SKILLWRAPPER: GENERATIVE PREDICATE INVENTION FOR SKILL ABSTRACTION

Anonymous authorsPaper under double-blind review

ABSTRACT

Generalizing from individual skill executions to solving long-horizon tasks remains a core challenge in building autonomous agents. A promising direction is learning high-level, symbolic representations of the low-level skills of the agents, enabling reasoning and planning independent of the low-level state space. Recent advances in foundation models have made it possible to generate symbolic predicates that operate on raw sensory inputs—a process we call generative predicate invention—to facilitate downstream abstraction learning. However, it remains unclear which formal properties the learned representations must satisfy, and how they can be learned to guarantee these properties. In this paper, we address both questions by presenting a formal theory of generative predicate invention for skill abstraction, resulting in symbolic operators that can be used for provably sound and complete planning. Within this framework, we propose SKILLWRAPPER, a method that leverages foundation models to actively collect robot data and learn human-interpretable, plannable representations of black-box skills, using only RGB image observations. Our extensive empirical evaluation in simulation and on real robots shows that SKILLWRAPPER learns abstract representations that enable solving unseen, long-horizon tasks in the real world with black-box skills.

1 Introduction

An autonomous agent operating in the real world must process low-level sensory and motor signals while reasoning about high-level objectives (Doncieux et al., 2018; Konidaris, 2019). Analogous to how humans can perform complex tasks, like cooking or cleaning, without reasoning about muscle-level control, agents should have internal models of their skills that abstract away nuanced activities on the lower level. Such models must capture the necessary conditions for a skill to be executed (e.g., "pouring a teapot requires holding it first") and the consequences of doing so (e.g., "pouring a teapot leaves it empty"). These two properties, known as preconditions and effects in the AI planning literature, enable compositional reasoning to identify long-horizon plans that can sequence lower-level skills to solve a task. Typically, these models must be specified manually. However, in real-world settings such skill representations may be nontrivial to acquire or engineer due to complex inter-skill constraints specific to the agent's embodiment. This calls for algorithms that learn symbolic transition models of black-box skills without hand specification, enabling agents to directly utilize those skills to solve long-horizon tasks with off-the-shelf AI planners.

Traditional approaches to skill abstraction often require factorizing the low-level state space to learn classifiers for each symbolic representation, relying heavily on hand-collected transition data (Konidaris et al., 2018). Recently, foundation models have enabled a new paradigm: generating semantically meaningful predicates directly from raw observations and directly evaluating their truth values on low-level observations (e.g., RGB images)—a process we refer to as *generative predicate invention*. Recent work has explored how foundation models can be used for predicate invention, by generating Python code to implement predicates (Liang et al., 2025) or sampling large predicate pools followed by sub-selection. However, these methods produce ad-hoc planning representations that cannot be guaranteed to solve a given task, and leave core questions on predicate invention unanswered: *what* properties should these learned abstractions satisfy, and *how* can they be learned to achieve these properties?

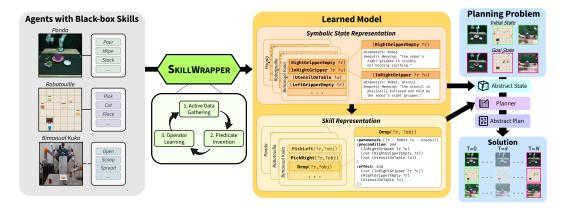


Figure 1: Brief overview of SKILLWRAPPER: given an agent equipped with named black-box skills, SKILLWRAPPER learns a plannable model or representation of skills (as PDDL operators) compatible with automated planners. These operators are comprised of predicates generated by the foundation model. With a novel planning problem described using images of the initial state and goal state, our system uses a foundation model to generate an abstract state representation, where the truth value of each invented predicate is determined by the model. It is important to note that SKILLWRAPPER is *agnostic* to the agent: we illustrate both real-world (robot) and simulated agents in this figure.

Our answer is twofold. First, we develop a formal theory of generative predicate invention for skill abstraction, precisely characterizing the conditions under which a learned skill abstraction will be provably *sound* and *complete* with respect to downstream planning. Building on this foundation, we introduce SKILLWRAPPER, a method explicitly designed to guarantee these theoretical criteria. SKILLWRAPPER uses foundation models in three ways: interactively collecting data in the environment, proposing predicates when the current model fails, and classifying predicate truth values based solely on RGB image observations. Using these data and predicates, SKILLWRAPPER learns symbolic representations of black-box skills that are both human-interpretable and directly usable for AI planning. Specifically, our approach is given a skill type signature as input and learns a symbolic robot skill model in PDDL (McDermott et al., 1998).

