, and *plate*) while judiciously ignoring irrelevant ones. It must also use its commonsense knowledge to reason that a *toaster can heat bread, but a pot of water cannot*.

The breakthrough of foundation models such as Large Language Models (LLM) and Vision-Language Models (VLM) has led to their increased usage in planning (Yao et al., 2023b; Li et al., 2024; Kong et al., 2024); yet, they come with critical limitations (Shojaee et al., 2025; Kambhampati et al., 2024). These models struggle to generate reliable long-horizon plans with sound causal reasoning about the interactions between the agent and its environment as they lack a mechanism to

¹In this work, we characterize *open-world* by partial observability and incomplete domain knowledge, distinguishing it from *open-world games*, which refers to open-ended tasks in expansive environments.




Figure 1: **Overview of the NL-PDDL framework: (Agent’s Actions)** The specification of actions available to the agent are provided in a natural language (NL) variant of PDDL that we introduce. **(Goal)** The user provides NL instructions that are automatically translated into an NL-PDDL goal specification. **(Initial State)** The initial state is given as an image. **(Planning Goal)** Open-world first-order planning proceeds top-down: starting from the goal, regressed subgoals are generated that must hold before each action can be applied (e.g., the goal *(bread) is heated* regresses to the action *toast (bread) using (toaster)*, which generates subgoals *the agent holds the (bread)* and *a toaster is available*). At each step, LLM-based entailment connects subgoals with the NL effects of candidate actions and reasons about object affordances (e.g., *a toaster can toast bread, but a pot of water cannot*). **(Selected Plan)** A VLM grounds the NL object names in a regressed subgoal to their corresponding entities in the initial state image that are then used to instantiate an actionable plan.

track state changes and project how actions alter the world (Valmeeckam et al., 2024). Moreover, they fail on complex goals involving multiple logical constraints (Goebel & Zips, 2025), and their black-box nature makes it hard to interpret or verify their generated plans (Aghzal et al., 2025). In contrast, classical planners, such as those based on the Planning Domain Definition Language (PDDL) (Aeronautiques et al., 1998), have traditionally served the role of planning in AI. While verifiably correct, these methods often presume all objects and relations are known in a perfect model that can be exhaustively grounded. These methods also cannot bridge misalignments between goals and action specifications (e.g., reasoning that the action *“toast the bread”* is able to achieve the goal *“heat the bread”*). Such limitations are too restrictive for the the open-world setting we address here.

To overcome the limitations of existing methods, we introduce NL-PDDL, a planning framework that combines the expressive flexibility of natural language (NL) with the formal guarantees of symbolic planning. An overview of the NL-PDDL workflow is provided in Figure 1 and its caption. Concretely, we make the following key contributions to open-world planning in embodied AI:

- We extend classical PDDL into NL-PDDL, a flexible representation that lets users specify goals and actions abstractly in NL. This flexibility reduces the specification burden (vs. rigid PDDL schemas) and tolerates semantically incomplete and syntactically imperfect domain and state descriptions while maintaining sound plan generation.
- We propose an open-world, regression-style NL-PDDL planner that uses LLM-based common-sense entailment over regressed subgoals to both infer the observations needed to achieve a target (e.g., *heat bread* → *find toaster*) and resolve misalignments between (sub)goals and action effects.

- 108 • We avoid exhaustive grounding via *lifted regression* (Reiter, 1991). Instead of reasoning about
109 specific objects (e.g., *pan3*, *toaster1*), lifted regression leverages the variables inherent in NL-
110 PDDL (e.g., “*there exists some x that can toast bread*”). Lifted plans are instantiated only when
111 suitable objects are found, making planning complexity independent of the number of objects and
112 abstractly capturing a lifted *conditional plan* of all initial state conditions that can achieve a goal.
- 113 • Across three diverse planning domains — Blocksworld (Valmeekam et al., 2022), ALFworld, and
114 ALFworld-Vision (Shridhar et al., 2021) — we show that lifted regression planning in NL-PDDL
115 achieves higher plan-success rates than strong baselines, remains robust as plan horizons increase,
116 and generalizes across both text and vision modalities.

118 2 METHODOLOGY


121 In this section, we present Natural Language PDDL (**NL-PDDL**), a framework for open-world,
122 goal-oriented planning by leveraging regression-style reasoning enhanced with LLM-based com-
123 monsense reasoning. While classical PDDL provides a sound framework for long-horizon planning,
124 it cannot facilitate commonsense reasoning unless such knowledge is explicitly encoded by a domain
125 designer. For example, a classical PDDL solver cannot infer that `isToasted(bread)` implies
126 `isHeated(bread)`, even though it is intuitively clear to humans that “*if bread is toasted, it is*
127 *likely heated*”. NL-PDDL builds on the formal framework of classical PDDL but replaces sym-
128 bolic predicates, variables, and objects with typed NL counterparts, enabling seamless integration
129 of LLMs to perform commonsense entailment inferences during planning. For a formal review of
130 PDDL and first-order regression planning extended here, we refer the reader to Appendix B.

131 **Natural Language Representation of PDDL.** In classical PDDL, predicates, objects, and
132 variables are represented with rigid symbols. NL-PDDL replaces these symbols with typed
133 NL counterparts. For instance, the symbolic predicate `isToasted(bread)` becomes
134 “*the (bread) is toasted*”|“*bread*”:“*food*”, while “*the (?o:food) is toasted*” denotes its lifted form.
135 Here, `?o:food` indicates that `(?o)` is of type `food`.

136 **NL-PDDL Problem Definition.** NL-PDDL problems are defined like classical PDDL problems,
137 but aimed at open-world planning. We define an NL-PDDL problem as a tuple $\mathcal{P} = \langle G, \mathcal{A}, \mathcal{F}, h \rangle$,
138 where G is a conjunctive formula expressing the agent’s goal, and \mathcal{F} is the set of lifted NL predicates
139 in the domain. \mathcal{A} is the set of actions $a(\vec{y})$, i.e., parameterized first-order operators. Each $a(\vec{y}) \in \mathcal{A}$
140 is defined by: preconditions $a(\vec{y}).\text{pre} \subseteq \mathcal{F}$, which must hold for the action to execute; add effects
141 $a(\vec{y}).\text{add} \subseteq \mathcal{F}$, which become true after the action is executed; and delete effects $a(\vec{y}).\text{del} \subseteq \mathcal{F}$, which
142 no longer hold after execution. The planning horizon $h \in \mathbb{N}$ is the maximum length of feasible plans.
143 The goal in NL-PDDL is to construct a *conditional plan* Π that maps subgoal formulas to action
144 sequences that can achieve them:

$$145 \Pi = [(\psi_1(\vec{x}_1), \langle a^{1,1}(\vec{x}^{1,1}), \dots, a^{1,n_1}(\vec{x}^{1,n_1}) \rangle), \dots, (\psi_k(\vec{x}_k), \langle a^{k,1}(\vec{x}^{k,1}), \dots, a^{k,n_k}(\vec{x}^{k,n_k}) \rangle)], \quad (1)$$

147 where each $\psi_i(\vec{x}_i)$ is a first-order formula representing a subgoal, and $\langle a^{i,1}(\vec{x}^{i,1}), \dots, a^{i,n_i}(\vec{x}^{i,n_i}) \rangle$
148 is the corresponding sequence of actions (i.e., a plan) with length $n_i \leq h$, and each $\vec{x}^{j,k} \subseteq \vec{x}^j$. As an
149 example, the plan derived by NL-PDDL for the problem in Figure 1 is presented in Figure 2.
150



160 Figure 2: The conditional plan derived by NL-PDDL for the problem in Figure 1.
161

162
163


2.1 FIRST ORDER REGRESSION IN NL-PDDL

164
165
166
167
168
169


Effect Axioms. Following the methodology for first-order regression defined in Appendix B, we first need to transform our NL-PDDL description into the *effect axioms* that are used to perform regression planning. Specifically, following the notation of Appendix B, we build axioms of the form of Equations 8 and 9, and provide a concrete example of $\gamma^+_{(\text{?o}) \text{ is toasted}}(\text{?o}, \text{?r})$ for action “*toast(?o) using (?r)*”, and predicate “*(?o) is toasted*” from the NL-PDDL description in Figure 3. In English, this says an object “*(?o)*” is toasted in the next state if the agent holds

170
171
172
173
174
175
176
177
178
179
180
181

Running Example:

Positive Effect Axiom for  with respect to 

$$\forall (?o): \text{"object", (?r): "object", a} \left(a = \text{"toast(?o) using (?r)"} \wedge \underbrace{\left(\text{"(?r) can toast (?o)"} \wedge \text{"the agent holds (?o)"} \right)}_{\gamma^+_{(\text{?o}) \text{ is toasted}}, \text{"toast(?o) using (?r)"}^{(\text{?o}, \text{?r})}} \Rightarrow \underbrace{\text{"(?o) is toasted"}_{\text{Add Effect}}}_{'} \right).$$

Negative Effect Axiom for  with respect to 

$$\forall (?o): \text{"object", (?r): "receptacle", a} \left(a = \text{"put(?o) on (?r)"} \wedge \underbrace{\left(\text{"the agent is near (?r)"} \wedge \text{"the agent holds (?o)"} \right)}_{\gamma^-_{\text{"the agent holds (?o)"}}, \text{"put(?o) on (?r)"}^{(\text{?o}, \text{?r})}} \Rightarrow \neg \underbrace{\text{"the agent holds (?o)"}_{\text{Delete Effect}}}_{'} \right).$$

182
183

Figure 3: Example of positive and negative effect axiom based on NL-PDDL problem in Figure 1.

184
185
186
187
188

it, and the agent toasts “*(?o)*” using object “*(?r)*” that can toast “*(?o)*”. To illustrate negative effect axioms, we construct $\gamma^-_{\text{"the agent holds (?o)"}, \text{"put(?o) on (?r)"}^{(\text{?o}, \text{?r})}}$ for action “*put(?o) on (?r)*” and predicate “*the agent holds (?o)*” in Figure 3. In English, this says if the agent holds an object “*(?o)*”, is near receptacle “*(?r)*”, and puts “*(?o)*” on “*(?r)*”, then the agent no longer holds “*(?o)*” in the next state.

189
190
191
192

These positive (negative) effect axioms are only constructed for predicates and actions if the predicate is in the add (delete) effect of the action in the NL-PDDL description. Cases that a predicate appears in the effects of multiple actions are handled via disjunction (see Appendix C for details).