We highlight the following contributions: (1) A formal theory of generative predicate invention for *provably sound and complete* skill abstraction; (2) SKILLWRAPPER, a principled system built on this framework that leverages foundation models to learn interpretable symbolic representations of black-box skills; and (3) an extensive empirical evaluation of the system, demonstrating effectiveness in simulation and on two real robots.

2 Problem Setting

In this section, we briefly discuss our problem setting while defining it formally in Appendix A. We consider an agent equipped with a finite set of *object-centric skills* Ω , modeled as black-box options (Sutton et al., 1999). The agent can execute any $\omega \in \Omega$ and determine whether it succeeds, but it does not possess a symbolic transition model of these skills. Without such a model, the agent cannot plan over a long horizon without reasoning at the low-level state space \mathcal{S} , which is continuous, high-dimensional, and impractical for classical search. The goal of *Skill Model Learning* is to acquire a symbolic abstraction of skills that enables efficient composition via classical planning.

Environment. An *environment* is a tuple $(S, \mathcal{T}, \Omega, T)$, where S is the continuous state space, \mathcal{T} is a finite set of object types, and Ω is the skill library. Each skill $\omega \in \Omega$ is parameterized by object types drawn from the set \mathcal{T} . The environment dynamics are governed by an unknown transition function $T: S \times \Omega \to S$. A *setting* is defined as (s_0, \mathcal{O}) , consisting of an initial state $s_0 \in S$ and a set of typed objects \mathcal{O} with $\tau(o) \subseteq \mathcal{T}$ for each $o \in \mathcal{O}$.

Black-box skills. A skill $\omega \in \Omega$ is a tuple $(\mathcal{I}_{\omega}, \pi_{\omega}, \beta_{\omega}, \theta_{\omega})$, where $\mathcal{I}_{\omega} \subseteq \mathcal{S}$ is the *initiation set*, π_{ω} is the option policy, $\beta_{\omega} \subseteq \mathcal{S}$ is the *termination set*, and $\theta_{\omega} = (\tau_{\omega}^1, \dots, \tau_{\omega}^k)$ are the type constraints

on its k parameters. A skill instance $\underline{\omega} = \omega(o_1, \dots, o_k)$ is valid for $o_i \in \mathcal{O}$ if $\tau_{\omega}^i \subseteq \tau(o_i)$ for all i. Executing ω from $s \in \mathcal{I}_{\omega}$ terminates in some $s' \in \beta_{\omega}$, affecting only the state of the bound objects.

Symbolic predicates and operators. To enable abstract reasoning, we introduce a finite set of predicates \mathcal{P} . Each $\sigma \in \mathcal{P}$ is parameterized by \mathcal{T} and has a classifier $\phi_{\sigma}: \mathcal{O}^n \times \mathcal{S} \longrightarrow \{0,1\}$ that evaluates its truth value for arguments in a state. Grounding predicates with valid objects induces a set of grounded predicates $\bar{\mathcal{P}}$, and hence an abstract state space $\bar{\mathcal{S}}=2^{\bar{\mathcal{P}}}$. The abstraction function $\Gamma: \mathcal{S} \to \bar{\mathcal{S}}$ maps each state to the set of grounded predicates true in that state, while the *grounding* function $\mathcal{G}: \mathcal{S} \to \mathcal{S}$ maps an abstract state to the states in which its grounded predicates hold.

An operator $a \in \mathcal{A}$ is defined as $(\omega_a, \theta_a, PRE_a, EFF_a)$, where $\omega_a \in \Omega$ is the associated skill, θ_a are typed operator parameters, $PRE_a \subseteq \mathcal{P}$ is the precondition set, and $EFF_a = (EFF_a^+, EFF_a^-)$ are the add and delete effects. Grounding all operators $a \in A$ with objects $o \in \mathcal{O}$ yields a set of abstract actions $\bar{a} \in \mathcal{A}$, each executable whenever its ground preconditions hold in the current abstract state.

Planning problems. Given an environment (S, T, Ω, T) , the agent's objective is to learn an abstract transition model $\mathcal{M}=(\mathcal{P},\mathcal{A})$, such that \mathcal{M} is sufficient for planning with Ω . Formally, a skill planning problem is $(s_0, \mathcal{O}, \mathcal{S}_q)$, where (s_0, \mathcal{O}) is a setting and $\mathcal{S}_q \subseteq \mathcal{S}$ are goal states. The corresponding abstract planning problem is $(\Gamma(s_0), \bar{\mathcal{A}}, \bar{\mathcal{S}}_a)$, where $\bar{\mathcal{S}}_a = \{\bar{s} \in \bar{\mathcal{S}} \mid \mathcal{G}(\bar{s}) \cap \mathcal{S}_a \neq \emptyset\}$. An abstract plan $\bar{\pi} = \langle \bar{a}_1, \dots, \bar{a}_n \rangle$ is valid if its execution under \mathcal{M} yields an abstract trajectory consistent with some feasible low-level trajectory under T, ending in a goal state.