193
194
195
196

Successor State Axioms. As described in Appendix B, successor state axioms (SSAs) are the workhorse of lifted regression planning in that they allow the replacement of a predicate in a post-action state with the necessary subgoals required to achieve it. Let $F'_{a_i}(\vec{y}_F)$ denote the value of a predicate $F(\vec{y}_F)$ in the next state after executing action $a_i(\vec{y})$. The SSA for $F(\vec{y}_F)$ is defined as:

197

$$F'_{a_i}(\vec{x}) \equiv \gamma^+_{F, a_i}(\vec{x}) \vee (F(\vec{x}) \wedge \neg \gamma^-_{F, a_i}(\vec{x})). \quad (2)$$

198
199
200

In English, an SSA states that a predicate F is true in the successor state *iff* it was made true by a positive effect, or it was already true and not made false by a negative effect.

201
202
203
204


We construct the SSAs for each action-predicate pair. While we leave complete formal details to the Appendix B due to space limitations, we provide the SSA for the predicate “*is toasted(?o)*” and action “*toast (?o) using (?r)*”, and the predicate “*the agent possesses(?o)*” and action “*put (?o) on (?r)*” in Figure 4 as examples to illustrate the process of SSA construction:

205
206
207
208
209
210
211
212

Running Example:

Successor State Axiom for  and  is:

$$\forall (?o): \text{"object", (?r): "object", a} \left[(\text{"(?o) is toasted"})' = (\text{"(?o) is toasted"}) \vee \left(\underbrace{\left(a = \text{"toast (?o) using (?r)"} \wedge (\text{"(?r) can toast (?o)"} \wedge \text{"the agent holds (?o)"}) \right)}_{\text{Positive Effect Axiom}} \right) \right]$$

Successor State Axiom for  and  is:

$$\forall (?o): \text{"object", (?r): "receptacle", a} \left[(\text{"the agent holds (?o)"}') \equiv (\text{"the agent holds (?o)"} \wedge \neg \underbrace{\left(a = \text{"put (?o) on (?r)"} \wedge (\text{"the agent is near (?r)"}) \right)}_{\text{Negative Effect Axiom}}) \right]$$

213
214
215

Figure 4: Examples of constructing SSAs based on NL-PDDL problem in Figure 1.

216
217

After constructing the SSAs, NL-PDDL recursively regresses each subgoal through applicable actions to construct the conditional plan Π . We explain the regression procedure in the next section.

216
217

2.1.1 REGRESSION OF POTENTIALLY MISALIGNED PREDICATES IN NL-PDDL

218
219
220
221

In regression planning with classical PDDL, goal predicates are only matched to action effects that share the same predicate names. In NL, such exact matches rarely occur due to variability in phrasing, instead an action effect may entail a goal literal, e.g., “the bread is toasted” entails (\vdash) “the bread is heated”. In this section, we extend lifted regression planning using LLM entailment (\vdash_{LLM}).

222
223
224
225
226
227
228
229

Regression of Positive NL Predicates. Let $P'(\vec{x})|\vec{x}:T_x$ be a predicate in the next state. We aim to regress it through an action a_i , to derive pre-action conditions such that if they hold, executing a_i results in a next state that entails $P'(\vec{x})$. To this end, we first identify a predicate $F(\vec{z}) \in \mathcal{F}|\vec{z}:T_z$ that can be unified with $P'(\vec{x})$. In PDDL, unification makes the two predicates syntactically identical through variable substitution, but NL-PDDL allows a *more general form of unification based on commonsense entailment*: predicates $F(\vec{z})$ and $P'(\vec{x})$ are unifiable *iff* they have equal arities and a substitution θ exists that induces an injective mapping between their variables such that:

230
231
232
233

- (i) For every pair of corresponding types $t_z \in T_z$ and $t_x \in T_x$ mapped by θ , $t_x \vdash_{\text{LLM}} t_z$.
- (ii) After applying θ , the substituted predicate $F(\vec{z})$ entails the substituted goal predicate $P'(\vec{x})$, i.e., $F(\vec{z})\theta \vdash_{\text{LLM}} P'(\vec{x})\theta$.²

234
235
236

A legal typed unification results in a substitution θ that is a mapping of variables to terms (e.g., $\theta = \{x/y\}$, so that $F(x)\theta \equiv F(y)$), which then permits a misaligned regression to proceed. Formally, let $F'_{a_i}(\vec{z}) \in \Phi$ denote the SSA formula for $F'(\vec{z})$ with respect to a_i under the unifier θ :

237
238

$$\text{Regr}_{\vdash}(P'(\vec{x}), a_i) \equiv F'_{a_i}(z)\theta. \quad (3)$$

239
240

With Regr_{\vdash} , we can now formally regress a predicate through an action with misaligned effects:


241
242
243
244
245
246

Figure 5: An example of a regression step of a positive predicate with goal–action misalignment.

247
248
249
250
251
252
253
254

In the example of Figure 5, to identify the subgoals necessary for the predicate “(Bread:food) is heated” to hold after an action is executed, we aim to find a predicate unifiable with it. The predicate “(?o:object) is toasted” qualifies, with $\theta = \{?o/ "bread"\}$, because (i) type consistency holds (“food” \vdash_{LLM} “object”), and (ii) after substitution, (“bread is toasted” \vdash_{LLM} “bread is heated”). This allows us to derive the regressed subgoal condition (given by the SSA) that can achieve “(bread) is heated”, which is how subgoal ψ_3 is derived from ψ_2 in the conditional plan example of Figure 2.

255
256
257
258
259

Regression of Negative NL predicates. As for negative predicates, we regress a negated predicate $\neg P'(\vec{x})|\vec{x}:T_x$ through an action a_i to derive preconditions that ensure $\neg P'(\vec{x})$ holds after a_i is executed. Thus, we identify a predicate $F(\vec{z}) \in \mathcal{F}|\vec{z}:T_z$ that can be unified with $P'(\vec{x})$. Let θ be the unifier inducing an injective mapping between predicate variables such that:

260
261
262
263
264

- (i) For every pair of corresponding types $t_x \in T_x$ and $t_z \in T_z$ mapped by θ , $t_x \vdash_{\text{LLM}} t_z$.
- (ii) After applying θ , the substituted goal predicate entails the substituted precondition predicate, i.e., $P'(\vec{x})\theta \vdash_{\text{LLM}} F(\vec{z})\theta$.

265
266
267

Critically note that the direction of entailment in check (ii) for negative NL predicates reverses from the case for positive NL predicates. Let $F'_{a_i}(\vec{z}) \in \Phi$ be the SSA formula for $F'(\vec{z})$, with respect to

268
269

²This relaxation from syntactic equivalence to entailment-based unification allows, for example, unifying “(?x:food) is on (?y:plate)” with “(?y:receptacle) is under (?x:object)”, as they are both binary and with $\theta = \{?x/?y, ?y/?x\}$, we have “food” \vdash_{LLM} “object”, “plate” \vdash_{LLM} “receptacle”, and “(?x) is on (?y)” \vdash_{LLM} “(?y) is under (?x)”.

270 a_i . Applying θ to this formula and negating the result gives the regressed condition for $\neg P'(\vec{x})$ with
 271 respect to a_i :

$$\text{Regr}_{\vdash}(\neg P'(\vec{x}), a_i) \equiv \neg F'_{a_i}(\vec{z})\theta. \quad (4)$$

273 We provide a worked example illustrating the regression of a negative NL predicate in Appendix D.
 274

275 **NL Regression with Multiple Entailments.** A predicate $P'(\vec{x})|\vec{x} : \vec{T}$ may be semantically en-
 276 tailed by multiple other predicates. For example, the goal “($?o$) is cooked” may be entailed by
 277 either “($?o$) is toasted” or “($?o$) is boiled”. In such cases, we aggregate all entailing predicates.
 278

279 Let $F_{\vdash}^+ = \{F_1(\vec{x}_1), \dots, F_m(\vec{x}_m) \mid F_i(\vec{x}_i) \in \mathcal{F} \wedge F_i(\vec{x}_i)\{\vec{x}_i/\vec{x}\} \vdash_{\text{LLM}} P'(\vec{x})\}$ be the set of all
 280 predicates that entail $P'(\vec{x})$. The aggregated regression of $P'(\vec{x})$ with respect to action a_i is:

$$\text{Regr}_{\vdash}(P'(\vec{x}), a_i) \equiv \bigvee_{j=1}^m F'_{a_i,j}(\vec{x}), \quad \text{where } F_j \in F_{\vdash}^+. \quad (5)$$

281 The same construction applies to negated predicates, with entailment checked in the reverse direc-
 282 tion. Let $F_{\vdash}^- = \{F_1(\vec{x}_1), \dots, F_m(\vec{x}_m) \mid F_i(\vec{x}_i) \in \mathcal{F} \wedge P'(\vec{x})\{\vec{x}_i/\vec{x}\} \vdash_{\text{LLM}} F_i(\vec{x}_i)\}$ be the set of all
 283 predicates entailed by $P'(\vec{x})$. The aggregated regression of $\neg P'(\vec{x})$ with respect to a_i is:

$$\text{Regr}_{\vdash}(\neg P'(\vec{x}), a_i) \equiv \bigwedge_{j=1}^m \neg F'_{a_i,j}(\vec{x}), \quad \text{where } F'_j \in F_{\vdash}^-. \quad (6)$$

291 2.1.2 REGRESSION OF FORMULAS

292 Full first-order formulas can now be regressed
 293 using Algorithm 1: the formula is rewritten
 294 in Disjunctive Normal Form (DNF), the pred-
 295 icates in each disjunct are regressed follow-
 296 ing the process in Section 2.1.1, and the re-
 297 sults are incrementally combined in DNF. The
 298 general regression process in the NL-PDDL
 299 framework is summarized in Appendix E.

300 3 EXPERIMENTS

301 We pose the following research questions:

Algorithm 1 REGRESSFORMULA(ψ' , $a_i(\vec{y}_i)$, \mathcal{F})

```

1: Input:  $\psi'$  in DNF ;  $a_i(\vec{y}_i)$ ;  $\mathcal{F}$ 
2: Output:  $\psi = \text{Regr}_{\vdash}(\psi', a_i)$  in DNF
3:  $\psi \leftarrow \emptyset$ 
4: for each disjunct  $C \in \psi$  do
5:    $\mathcal{D} \leftarrow \top$ 
6:   for each literal  $F_j(\vec{x}_j)$  in  $C$  do
7:     Construct:  $F_{\vdash}^+, F_{\vdash}^-$ 
8:      $F'_{a_i,j}(\vec{x}) \leftarrow \text{Regr}_{\vdash}(F_j(\vec{x}), a_i)$ 
9:      $\mathcal{D} \leftarrow \mathcal{D} \wedge F'_{a_i,j}(\vec{x})$ 
10:     $\mathcal{D} \leftarrow \text{ConvertToDNF}(\mathcal{D})$ 
11:   end for
12:    $\psi \leftarrow \psi \vee \mathcal{D}$ 
13: end for
14: return  $\psi$ 

```

305 **RQ1: Overall Planning Performance.** How does NL-PDDL³ perform in terms of success rate
 306 and computational cost across modalities (e.g., text, images), compared to baselines?

307 **RQ2: Reasoning over Language Misalignment.** How well does NL-PDDL handle planning
 308 tasks where NL descriptions of the action specifications and the goal are misaligned?

310 **RQ3: Effect of Task Complexity.** How does NL-PDDL compare to baselines as task complexity
 311 increases, measured by optimal plan length and the number of constraints in the goal?

313 3.1 EXPERIMENTAL SETUP

314 **Open-World Tasks.** We conduct open-world planning experiments on *ALFWorld Text* (Shridhar
 315 et al., 2021) and *ALFWorld Vision* (cf. Figure 6) to evaluate performance across modalities. Agents
 316 must operate under partial observability and incomplete task-relevant knowledge to complete com-
 317 mon household tasks in virtual home environments. It supports a rich space of interactions, requiring
 318 the agents to reason about object relations and action affordances without hardcoded annotations.
 319 Following prior work (Yang et al., 2024; Shridhar et al., 2021; Yao et al., 2023b; Wang et al., 2023),
 320 we evaluate NL-PDDL on 135 out-of-distribution tasks with 50 movement budget (cf. Appendix G).
 321 We also extend both ALFWorld benchmarks by introducing misaligned language descriptions of the
 322 agent’s action model and goals (cf. Appendix H).

323 ³https://anonymous.4open.science/r/nl_pddl_planner-4C67/README.md




Figure 6: An Example of a Visual ALFWorld Task.

Table 1: (RQ1) Performance of different methods in two **open-world** datasets **without misalignment** between goal and model.[†]Reported performance in the original work.

Open-World Without Misalignment						
Datasets	Methods	# Tok.	# Expert Trajectory	Fine-Tuning	SR (%)	
ALFWorld Text	Direct LLM-based Planner	GPT-4o	1,358,934	0	✗	21%
		Gemini-2.0-Flash	1,176,406	0	✗	16%
		LLaMA-3.1	1,193,640	0	✗	19%
	Reflective LLM-based Planner	ReAct (w examples)	5,507,266	0	✗	81%
		ReAct (w model)	4,242,560	0	✗	34%
		Reflexion-3 [†]	NA	0	✗	83%
		Reflexion-10 [†]	NA	0	✗	91%
		DEPS [†]	NA	0	✗	76%
	Fine-tuned LM	BUTLER	NA	100,000	✓	26%
	Ours	NL-PDDL	443,407	0	✗	94%
ALFWorld Vision	Direct VLM-based Planner	GPT-4o	1,823,539	0	✗	8%
		Gemini-2.0-Flash	1,440,025	0	✗	2%
		LLaMA-3.1	1,505,371	0	✗	2%
	Fine-tuned VLM	LLaMA-Adapter [†]	NA	170,000	✓	13%
		InstructBLIP [†]	NA	170,000	✓	22%
		EMMA-3 [†]	NA	15,237	✓	37%
		EMMA-10 [†]	NA	15,237	✓	82%
	Ours	NL-PDDL	679,148	0	✗	84%

Table 2: (RQ1) Performance comparison across **closed-world** Blocksworld variants **without misalignment**. Direct LLM-based planners fail to generate valid plans for Mystery and Randomized Blocksworld. NL-PDDL demonstrates consistent and robust performance across all variants.

Closed-World Without Misalignment						
Method	Blocksworld		Mystery Blocksworld		Randomized Blocksworld	
	# Tok.	SR (%)	# Tok.	SR (%)	# Tok.	SR (%)
GPT-4o	982,602	34%	834,890	0%	835,974	0%
Gemini-2.0 Flash	928,024	18%	834,060	1%	835,824	0%
LLaMA 3.1	963,565	44%	841,791	0%	842,622	0%
Fast Downward	N/A	100%	N/A	100%	N/A	100%
NL-PDDL	0	70%	0	70%	0	70%

Closed-World Tasks. *Blocksworld* is a widely used closed-world planning domain in which the task is to rearrange colored blocks, initially stacked or on a table, to achieve a given goal. We adopt a variant introduced in Valmeeckam et al. (2022) by converting goal specifications into NL descriptions for NL-based planners. We include two additional variants to probe brittleness to surface form and lexical priors by replacing object and goal names in the NL descriptions with random English words (*Mystery Blocksworld*) and random strings (*Randomized Blocksworld*). Beyond these, we introduce *Misalignment Blocksworld*, where the NL descriptions of action model and goals are contextually meaningful but intentionally misaligned, to assess the commonsense entailment capabilities of planners (cf. Appendix H for details).

Baselines and Evaluation Metrics. We compare NL-PDDL with SOTA baselines in open- and closed-world domains. In the open-world settings (i.e., ALFWorld Text and Vision), we compare with (i) **Direct** LLM/VLM-based planners: GPT-4o (Hurst et al., 2024), Gemini-2.0 Flash (Team et al., 2023), and LLaMA-3.1 (Grattafiori et al., 2024), (ii) **Reflective** LLM-based planners: Re-

378 Table 3: **(RQ2)** Performance across settings **with misalignment** in open-world ALFWorld (Text-
 379 Vision) and closed-world Misalignment Blocksworld. Reported-results-only methods are omitted
 380 since prior works did not consider misalignment.

Open-World With Misalignment							
Category	Method	ALFWorld Text		ALFWorld Vision		Blocksworld	
		# Tok.	SR (%)	# Tok.	SR (%)	# Tok.	SR (%)
Direct LLM	GPT-4o	1,440,470	17% (↓5%)	1,811,404	7% (↓1%)	937,905	27% (↑7%)
	Gemini-2.0 Flash	1,251,533	15% (↓1%)	1,262,016	5% (↑3%)	939,950	23% (↑5%)
	LLaMA-3.1	1,310,664	15% (↓4%)	1,480,392	2% (↓0%)	950,490	42% (↓3%)
ReAct	w/ examples	5,215,612	79% (↓12%)	N/A	N/A	N/A	N/A
	w/ model	4,428,914	23% (↓11%)	N/A	N/A	N/A	N/A
Classical Planner	Fast Downward	N/A	N/A	N/A	N/A	0	0%
Ours	NL-PDDL	501,049	91% (↓3%)	745,365	80% (↓4%)	11,656	70%

391 Act (Yao et al., 2023b), Reflexion (Shinn et al., 2023), and DEPS (Wang et al., 2023), and (iii)
 392 **Fine-tuned LLM/VLMs:** BUTLER (Shridhar et al., 2021), LLaMA-Adapter (Gao et al., 2023),
 393 InstructBLIP (Dai et al., 2023), and EMMA (Yang et al., 2024), all trained on ALFWorld expert
 394 demonstrations. In the closed-world setting, we compare against Direct LLM-based planners and
 395 the classical Fast Downward (Helmert, 2006) planner, which serves as an upper bound on planning
 396 performance. The detailed implementation of NL-PDDL on the ALFWorld and Blocksworld bench-
 397 marks is included in Appendix I. Methods are evaluated under a feasible maximum runtime. For
 398 evaluation, we consider two aspects: task *Success Rate* (SR), and associated cost, as captured by the
 399 # *Tokens* Usage (# Tok.) in LLM.

4 RESULTS AND DISCUSSION

400 **RQ1.** In open-world tasks (cf. Table 1), NL-PDDL achieves superior SR on both ALFWorld Text
 401 and Vision compared to all baselines. On ALFWorld Text, NL-PDDL achieves a 94% SR with
 402 only 443K tokens, whereas ReAct consumes over ten times more tokens while reaching a lower SR.
 403 Although Reflexion-10 achieves performance comparable to NL-PDDL, it depends on repeating
 404 ten trials of the same task, which is infeasible for many embodied AI applications. Impressively,
 405 NL-PDDL even outperforms all fine-tuned VLMs that rely on thousands of expert trajectories in
 406 visual ALFWorld. Our results highlight its generalizability across different modalities for open-
 407 world tasks. In closed-world tasks (cf. Table 2), NL-PDDL demonstrates robust performance with
 408 a consistent 70% SR at 0 tokens. In contrast, direct LLM-based planners entirely fail to solve the
 409 Mystery and Randomized variants, performing with less than 1% SR despite substantial token usage.
 410 Fast Downward, a closed-world planner, achieves a 100% success rate on the three Blocksworld
 411 domains. However, it is brittle under goal–action misalignment, as we demonstrate in RQ3.

412 **RQ2.** In Table 3, we evaluate the performance of different methods on open- and closed-world
 413 tasks with misalignment between the NL descriptions of action models and goals. NL-PDDL
 414 demonstrates clear robustness in the face of such misalignments. In open-world tasks, NL-PDDL
 415 maintains a 91% SR on ALFWorld Text and 80% on ALFWorld Vision. By contrast, ReAct,
 416 whether explicitly prompted to focus on action models or supported with few-shot examples, suf-
 417 fers notable degradation, underscoring its sensitivity to misalignment. The closed-world Misalign-
 418 ment Blocksworld variant highlights the gap: Fast Downward, though highly effective in purely
 419 symbolic and aligned settings, does not perform commonsense entailment and fails under goal–
 420 action schema misalignment; in contrast, NL-PDDL maintains performance comparable to other
 421 Blocksworld variants, evidencing robust misalignment handling.

422 **RQ3.** We evaluate planner performance relative to optimal plan depth. NL-PDDL achieves per-
 423 fect success up to depth 6, with performance declining to 84% at depth 8 and fails beyond depth
 424 10 under the runtime limit. By contrast, LLM-based planners exhibit consistent deterioration across
 425 all depths. NL-PDDL’s decline stems from scalability limitations under the predefined maximum
 426 runtime, while LLM-based planners fail due to their inability to reason over an extended horizon.
 427 The Fast Downward planner maintains perfect performance when the action model and goals are
 428 fully aligned, but fails on the Misalignment Blocksworld where commonsense entailment reason-
 429 ing is required. In open-world tasks where optimal depth is inconsistent, we measure NL-PDDL’s
 430




Figure 7: (RQ3) SR by goal complexity and optimal plan depth in ALFWorld and Blocksworld across planners

performance against goal complexity. NL-PDDL remains stable, showing only a 3% average drop across ALFWorld Text and Vision, while LLM/VLM-based planners degrade much more sharply, underscoring their sensitivity to increasing goal complexity. Together, these results demonstrate the robustness of NL-PDDL with respect to goal complexity across open- and closed-world tasks.