Problem statement. The Skill Model Learning problem is: given experience in the form of state-skill-instance-next-state tuples $\langle s, \omega, s' \rangle$, learn a model $\mathcal{M} = (\mathcal{P}, \mathcal{A})$ such that every plan found by a complete symbolic planner over \mathcal{M} corresponds to a feasible low-level skill plan.

3 SKILLWRAPPER

108

109

110 111

112

113

114

115 116

117

118

119

120

121 122

123

124

125

126

127 128

129

130

131 132

133 134

135

136

137

138

139

141

142

143

144

145

146 147

158 159

161

In this section, we introduce SKILLWRAPPER, a novel approach that autonomously learns symbolic representations for black-box skills using the concepts defined in Appendix C. To produce a valid abstract model for skill planning, SKILLWRAPPER iterates through a three-step process: 1) actively proposing and executing exploratory sequences of skills to collect data on the initiation and termination set of each skill, 2) incrementally building a set of predicates from scratch by contrasting positive and negative examples, and then 3) constructing valid operators using these invented predicates, from which further exploratory skill sequences can be proposed. This procedure is outlined in Algorithm 1. As SKILLWRAPPER continues to collect data, add predicates, and update its planning model, it learns a progressively more accurate abstract transition model. The resulting skill representations, or *operators*, can be used with an off-the-shelf classical planner to solve task planning problems specified using natural language or images. We delve into the core components of our system (lines 4-9 of Algorithm 1) in the following subsections. Lastly, we provide strong theoretical results for the soundness and completeness of SKILLWRAPPER in Section 3.4.

Algorithm 1 SKILLWRAPPER

```
148
                   1: Input: Set of skills \Omega, number of iterations m \in \mathbb{N}_1
149
                   2: Output: Abstract transition model \mathcal{M} = (\mathcal{P}, \mathcal{A})
150
                   3: \mathcal{D}, \mathcal{P}, \mathcal{A} \leftarrow \emptyset
151
                   4: for i \in \{1, \dots, m\} do
152
                                \langle \underline{\omega}_1, \dots, \underline{\omega}_k \rangle \leftarrow \mathsf{ProposeSkillSequence}(\Omega, \mathcal{D}, \mathcal{P})
153
                                \mathcal{D} \leftarrow \mathcal{D} \cup \text{EXECUTESKILLS}(\underline{\omega}_1, \dots, \underline{\omega}_k)
                   7:
                                \mathcal{P} \leftarrow \text{InventPredicates}(\Omega, \mathcal{D}, \mathcal{P}, \mathcal{A})
154
                                \mathcal{A} = LearnOperators(\hat{\mathcal{D}}, \hat{\mathcal{P}})
                   9: end for
156
                 10: return \mathcal{M}
157
```

3.1 ACTIVE DATA COLLECTION

To collect data for learning, our method commands the agent to execute its skills in the world and then collects the resulting transitions. Each command is a sequence of skill instances $\langle \underline{\omega}_1, \dots, \underline{\omega}_k \rangle$. While executing these commands, the agent collects a dataset \mathcal{D} of transitions of the form $\langle s, \underline{\omega}, s' \rangle$, where $s, s' \in \mathcal{S}$. These transitions can potentially answer two questions:

- 1. Executability: Can the skill instance $\underline{\omega}$ be executed from state s?
- 2. Skill Dynamics: If $s \neq s'$, what has changed in the environment due to executing $\underline{\omega}$?

We guide the exploration of skill preconditions and effects using a foundation model, which proposes skill sequences in natural language. Rather than naively sampling these sequences from the foundation model's token distribution, we apply two additional heuristics to bias the system toward promising sequences (see Appendix B.1 for details). SKILLWRAPPER uses failure transitions to improve the learned abstract model by eliminating unnecessary preconditions.

3.2 PREDICATE INVENTION

162

163

164

165 166

167

168

169

170

171

172 173

174 175

176

177 178

179

180

181

182

183

185

187

188

189

190

191 192 193

194

195

196

197

199

200

201

202

203204

205

206

207

208

209

210

211212213

214

215

We now present our predicate invention algorithm, which, unlike prior work (Silver et al., 2023; Liang et al., 2025), does not require an initial set of predicates to bootstrap the invention process.