5 RELATED WORK

Eliciting Stronger Reasoning from LLMs: As LLMs scale, they remain prone to hallucinations and logical fallacies (Tommoy et al., 2024; Zhang et al., 2025) resulting in new prompting strategies to mitigate these concerns. Step-by-step methods (Wei et al., 2022; Kojima et al., 2022; Yao et al., 2023a) encourage incremental reasoning by breaking problems into sub-steps or branching search trajectories. Reflective methods such as ReAct (Yao et al., 2023b) and Reflexion (Shinn et al., 2023) add self-reflection, critiquing, and revision. Yet these methods often fail on tasks with several logical constraints or long-horizon dependencies (Valmeekam et al., 2023b; 2022; Song et al., 2025); are sensitive to prompt design and few-shot examples, iteration dependent (Pan et al., 2023); and remain unverifiable since the LLM acts as a black box without guarantees of soundness (Shanahan, 2024).

Planning with LLMs and Symbolic Planners: A separate body of work integrates LLMs with classical planners. One approach treats LLMs as planners that generate action sequences from symbolic task descriptions (Silver et al., 2022; Song et al., 2023; Silver et al., 2024). Another uses LLMs as *translators*, converting NL problems and domains into structured formats for symbolic solvers (Guan et al., 2023; Kambhampati et al., 2024; Oswald et al., 2024; Yang et al., 2023). Both approaches face limitations: LLM-as-planner approaches are unreliable due to reasoning errors and hallucinations, while translator-based methods often introduce faulty predicates, actions, or logical mappings of NL descriptions that cause downstream planning failures. Translation also imposes an expressivity bottleneck, restricting applicability to problems encodable in symbolic form. Finally, most assume closed-world settings, making them inappropriate for open-world planning.

Our Point of Departure: We introduce a novel NL-PDDL hybrid natural language variation of PDDL to facilitate the use of powerful commonsense (V)LLM reasoning during planning, and blend it with symbolic *lifted* regression over NL-PDDL to address open-world, long-horizon planning.

We further discuss NL-to-PDDL translation approaches and clarify their distinctions from NL-PDDL in Appendix J.

6 CONCLUSION

We proposed NL-PDDL, a framework for open-world planning that extends symbolic PDDL to support NL, improving accessibility for non-experts. NL-PDDL facilitates the integration of LLM commonsense knowledge into open-world regression planning to address goal-action misalignment,

486 while retaining soundness and verifiability. Empirically, we showed that lifted regression planning
 487 in NL-PDDL achieves higher plan-success rates in contrast to existing strong baselines, remains
 488 robust as plan horizons increase, and generalizes well across both text and vision modalities.
 489

490 **REFERENCES**
 491

492 Constructions Aeronautiques, Adele Howe, Craig Knoblock, ISI Drew McDermott, Ashwin Ram,
 493 Manuela Veloso, Daniel Weld, David Wilkins Sri, Anthony Barrett, Dave Christianson, et al.
 494 Pddl—the planning domain definition language. *Technical Report, Tech. Rep.*, 1998.

495 Mohamed Aghzal, Erion Plaku, Gregory J Stein, and Ziyu Yao. A survey on large language models
 496 for automated planning. *arXiv preprint arXiv:2502.12435*, 2025.

497 Pierre Carbonnelle. pydatalog: Logic programming in python. <https://sites.google.com/site/pydatalog/home> or <https://github.com/pcarbonn/pyDatalog>,
 498 2024.

501 Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li,
 502 Pascale N Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models
 503 with instruction tuning. *Advances in neural information processing systems*, 36:49250–49267,
 504 2023.

505 Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu,
 506 Conghui He, Xiangyu Yue, et al. Llama-adapter v2: Parameter-efficient visual instruction model.
 507 *arXiv preprint arXiv:2304.15010*, 2023.

509 Alfonso Emilio Gerevini. An introduction to the planning domain definition language (pddl): Book
 510 review. *Artificial Intelligence*, 280:103221, 2020.

511 Kai Goebel and Patrik Zips. Can llm-reasoning models replace classical planning? a benchmark
 512 study. *arXiv preprint arXiv:2507.23589*, 2025.

514 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 515 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd
 516 of models. *arXiv preprint arXiv:2407.21783*, 2024.

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

520 Patrik Haslum, Nir Lipovetzky, Daniele Magazzeni, Christian Muise, Ronald Brachman, Francesca
 521 Rossi, and Peter Stone. *An introduction to the planning domain definition language*, volume 13.
 522 Springer, 2019.

524 Malte Helmert. The fast downward planning system. *Journal of Artificial Intelligence Research*, 26:
 525 191–246, 2006.

526 Yongfeng Huo, Jing Tang, Yinghui Pan, Yifeng Zeng, and Langcai Cao. Learning a planning domain
 527 model from natural language process manuals. *IEEE Access*, 8:143219–143232, 2020.

529 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-
 530 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint*
 531 *arXiv:2410.21276*, 2024.

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

535 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large
 536 language models are zero-shot reasoners. *Advances in neural information processing systems*,
 537 35:22199–22213, 2022.

539 Xiangrui Kong, Wenxiao Zhang, Jin Hong, and Thomas Braunl. Embodied ai in mobile robots:
 Coverage path planning with large language models. *arXiv preprint arXiv:2407.02220*, 2024.

540 Manling Li, Shiyu Zhao, Qineng Wang, Kangrui Wang, Yu Zhou, Sanjana Srivastava, Cem Gokmen,
 541 Tony Lee, Erran Li Li, Ruohan Zhang, et al. Embodied agent interface: Benchmarking llms for
 542 embodied decision making. *Advances in Neural Information Processing Systems*, 37:100428–
 543 100534, 2024.

544 Shivam Miglani and Neil Yorke-Smith. Nltopddl: one-shot learning of pddl models from natu-
 545 ral language process manuals. In *Proc. of the ICAPS Workshop on Knowledge Engineering for*
 546 *Planning and Scheduling (KEPS). ICAPS*, 2020.

547 Christian Muise, Florian Pommerening, Jendrik Seipp, and Michael Katz. Planutils: Bringing plan-
 548 ning to the masses. In *ICAPS 2022 System Demonstrations*, 2022.

549 James Oswald, Kavitha Srinivas, Harsha Kokel, Junkyu Lee, Michael Katz, and Shirin Sohrabi.
 550 Large language models as planning domain generators. In *Proceedings of the International Con-*
 551 *ference on Automated Planning and Scheduling*, volume 34, pp. 423–431, 2024.

552 Liangming Pan, Alon Albalak, Xinyi Wang, and William Yang Wang. Logic-lm: Empower-
 553 ing large language models with symbolic solvers for faithful logical reasoning. *arXiv preprint*
 554 *arXiv:2305.12295*, 2023.

555 Raymond Reiter. Equality and domain closure in first-order databases. *Journal of the ACM (JACM)*,
 556 27(2):235–249, 1980.

557 Raymond Reiter. The frame problem in the situation calculus: A simple solution (sometimes)
 558 and a completeness result for goal regression. *Artificial intelligence and mathematical theory of*
 559 *computation: papers in honor of John McCarthy*, 50:359–380, 1991.

560 Murray Shanahan. Talking about large language models. *Communications of the ACM*, 67(2):68–79,
 561 2024.

562 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:
 563 Language agents with verbal reinforcement learning. *Advances in Neural Information Processing*
 564 *Systems*, 36:8634–8652, 2023.

565 Parshin Shojaee, Iman Mirzadeh, Keivan Alizadeh, Maxwell Horton, Samy Bengio, and Mehrdad
 566 Farajtabar. The illusion of thinking: Understanding the strengths and limitations of reasoning
 567 models via the lens of problem complexity. *arXiv preprint arXiv:2506.06941*, 2025.

568 Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Cote, Yonatan Bisk, Adam Trischler, and Matthew
 569 Hausknecht. Alfworld: Aligning text and embodied environments for interactive learning. In
 570 *International Conference on Learning Representations*, 2021.

571 Tom Silver, Varun Hariprasad, Reece S Shuttleworth, Nishanth Kumar, Tomás Lozano-Pérez, and
 572 Leslie Pack Kaelbling. Pddl planning with pretrained large language models. In *NeurIPS 2022*
 573 *foundation models for decision making workshop*, 2022.

574 Tom Silver, Soham Dan, Kavitha Srinivas, Joshua B Tenenbaum, Leslie Kaelbling, and Michael
 575 Katz. Generalized planning in pddl domains with pretrained large language models. In *Pro-*
 576 *ceedings of the AAAI conference on artificial intelligence*, volume 38, pp. 20256–20264, 2024.

577 Pavel Smirnov, Frank Joublin, Antonello Ceravola, and Michael Gienger. Generating consistent
 578 pddl domains with large language models. *arXiv preprint arXiv:2404.07751*, 2024.

579 Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su.
 580 Llm-planner: Few-shot grounded planning for embodied agents with large language models. In
 581 *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 2998–3009, 2023.

582 Peiyang Song, Pengrui Han, and Noah Goodman. A survey on large language model reasoning
 583 failures. In *2nd AI for Math Workshop@ ICML 2025*, 2025.

584 Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut,
 585 Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly
 586 capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

594 SM Tonmoy, SM Zaman, Vinija Jain, Anku Rani, Vipula Rawte, Aman Chadha, and Amitava Das.
 595 A comprehensive survey of hallucination mitigation techniques in large language models. *arXiv*
 596 *preprint arXiv:2401.01313*, 6, 2024.

597

598 Karthik Valmeekam, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. Large lan-
 599 guage models still can't plan (a benchmark for llms on planning and reasoning about change). In
 600 *NeurIPS 2022 Foundation Models for Decision Making Workshop*, 2022.

601

602 Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kamb-
 603 hampati. Planbench: An extensible benchmark for evaluating large language models on planning
 604 and reasoning about change, 2023a. URL <https://arxiv.org/abs/2206.10498>.

605

606 Karthik Valmeekam, Matthew Marquez, Sarath Sreedharan, and Subbarao Kambhampati. On the
 607 planning abilities of large language models-a critical investigation. *Advances in Neural Infor-
 608 mation Processing Systems*, 36:75993–76005, 2023b.

609

610 Karthik Valmeekam, Kaya Stechly, and Subbarao Kambhampati. Llms still can't plan; can lrms? a
 611 preliminary evaluation of openai's o1 on planbench. *arXiv preprint arXiv:2409.13373*, 2024.

612

613 Zihao Wang, Shaofei Cai, Guanzhou Chen, Anji Liu, Xiaojian Ma, and Yitao Liang. Describe,
 614 explain, plan and select: Interactive planning with large language models enables open-world
 615 multi-task agents. *arXiv preprint arXiv:2302.01560*, 2023.

616

617 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 618 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in
 619 neural information processing systems*, 35:24824–24837, 2022.

620

621 Yaqi Xie, Chen Yu, Tongyao Zhu, Jinbin Bai, Ze Gong, and Harold Soh. Translating natural lan-
 622 guage to planning goals with large-language models. *arXiv preprint arXiv:2302.05128*, 2023.

623

624 Yijun Yang, Tianyi Zhou, Kanyue Li, Dapeng Tao, Lusong Li, Li Shen, Xiaodong He, Jing Jiang,
 625 and Yuhui Shi. Embodied multi-modal agent trained by an llm from a parallel textworld. In
 626 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 26275–
 627 26285, 2024.

628

629 Zhun Yang, Adam Ishay, and Joohyung Lee. Coupling large language models with logic program-
 630 ming for robust and general reasoning from text. *arXiv preprint arXiv:2307.07696*, 2023.

631

632 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik
 633 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Ad-
 634 vances in neural information processing systems*, 36:11809–11822, 2023a.

635

636 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.
 637 React: Synergizing reasoning and acting in language models. In *International Conference on
 638 Learning Representations (ICLR)*, 2023b.

639

640

641 Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao,
 642 Yu Zhang, Yulong Chen, et al. Siren's song in the ai ocean: A survey on hallucination in large
 643 language models. *Computational Linguistics*, pp. 1–46, 2025.

644

645 A ETHICS, REPRODUCIBILITY, LLM USAGE STATEMENT

646 A.1 ETHICS STATEMENT

647 We adhere to the ICLR Code of Ethics. This work does not involve human or animal subjects, and no
 648 identifiable personal information was used in developing this paper. Therefore, there are no privacy
 649 or security concerns. We are committed to maintaining integrity and transparency throughout the
 650 research process.

648 A.2 REPRODUCIBILITY STATEMENT
649650 We have made every effort to ensure that the methodology presented in this work is reproducible.
651 All code used in this study has been released in an anonymous repository to support reproducibility
652 and verifiability. Full descriptions of the methods and experimental setup are provided in both the
653 main text and the appendix to facilitate transparent evaluation and replication of our results.
654655 A.3 THE USE OF LARGE LANGUAGE MODELS (LLMs)
656657 Large Language Models (LLMs) were used to aid or polish the writing of this paper. Specifically,
658 we employed LLMs to improve the readability and clarity of the manuscript.
659660 B FUNDAMENTALS OF REGRESSION PLANNING IN PDDL
661663 In this section, we review PDDL and its representation of planning problems, then introduce the first-
664 order regression approach to planning. The corresponding operations for our proposed NL-PDDL
665 are presented in Section 2.1 of the main paper.
666

667 B.1 PDDL PLANNING REPRESENTATION.

669 A planning problem is defined as the tuple $\mathcal{P} = \langle \mathcal{F}, \mathcal{O}, \mathcal{A}, s_0, G \rangle$, where:670

- \mathcal{F} is the set of first-order predicates that describe properties of objects or relations be-
671 tween them, whose truth values may change as actions are applied, such as `boiled(x)`,
672 `can_toast(x, y)`, etc.
- \mathcal{O} is the set of objects in the domain, which serve as the constants over which predicates
673 and actions are instantiated, (e.g., `potato`, `plate`, etc.).
- \mathcal{A} is the set of actions $a_i(\vec{y})$, i.e., parameterized first-order operators. Each action schema
674 is defined by:
 - *Preconditions*: Denoted by $a_i(\vec{y})\text{.pre} \subset \mathcal{F}$ is the set of predicates that must hold for
675 the action to be applicable.
 - *Add effects*: Denoted by $a_i(\vec{y})\text{.add} \subset \mathcal{F}$ is the set of predicates that become true
676 once the action is executed.
 - *Delete effects*: Denoted by $a_i(\vec{y})\text{.del}$ is the set of predicates that are no longer true
677 after the action is executed.

678 For instance, the action `pickup(?o)` has the preconditions `hand_empty` and
679 `near(?o)`, the add effect `possess(?o)`, and the delete effect `hand_empty`.
680681

- s_0 is the initial state, given as a set of ground predicates that hold true before any action is
682 executed.
- G is the goal condition, specified as a set of predicates that must be satisfied in a terminal
683 state.

685 In classical planning, problems are often specified in the *Planning Domain Definition Language*
686 (*PDDL*), which is divided into two parts: the *domain* and the *problem instance*. The domain defines
687 the agent’s *action model* together with the set of predicates \mathcal{F} and the set of action schemas \mathcal{A} . The
688 *problem instance* specifies a finite object set \mathcal{O} , an initial state s_0 and G given as a set of ground
689 predicates.
690691 The objective of solving a planning problem is to find a *conditional plan*, i.e., a finite sequence of
692 actions $\pi = \langle a^{(1)}, \dots, a^{(n)} \rangle$, such that, starting from s_0 , executing each action in π sequentially
693 results in a state s_n that satisfies the goal G . Here, the superscript (i) in $a^{(i)}$ indicates that a is the
694 action executed at the i -th step of the plan, i.e., it represents the horizon index of the action within
695 the sequence.
696

756 These axioms are essential for accurately regressing a goal through an action: to determine the
 757 conditions that must have held before the action, one needs to know exactly how each action modifies
 758 the state. To obtain these axioms, we make the following assumptions.
 759

760 • No new variables outside of action parameter are introduced in preconditions and effects.
 761 • No quantifiers are used in the action’s preconditions and effects.
 762 • All unquantified variables are implicitly universally quantified.
 763

764 We assume positive and negative effect axioms can be specified by considering all of the ways in
 765 which each action can affect each predicate. Let $F(\vec{y}_F)$ be a predicate in the current state, $\gamma_{F,a_i}^+(\vec{y}, s)$
 766 is a first-order formula such that, if it holds in the current state s , then $F'(\vec{y}_F)$ holds after executing
 767 $a_i(\vec{y})$. Similarly, let $\gamma_{F,a_i}^-(\vec{y}, s)$ be a first-order formula such that, if it holds in the current state s ,
 768 then $\neg F'(\vec{y}_F)$ holds after executing $a_i(\vec{y})$. We write $\gamma_{F,a_i}^+(\vec{y})$ and $\gamma_{F,a_i}^-(\vec{y})$ as shorthands for the
 769 state-dependent conditions $\gamma_{F,a_i}^+(\vec{y}, s)$ and $\gamma_{F,a_i}^-(\vec{y}, s)$ in the current state, respectively, and define
 770 the normal form of effect axioms as following:
 771

772
$$\forall \vec{y} : \vec{T} [\gamma_{F,a_i}^+(\vec{y}) \Rightarrow F'_{a_i}(\vec{y}_F)], \quad \forall \vec{y} : \vec{T} [\gamma_{F,a_i}^-(\vec{y}) \Rightarrow \neg F'_{a_i}(\vec{y}_F)]. \quad (7)$$

773 Here, $\vec{y} : \vec{T}$ indicates that each variable in \vec{y} is assigned its type by a corresponding element in \vec{T} .
 774 Given a PDDL problem \mathcal{P} with an action set \mathcal{A} , we can derive effect axioms in the aforementioned
 775 normal form, action by action. Consider an action $a(\vec{y}) \in \mathcal{A}$, and let $F(\vec{y}_F)$ be a predicate in its
 776 positive effect set, i.e., $F(\vec{y}_F) \in a(\vec{y}).\text{add}$. The predicate parameters \vec{y}_F correspond to a permutation
 777 of a subset of the action parameters \vec{y} . We use $\vec{y}_{\setminus F}$ to denote the remaining action parameters that
 778 do not appear in \vec{y}_F which are existentially quantified. With this setup, we construct the following
 779 implication:
 780

781
$$\text{for each } F(\vec{y}_F) \in a_i(\vec{y}).\text{add} : \forall \vec{y} : \vec{T}, a [a = a_i(\vec{y}) \wedge \underbrace{\bigwedge_{\text{Pre}_j(\vec{y}_j) \in a_i(\vec{y}).\text{pre}} \text{Pre}_j(\vec{y}_j)}_{\gamma_{F,a_i}^+(\vec{y})}] \Rightarrow F'_{a_i}(\vec{y}_F).$$

782
$$(8)$$

783 Similarly, for the predicates in the delete effects of the action, we have:
 784

785
$$\text{for each } F(\vec{y}_F) \in a_i(\vec{y}).\text{del} : \forall \vec{y} : \vec{T}, a [a = a_i(\vec{y}) \wedge \underbrace{\bigwedge_{\text{Pre}_j(\vec{y}_j) \in a_i(\vec{y}).\text{pre}} \text{Pre}_j(\vec{y}_j)}_{\gamma_{F,a_i}^-(\vec{y})}] \Rightarrow \neg F'_{a_i}(\vec{y}_F).$$

786
$$(9)$$

787 We can combine multiple predicates with the same name via a disjunction. The detailed methodology
 788 is outlined in Appendix C.
 789

790 **Positive effect axiom.** For $\text{isHot}(\text{?o})$ with respect to the action $\text{toast}(\text{?o}, \text{?r})$:

791
$$\forall \text{?o:object}, \text{?r:object}, a \left(\underbrace{a = \text{toast}(\text{?o}, \text{?r}) \wedge \text{can_toast}(\text{?o}, \text{?r}) \wedge \text{possess}(\text{?o})}_{\gamma_{\text{isHot}, \text{toast}}^+(\text{?o}, \text{?r})} \Rightarrow \text{isHot}'(\text{?o}) \right).$$

792 **Negative effect axiom.** For $\text{possess}(\text{?o})$ with respect to the action $\text{puton}(\text{?o}, \text{?r})$:

793
$$\forall \text{?o:object}, \text{?r:receptacle}, a \left(\underbrace{a = \text{puton}(\text{?o}, \text{?r}) \wedge \text{near}(\text{?r}) \wedge \text{possess}(\text{?o})}_{\gamma_{\text{possess}, \text{puton}}^-(\text{?o}, \text{?r})} \Rightarrow \neg \text{possess}'(\text{?o}) \right).$$

810 **Succesor State Axioms** When a PDDL problem is transformed into the above normal form, Reiter
 811 (1991) showed that, under the assumptions of the *Unique Names Axioms* (Reiter, 1980) and the
 812 *Explanation Closure Axioms* (Reiter, 1991), we can construct *Successor State Axioms* (SSAs) that
 813 capture how a predicate may change or persist as the agent interacts with the environment. Let
 814 $F'_{a_i}(\vec{y}_F)$ denote the value of a predicate $F(\vec{y}_F)$ in the next state after executing action $a_i(\vec{y})$. The
 815 SSA for $F(\vec{y}_F)$ is defined as:

$$817 \quad \forall \vec{y} : \vec{T} [F'_{a_i}(\vec{y}_F) \equiv \gamma_{F,a_i}^+(\vec{y}) \vee (F(\vec{y}_F) \wedge \neg \gamma_{F,a_i}^-(\vec{y})]. \quad (10)$$

819 Intuitively, SSAs state that a predicate $F'(\vec{y}_F)$ in the next state can be true either because it is made
 820 true by an action $a_i(\vec{y})$, as specified by $\gamma_{F,a_i}^+(\vec{y})$, or because it was already true in the previous state,
 821 $F(\vec{y}_F)$, and the action does not make it false, i.e., $\neg \gamma_{F,a_i}^-(\vec{y})$.

$$823 \quad \forall ?o:object, ?r:receptacle, a$$

$$824 \quad \left[\text{possess}'(?o) \equiv \underbrace{(a = \text{pickup}(?o) \wedge \text{hand_empty} \wedge \text{near}(?o))}_{\gamma_{\text{possess}, \text{pickup}}^+ (?o)} \right.$$

$$825 \quad \left. \vee \left(\text{possess}(?o) \wedge \underbrace{\neg(a = \text{puton}(?o, ?r) \wedge \text{near}(?r))}_{\gamma_{\text{possess}, \text{puton}}^- (?o, ?r)} \right) \right].$$

832 **First-Order Regression** Let ψ' denote a first-order state description that holds *after* executing an
 833 action a_i . The regression operator $Regr(\psi', a_i)$ “backprojects” ψ' to compute a logical formula
 834 ψ that must hold *before* the execution of a_i . Fortunately, SSAs provide us a logically equivalent
 835 pre-action condition for each predicate $F'_{a_i}(\vec{x}) \in \psi'$ with respect to an action $a_i \in \mathcal{A}$. Regression of
 836 the entire formula ψ' is performed by recursively replacing each post-action predicate $F'(\vec{x})$ with
 837 its corresponding precondition formula $F'_{a_i}(\vec{x})$, as defined by the appropriate SSA.

$$839 \quad 840 \quad Regr(F'(\vec{x}), a_i) \equiv F'_{a_i}(\vec{x}) \quad (11)$$

$$842 \quad Regr(\text{isHot} (?o), \text{toast} (?o, ?r)) \equiv \text{can_toast} (?o, ?r) \wedge \text{possess} (?o)$$

844 SSAs lay the foundation of lifted regression planning, since using SSAs, we are able to replace
 845 a predicate in a post-action state with the subgoals that are required to hold before the action is
 846 executed to achieve the predicate. We provide examples of this functionality in Section 2.1.

848 C COMBINING EFFECT AXIOMS

850 **Combining Effect Axioms(Appendix)** For predicates that appear in multiple action effects, we
 851 need to combine them into a single effect axiom. For instance, we have two antecedent condition
 852 formulae $C_1(\vec{x}_1, \vec{y}_1)$ and $C_2(\vec{x}_2, \vec{y}_2)$ where:

$$854 \quad \forall \vec{x}_1 : \vec{T}_{x_1}, \vec{y}_1 : \vec{T}_{y_1} [C_1(\vec{x}_1, \vec{y}_1) \Rightarrow F(\vec{x}_1)], \quad (12)$$

$$855 \quad \forall \vec{x}_2 : \vec{T}_{x_2}, \vec{y}_2 : \vec{T}_{y_2} [C_2(\vec{x}_2, \vec{y}_2) \Rightarrow F(\vec{x}_2)]. \quad (13)$$

857 let $\theta = \{\vec{x}_2 \mapsto \vec{x}_1, \vec{y}_2 \mapsto \vec{y}_1\}$ be the most general unifier (MGU) of C_1 , and C_2 , i.e., the substitution
 858 that unifies the variables of C_2 with those of the C_1 without introducing unnecessary restrictions.
 859 We then apply θ to $C_2(\vec{x}_2, \vec{y}_2)$, i.e., $\text{SUBST}(\theta, C_2(\vec{x}_2, \vec{y}_2)) = C_2(\vec{x}_1, \vec{y}_1)$, and form the following
 860 implication:

$$862 \quad \forall \vec{x}_1 : \vec{T}_{x_1}, \vec{y}_1 : \vec{T}_{y_1} [C_1(\vec{x}_1, \vec{y}_1) \vee C_2(\vec{x}_1, \vec{y}_1)] \Rightarrow \gamma_{F,a}^+(\vec{x}_1, \vec{y}_1), \quad (14)$$

863 where $\text{SUBST}(\theta, p)$ denotes the formula obtained by applying substitution θ to p .

```

864
865    $\forall a : \mathcal{A}, ?o : \text{object}, ?r : \text{object} [$ 
866      $(a = \text{bake}(?o, ?r) \wedge \text{baking\_device}(?r) \wedge \text{can\_bake}(?o, ?r) \wedge \text{possess}(?o))$ 
867
868      $\vee$ 
869      $(a = \text{boil}(?o, ?r) \wedge \text{boiling\_device}(?r) \wedge \text{full\_of\_water}(?r) \wedge \text{possess}(?o))]$ 
870
871      $\Rightarrow \text{isHot}'(?o)$ 
872
873
874
875
876
877
```

D EXAMPLE FOR REGRESSION OF NEGATIVE NL PREDICATES




Figure 8: An example of a regression step of a negative predicate with goal–action misalignment.

The process of regressing a negative NL predicate is summarized in Figure 8. To identify the subgoals necessary for the predicate $\neg \text{"the agent possesses(bread:food)"}$ to hold after an action is executed, we aim to find a predicate that is unifiable with it. The predicate $\neg \text{"the agent holds(?o:object)"}$ qualifies for this, under the substitution $\theta = \{?o/\text{"bread"}\}$, because **(i)** type consistency condition holds ($\text{"food"} \vdash_{LLM} \text{"object"}$), and **(ii)** once the substitution is applied, ($\text{"the agent possesses(bread)"} \vdash_{LLM} \text{"the agent holds(bread)"}$). Note that the direction of the entailment in condition **(ii)** is the opposite of the case of positive predicates. The post-action predicate $\neg \text{"the agent possesses(bread:food)"}$ can thus be unified with the post-action predicate $\neg \text{"the agent holds(?o:object)"}$, which is in turn equivalent to its SSA with the action "pick up (?o)" , which yields the new subgoal.

E NL-REGRESSION ALGORITHM

Algorithm 2 presents the regression process in the NL-PDDL framework. Starting from G , we standardize actions and regress the current subgoal ψ' through them. The result is simplified and if valid, the action is appended to the plan, forming (ψ, \vec{a}) . This process is repeated to depth h , yielding the final plan Π which consists of all such pairs (ψ, \vec{a}) where ψ is logically equivalent to the original goal G .

918 **Algorithm 2** NL-REGRESSION ALGORITHM

```

919
920 1: Input:  $T = \langle G, \mathcal{A}, h, \mathcal{F} \rangle$ 
921 2: Output:  $\Pi \subseteq \{(\psi, \vec{a})\}$  with each  $\psi$  in DNF
922 3:  $\Phi = \{F'_{a_i}(\vec{x}) \mid F(\vec{x}) \in \mathcal{F}, a_i(\vec{y}) \in \mathcal{A}\}$ .
923 4:  $Frontier \leftarrow \{(G, [], 0)\}$ 
924 5:  $\Pi \leftarrow \{(G, [], \emptyset)\}$ 
925 6: while  $Frontier \neq \emptyset$  do
926 7:   Extract  $(\psi', \vec{a}, d)$  from  $Frontier$ 
927 8:   if  $d \geq h$  then
928 9:     continue
929 10:   end if
930 11:   for each  $a_i(\vec{y}) \in \mathcal{A}$  do
931 12:      $a_i(\vec{y}) \leftarrow \text{STANDARDIZE}(a_i(\vec{y}))$  //See Appendix F.1
932 13:      $\psi \leftarrow \text{REGRESSFORMULA}(\psi', a_i(\vec{y}), \mathcal{F})$ 
933 14:     if  $\psi$  is evaluated to False then
934 15:       continue
935 16:     end if
936 17:      $\psi \leftarrow \text{SIMPLIFY}(\psi)$  //See Appendix F.2
937 18:      $\vec{a}' \leftarrow \vec{a} \parallel [a_i(\vec{y})]$ 
938 19:      $Frontier \leftarrow Frontier \cup \{(\psi, \vec{a}', d+1)\}$ 
939 20:      $\Pi \leftarrow \Pi \cup \{(\psi, \vec{a}')\}$ 
940 21:   end for
941 22: end while
942 23: return  $\Pi$ 

```

943 **F** OPERATIONS IN FIRST-ORDER REGRESSION944 **F.1** STANDARDIZATION

945 Let φ be a first-order formula. The *standardization* of φ , denoted $\text{STANDARDIZE}(\varphi)$, is obtained by
946 renaming the bound variables in φ with fresh variables so that no two distinct quantifiers in φ bind
947 the same variable symbol.

948 Formally, if x_1, \dots, x_k are the bound variables in φ , then

$$949 \text{STANDARDIZE}(\varphi) = \varphi[\rho(x_1), \dots, \rho(x_k)],$$

950 where ρ is a bijective renaming function mapping each x_i to a fresh variable x'_i such that $x'_i \notin \text{Var}(\varphi)$
951 for all i . Here, $\text{Var}(\varphi)$ denotes the set of variables occurring in the formula φ .

952 **F.2** SIMPLIFICATIONS

953 Regression typically generates a large, expanded DNF formula. To keep the resulting plan compact
954 and interpretable, we apply a set of simplification rules that maintain logical equivalence while
955 eliminating redundancy. These simplifications enable the planner to express subgoals clearly and
956 prevent repeated formulas from arising during regression.

957 **F.2.1** EQUALITY-BASED QUANTIFIER ELIMINATION

958 We eliminate quantified variables that are equal to an object by replacing them with the object, i.e.,

$$959 \exists x : T, [x = x^* \wedge F(x)] \equiv F(x^*)$$

960 where x^* is an object. This transformation preserves logical equivalence while eliminating redundant
961 quantification. It is often applied during regression, where it is triggered by the need to unify
962 variables between predicates and action parameters. This rule is especially valuable when the re-
963 gressed formula includes equality constraints that enforce correspondences between variables.

964 **F.2.2** CONTRADICTION EVALUATION

965 Contradictions often arise within regressed formulas and can be simplified to \perp (false). We consider
966 two main cases:

972 • **Conjunctive Formulas.** A conjunction that contains mutually exclusive literals or unsat-
 973 isifiable equality conditions simplifies to false. For example:
 974

$$\begin{aligned} F(x) \wedge \neg F(x) &\rightarrow \perp, \\ x = y \wedge x \neq y &\rightarrow \perp. \end{aligned}$$

977 These situations typically occur when inconsistent conditions are introduced through re-
 978 gression or substitution.

979 **Disjunctive Formulas.** In disjunctions, the false literal \perp can be eliminated since it does
 980 not affect the overall satisfiability:
 981

$$\perp \vee F(x) \vee \dots \equiv F(x) \vee \dots$$

983 F.2.3 NO NO-OP ASSUMPTION.

985 We assume that the domain does not contain *no-op actions*. A no-op action is an action that leaves
 986 the state unchanged, so we assume that every action makes a meaningful change to the state. Hence,
 987 any disjunct in the regressed formula that is logically identical to the original goal can be safely
 988 eliminated.

$$\text{Regr}(G(x)) \equiv \bigvee_i \Psi_i^{DNF} \vee G(x) \equiv \bigvee_i \Psi_i'^{DNF}$$

991 This guarantees that plans are constructed only from informative regressions.

993 F.2.4 DNF SUBSUMPTION.

994 Let ϕ be a DNF formula:
 995

$$\phi = \bigvee_{i=1}^n C_i,$$

998 in which each clause C_i is a conjunction of literals. We say clause C_i *subsumes* another clause C_j if
 999

$$C_i \Rightarrow C_j.$$

1001 In this simplification technique, we eliminate any clause C_j for which there exists another clause C_i
 1002 such that $C_i \Rightarrow C_j$. Formally, if
 1003

$$\exists i \neq j, \text{ such that } C_i \Rightarrow C_j,$$

1004 the formula
 1005

$$\phi = \bigvee_{i=1}^n C_i$$

1008 is simplified to
 1009

$$\phi' = \bigvee_i C_i \text{ such that } \forall j, k : (j \neq k) \Rightarrow \neg(C_j \Rightarrow C_k)$$

1011 For example, with DNF subsumption, we have
 1012

$$[F_1(\vec{x}_1) \wedge F_2(\vec{x}_2)] \wedge [F_1(\vec{x}_1) \wedge F_2(\vec{x}_2) \wedge F_3(\vec{x}_3)] \equiv [F_1(\vec{x}_1) \wedge F_2(\vec{x}_2)].$$

1015 This simplification allows the planner to just maintain the most general subgoal formulas while
 1016 preserving correctness.

1017 F.2.5 DUPLICATE DETECTION.

1019 To prevent inefficiency during plan generation, a process called duplicate detection ensures the plan-
 1020 ner doesn't repeatedly visit the same subgoal. Exploring logically equivalent but distinct branches
 1021 would lead to wasted computation and overly complex policies.

1022 A regressed formula is marked as a duplicate if it is structurally identical to one already seen, ignor-
 1023 ing differences in how variables are named or how the conjuncts are ordered.

1025 To enable this duplicate detection, every conjunctive formula is standardized by performing the
 1026 following steps:

- *Canonicalizing order*: Sorting all predicates and their arguments into a fixed, canonical sequence.
- *Consistent variable mapping*: Assigning variables consistently across the formula (e.g., uniformly replacing x, y, z with v_1, v_2, v_3 in corresponding positions).
- *Flattening*: Removing nested conjunctions where possible.

G ALFWORLD DETAILS

G.1 ALFWORLD TASKS

ALFWorld simulates a typical household environment and focuses on daily embodied AI tasks. Table 4 lists the supported task types along with their corresponding goal templates.

Table 4: Task-types and templated goal descriptions in ALFWorld.

Task-type	Templates
Pick & Place	(a) put a {obj} in {recep}. (b) put some {obj} on {recep}.
Examine in Light	(a) look at {obj} under the {lamp}. (b) examine the {obj} with the {lamp}.
Clean & Place	(a) put a clean {obj} in {recep}. (b) clean some {obj} and put it in {recep}.
Heat & Place	(a) put a hot {obj} in {recep}. (b) heat some {obj} and put it in {recep}.
Cool & Place	(a) put a cool {obj} in {recep}. (b) cool some {obj} and put it in {recep}.
Pick Two & Place	(a) put two {obj} in {recep}. (b) find two {obj} and put them {recep}.

- **Pick & Place** (e.g., “put a plate on the coffee table”) — the agent must find an object of the desired type, pick it up, find the correct location to place it, and put it down there.
- **Examine in Light** (e.g., “examine a book under the lamp”) — the agent must find an object of the desired type, locate and turn on a light source with the desired object in-hand.
- **Clean & Place** (e.g., “clean the knife and put in the drawer”) — the agent must find an object of the desired type, pick it up, go to a sink or a basin, wash the object by turning on the faucet, then find the correct location to place it, and put it down there.
- **Heat & Place** (e.g., “heat a mug and put on the coffee table”) — the agent must find an object of the desired type, pick it up, go to a microwave, heat the object by turning on the microwave, then find the correct location to place it, and put it down there.
- **Cool & Place** (e.g., “put a cool bottle on the countertop”) — the agent must find an object of the desired type, pick it up, go to a fridge, put the object inside the fridge and cool it, then find the correct location to place it, and put it down there.
- **Pick Two & Place** (e.g., “put two pencils in the drawer”) — the agent must find an object of the desired type, pick it up, find the correct location to place it, put it down there, then look for another object of the desired type, pick it up, return to the previous location, and put it down there with the other object.

1080
1081

G.2 EXAMPLE TRAJECTORY

1082

```

1083 You are in the middle of a room. Looking quickly around you, you see a
1084 drawer 15, a drawer 19,
1085 a drawer 7, an armchair 1, ... and a drawer 10.
1086 Your task is to: find two remotecontrol and put them in armchair.
1087 > go to sidetable 2
1088 You arrive at loc 34. On the sidetable 2, you see a remotecontrol 1.
1089 > take remotecontrol 1 from sidetable 2
1090 You pick up the remotecontrol 1 from the sidetable 2.
1091 > go to armchair 1
1092 You arrive at loc 1. On the armchair 1, you see nothing.
1093 > put remotecontrol 1 in/on armchair 1
1094 You put the remotecontrol 1 in/on the armchair 1.
1095 > go to sofa 1
1096 You arrive at loc 2. On the sofa 1, you see a newspaper 1, a pillow 1, and
1097 a remotecontrol 2.
1098 > take remotecontrol 2 from sofa 1
1099 You pick up the remotecontrol 2 from the sofa 1.
1100 > go to armchair 1
1101 You arrive at loc 1. On the armchair 1, you see a remotecontrol 1.
1102 > put remotecontrol 2 in/on armchair 1
1103 You won!
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
```

H ENTAILMENT DETAILS

H.1 ENTAILMENT DESIGN

In this section we outline how we misaligned the action model and the goal for both AFLWorld and Blocksworld domain.

AFLWorld Entailment Predicates Instead of using symbolic predicates such as, we define each action effect and precondition directly in NL and allow multiple entailed variants. Hard string matching is replaced with entailment sets as follows:

```

handEmpty : {"the agent's hand is empty", "the agent is not holding anything"}
in( $o, r$ ) : {" $o$  is in  $r$ ", " $o$  is inside  $r$ ", " $o$  is stored in  $r$ "}
hot( $o$ ) : {" $o$  is heated", " $o$  is hot", " $o$  is baked"}
washed( $o$ ) : {" $o$  is washed", " $o$  is clean", " $o$  is cleaned"}
cooled( $o$ ) : {" $o$  is chilled", " $o$  is cool", " $o$  is cooled"}
holding( $o$ ) : {"the agent is holding  $o$ ", "the agent possesses  $o$ ", "the agent has  $o$  in possessing"}
on( $r$ ) : {" $r$  is turned on", " $r$  is on", " $r$  is switched on"}
```

Blocksworld Entailment Predicates). We express each predicate directly in NL and allow multiple entailed variants rather than relying on canonical symbols. The entailment sets are:

```

handEmpty : {"the agent is not holding any objects", "the agent is not holding anything"}
clear( $b$ ) : {"block  $b$  has nothing on top of it", "no block is on  $b$ "}
onTable( $b$ ) : {"block  $b$  sits directly on the table", " $b$  is on the table"}
on( $b_1, b_2$ ) : {"block  $b_1$  is directly above block  $b_2$ ", " $b_1$  on  $b_2$ "}
holding( $b$ ) : {"the agent possesses block  $b$ ", "the agent is holding  $b$ "}
```

H.2 NATURAL LANGUAGE DESCRIPTION OF ACTION MODELS

We convert each action schema into a natural language description that specifies its preconditions and effects. These NL descriptions are provided to all LLM- and VLM-based models. To test

1134 alignment, we intentionally modify the descriptions to use phrases that entail the goal but do not
 1135 exactly match the goal string, thereby requiring explicit entailment reasoning. Below we provide
 1136 examples of the NL descriptions of both the ALFWorld and Blocksworld action models.
 1137

1138 **Natural Language ALFWorld Action Model :**

1140
 1141 The pickup action requires the agent’s hand to be empty and results in the
 1142 agent holding the target object. The put action requires the agent to
 1143 already be holding an object and allows the object to be placed inside or
 1144 on top of another receptacle, after which the agent’s hand becomes empty.
 1145 The heat action requires the agent to be holding an object and to be at a
 1146 heating device, producing the effect that the object becomes hot or baked.
 1147 The wash action requires the agent to be holding the object and to be at a
 1148 washing device, resulting in the object becoming washed. The chill action
 1149 requires the agent to be holding the object and to be at a chilling device,
 1150 resulting in the object becoming chilled. The light action requires the
 1151 agent to be holding the object and to be at a lighting device, producing
 1152 the effect that the object becomes illuminated. Finally, device state can
 1153 be toggled by turn on and turn off, which require the agent to be at the
 1154 device and result in it being switched on or off
 1155

1156 **Natural Language Blocksworld Action Model :**

1157
 1158 The pickup action requires the agent’s hand to be empty and results in the
 1159 agent holding the target object. The put action requires the agent to
 1160 already be holding an object and allows the object to be placed inside or
 1161 on top of another receptacle, after which the agent’s hand becomes empty.
 1162 The heat action requires the agent to be holding an object and to be at a
 1163 heating device, producing the effect that the object becomes hot or baked.
 1164 The wash action requires the agent to be holding the object and to be at a
 1165 washing device, resulting in the object becoming washed. The chill action
 1166 requires the agent to be holding the object and to be at a chilling device,
 1167 resulting in the object becoming chilled. The light action requires the
 1168 agent to be holding the object and to be at a lighting device, producing
 1169 the effect that the object becomes illuminated. Finally, device state can
 1170 be toggled by turn on and turn off, which require the agent to be at the
 1171 device and result in it being switched on or off
 1172

1173 **I DESIGN DETAILS OF NL-PDDL ALFWORLD AND BLOCKSWORLD
 1174 IMPLEMENTATION**

1175 **I.1 ALFWORLD**

1176 In addition to the core NL-PDDL planner, our ALFWorld agent consists of an VLM-based obser-
 1177 vation parser, a knowledge base (KB), and an LLM-based object grounder. The observation parser
 1178 extracts a list of object names from either text or image input. We use Gemini-2.0-Flash (Team
 1179 et al., 2023) for image observations parsing. After obtaining the object names, we use GPT-4o to
 1180 instantiate these into NL-predicates and update the agent’s KB. The KB is implemented using Py-
 1181 Datalog (Carbonnelle, 2024). Whenever new knowledge is acquired, the agent queries subgoals
 1182 generated by the NL-PDDL regression planner to determine a feasible plan to execute.

1183 **Observation Parser** The observation parser is responsible for extracting candidate object names
 1184 and their types from either textual descriptions or visual observations. In the ALFWorld Text setting,
 1185 object names are directly provided in the environment description, so we simply parse this list to
 1186 extract object references. In the ALFWorld Vision setting, the parser receives RGB frames from a
 1187 simulated first-person camera. Here, we use Gemini-2.0-Flash (Team et al., 2023), a VLM capable
 1188 of open-vocabulary object detection, to generate bounding boxes around objects of interest. Gemini
 1189 is prompted with relevant object descriptions (e.g., “Find a cup”, “Find something that can wash the

1188 cup") derived from subgoals generated by the NL-PDDL planner, and it returns a list of detected
 1189 objects with bounding boxes, labels, and confidence scores. The resulting object list serves as the
 1190 input to the grounding stage but does not by itself produce predicates or symbolic facts.
 1191
 1192
 1193

1194 **Object Grounder** The object grounder bridges the gap between raw object names and the struc-
 1195 tured predicates required by the NL-PDDL planner. Given the list of parsed object tokens and
 1196 their textual or visual context, we use GPT-4o to generate grounded natural language predicates
 1197 that align with our lifted planning formalism. For example, given the input "mug is cleaned," the
 1198 grounder infers the predicate `clean(mug_1)`. The grounder also infers action-relevant relations
 1199 such as `canClean(sink_1, apple_1)` when provided with relevant context. This ensures that
 1200 grounded facts are consistent with both the agent's environment and the symbolic schema expected
 1201 by the planner.
 1202
 1203

1204 **Knowledge Base** The knowledge base (KB) stores grounded predicates and supports logical rea-
 1205 soning over them. We implement the KB using PyDatalog (Carbonnelle, 2024), a lightweight logic
 1206 programming library for Python that supports declarative predicate logic, variable binding, and rule-
 1207 based inference. During plan execution, the planner queries the KB to check whether a lifted subgoal
 1208 can be satisfied with the current facts. If the KB entails a subgoal, we bind the variables in both the
 1209 subgoal and the associated actions, and then execute the corresponding plan. This allows the planner
 1210 to maintain consistency between abstract subgoals and grounded execution.
 1211
 1212

1213 I.2 BLOCKSWORLD

1214
 1215 For the Blocksworld domain, we convert the original PDDL goal into a corresponding NL-PDDL
 1216 formulation. We follow the methodology of Valmeekam et al. (2023a) to ensure consistency be-
 1217 tween symbolic and natural language representations. The lifted regression planner then produces
 1218 a sequence of subgoals with a fixed horizon of 10, which are grounded and validated sequentially.
 1219 Each subgoal and the agent's initial state are translated into PyDatalog and checked for logical con-
 1220 sistency with the original PDDL specification. Once all subgoals are validated, we use the standard
 1221 PDDL validator from Muise et al. (2022) to ensure correctness of the complete plan.
 1222
 1223

1224 I.3 LLM PROMPT FOR TYPE CONSISTENCY CHECK

```

1228     ROLE:  

1229        You are a helper agent in a common household setting.  

1230        You are checking TYPE ENTAILMENT between two predicates' term types.  

1231     INSTRUCTION:  

1232        1. If the candidate's term type set implies or matches the target's term  

1233           type set for each corresponding term, answer YES.  

1234        2. Consider synonyms and common-sense subtype relations (e.g.,  

1235           'vegetable' entails 'food').  

1236        3. If information is unknown, be conservative and answer NO unless it's  

1237           very likely.  

1238     INPUT:  

1239        - Target term types: ${predicates_in_action_model_type$}  

1240        - Candidate term types: ${misaligned_goal_type$}  

1241     OUTPUT FORMAT:  

1242        - Line 1: exactly YES or NO.  

1243        - Line 2: Reason.  

1244     RESPONSE:

```

```

1242
1243     ROLE:
1244         You are a helper agent in a common household setting.
1245     QUESTION:
1246         - if you know Predicate 2 "{$action_model_predicate$}" is true, can you
1247             imply Predicate 1 "{$misaligned_goal_predicate$}" is true?.
1248         - Respond with exactly "YES" if you think the statement is generally
1249             implied
1250         - Respond with "NO" if you think the statement is generally false
1251     INPUT:
1252         - Predicate 1: "{$misaligned_goal_predicate$}"
1253         - Predicate 2: "{$action_model_predicate$}"
1254     INSTRUCTION:
1255         1. Use the definition of the predicates to determine if Predicate 2
1256             implies Predicate 1.
1257         2. You know the following background to determine the specific
1258             information of the objects within Predicate 1 and Predicate 2: $goal
1259             predicates$.
1260         3. When determining the response, consider the meaning of the Predicate
1261             1 and Predicate 2 with the type of the specific object each refers to in
1262             common contexts.
1263         4. Be creative and think outside the box. If there is just a typo
1264             between the two predicates, you should say Yes.
1265     OUTPUT FORMAT:
1266         - Line 1: exactly YES or NO.
1267         - Line 2: Reason.
1268     RESPONSE:
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

```

J WORKS ON TRANSLATING NL TO PDDL

PDDL has long been established as the primary standard for defining planning domains and problems in AI Gerevini (2020); Haslum et al. (2019). However, authoring accurate PDDL domains and problem specifications is a resource-intensive process that requires human expertise. To facilitate this process, a growing body of work has explored translating NL descriptions of planning domains and problems into PDDL.

Early efforts include NLtoPDDL Miglani & Yorke-Smith (2020), a pipeline that leverages readily available NL data by combining pre-trained contextual embeddings with Deep Reinforcement Learning (DRL) techniques previously used to extract structured plans from NL. Another example is FPTCP Huo et al. (2020), which constructs a ternary template of NL sentences to extract actions and their associated objects in a human-in-the-loop framework.

With the emergence of large language models (LLMs), their strong language understanding and commonsense reasoning abilities have been used to further facilitate PDDL construction from NL. For instance, Xie et al. (2023) uses the in-context learning capabilities of LLMs to translate NL domains and problems into PDDL via few-shot examples, demonstrating that using LLMs solely for translation and delegating the planning step to a PDDL solver yields superior results compared to direct LLM-based planning. Similarly, Smirnov et al. (2024) employs an LLM to generate PDDL plans, but also introduces LLM-based consistency checks and error-correction loops to improve plan quality. More recently, advanced approaches such as Ada ? have been proposed. Rather than translating a pre-existing instruction or domain manual into a single PDDL file, Ada interactively learns adaptive planning representations. In this framework, high-level action abstractions and low-level controllers are jointly adapted to a given domain of planning tasks, guided by NL inputs.

Although these seminal works significantly improve the usability of PDDL, they are all limited to translating NL inputs into PDDL and then relying on an existing PDDL planner—a process that can be error-prone due to the inherent limitations of LLMs and the auxiliary techniques involved. In contrast, NL-PDDL directly extends the PDDL framework to natively support NL specifications, thereby eliminating the need for NL-to-PDDL translation. Moreover, prior approaches primarily focus on generating and refining closed-world domains, whereas NL-PDDL supports open-world planning with incomplete domain knowledge.

1296 **K NUMBER OF UNDEFINED PREDICATES IN ALFWORLD**
12971298 In practical scenarios, misalignment between an agent’s action model and the underlying task goals
1299 is common. Consequently, it is essential for a symbolic planner to incorporate LLM-based reasoning
1300 when necessary—a capability that NL-PDDL is explicitly designed to provide.1301 As shown in Table 5, In the ALFWorld domain, a total of 583 unique predicates appear across all
1302 tasks, with an average of 52.26 predicates per task. Yet the action model defines only 22 predicates—
1303 approximately 3.7% of all predicates observed in our experiments. The remaining predicates
1304 arise dynamically from the user’s instructions and the agent’s interaction with the environment.1305 Because these predicates are not predefined, existing symbolic planning approaches that require
1306 a fixed, closed-world predicate set become infeasible. In contrast, NL-PDDL is the only formal
1307 framework that removes the need for a predefined predicate vocabulary by leveraging LLM-based
1308 entailment reasoning to interpret and ground predicates on demand.1309
1310 Table 5: Predicate statistics for the ALFWorld domain.
1311

1312 Statistic	1313 Value
1314 Total Number of Predicates	1315 583
1315 Predefined Predicates	1316 22
1316 Predicates per Task	52.26

1317 **L ENTILAMENT CHECK**
13181319 We conducted an additional ablation study (cf. Table 6) to evaluate the correctness of the LLM when
1320 performing predicate-level affordance reasoning over natural-language predicates. We randomly
1321 selected 20 predicates from ALFWorld and asked a human annotator to manually provide three
1322 sentences that should be entailed by each predicate and three sentences that should not be entailed.
1323 This yielded 120 entailment–non-entailment pairs in total. We then evaluated all pairs using the
1324 NL-PDDL entailment procedure to assess the model’s precision and robustness in predicate-level
1325 reasoning.1326
1327 Table 6: Classification statistics for the ALFWorld predicate evaluation.
1328

1329 Metric	1330 Number of Instances
1331 True Positive	60
1332 True Negative	58
1333 False Positive	0
1334 False Negative	2
1334 Total Accuracy	98.3

1336 **M LLM CALLS AND LATENCY**
13371338 We randomly select 20 goals from *Misaligned BlockWorld* and *Misaligned ALFWorld* for each depth
1339 and report the mean and 95% confidence interval of the total number of LLM calls in Table 7. These
1340 totals include both local cache hits (i.e., we store all previously encountered LLM entailments in a
1341 global cache shared across goals) and actual API calls.1342
1343 Table 7: LLM call and entailment statistics across domains.
1344

1345 Domain	1346 ALFWORLD with Misalignment	1347 Misalignment BlocksWorld
1346 LLM API Call	50.4 ± 12.9	73.2 ± 23.4
1347 Total Entailments Reasoning	221.5 ± 155.6	1263 ± 435.7
1348 Avg Latency per Call	1.8648 sec	1.2443 sec