Algorithm 2 Invent Predicates

```
1: Input: Skill set \Omega, skill execution traces \mathcal{D}(\omega) = \{\langle s, s' \rangle\}_{\omega}, predicate set \mathcal{P}, operator set \mathcal{A}, and and
        abstraction function \Gamma.
       Output: Predicate set \mathcal{P}
 3: for \omega \in \Omega do
                while \exists \langle s_i, \underline{\omega}, s_i' \rangle, \langle s_j, \underline{\omega}, s_j' \rangle \in \mathcal{D} s.t. \Gamma(s_i) = \Gamma(s_j), s_i \in \mathcal{I}_{\underline{\omega}}, s_j \notin \mathcal{I}_{\underline{\omega}} do
  5:
                        P \leftarrow new\_predicate
                       \mathcal{P} \leftarrow \mathcal{P} \cup \mathbf{P} if ScorePrecond(\mathcal{A}, \omega, \mathcal{D})
 6:
                end while
 7:
                while \exists \langle s_i, \underline{\omega}, s_i' \rangle, \langle s_j, \underline{\omega}, s_j' \rangle \in \mathcal{D} s.t. \bar{\Delta}(s_i, s_i') = \bar{\Delta}(s_j, s_j'), s_i \in \mathcal{I}_{\omega}, s_j \notin \mathcal{I}_{\omega} do
  8:
 9:
                       P \leftarrow new\_predicate
                        \mathcal{P} \leftarrow \mathcal{P} \cup \mathbf{P} \text{ if } \mathbf{SCOREEFF}(\mathcal{A}, \omega, \mathcal{D})
10:
                end while
11:
12: end for
13: return \mathcal{P}
```

Conditions for predicate invention. SKILLWRAPPER identifies two conditions under which the current predicate vocabulary is insufficient, based on the desired properties of soundness and completeness (discussed fully in Appendix C). In these cases, the system must invent new predicates to resolve incongruities between the observed data and the current abstract transition model.

The first condition arises when SKILLWRAPPER detects that the symbolic vocabulary cannot express a necessary precondition for a skill. Concretely, this occurs when two transitions of a skill, one successful and one failed, begin from the same abstract state under the current predicates. Because the vocabulary cannot distinguish between these initial states, an additional predicate is required. Formally, this condition is expressed as:

$$\exists \langle s_i, \underline{\omega}, s_i' \rangle, \langle s_j, \underline{\omega}, s_j' \rangle \in \mathcal{D} \text{ s.t. } \Gamma(s_i) = \Gamma(s_j), \text{ but } s_i \in \mathcal{I}_{\underline{\omega}} \text{ while } s_j \notin \mathcal{I}_{\underline{\omega}}.$$

The second condition used by SKILLWRAPPER to trigger predicate invention is based on inconsistencies in observed skill effects. Specifically, this occurs when two transitions of the same skill produce identical abstract effects, despite one succeeding and the other failing. In our deterministic setting, this reduces to a successful skill execution producing no effects, though the condition naturally extends to stochastic settings with mid-execution failures. Formally, letting $\bar{\Delta}(s,s') = (\Gamma(s) \setminus \Gamma(s'), \Gamma(s') \setminus \Gamma(s))$, we express the condition as:

$$\exists \langle s_i,\underline{\omega},s_i'\rangle, \langle s_j,\underline{\omega},s_j'\rangle \in \mathcal{D} \text{ s.t. } \bar{\Delta}(s_i,s_i') = \bar{\Delta}(s_j,s_j'), \text{ yet } s_i \in \mathcal{I}_{\underline{\omega}} \text{ and } s_j \notin \mathcal{I}_{\underline{\omega}}.$$

Empirical predicate selection. Although foundation models provide a strong prior on which information is skill-relevant for predicate construction, they may still produce errors and hallucinations. To ensure robustness, we introduce a scoring function that estimates the usefulness of a candidate predicate by adding it to the vocabulary and learning hypothetical operators. For each

successful transition in \mathcal{D} , there must exist at least one operator with preconditions satisfied by the transition's initial state; for each failed transition, no such operator may exist. Effect evaluation follows the same principle. After evaluating all transitions, we decide whether to add the candidate predicate based on an empirical threshold (Details of Algorithm 6 can be found in the appendix.)

3.3 OPERATOR LEARNING

Our operator learning procedure extends the *associative model learning* paradigm (Arora et al., 2018) to the setting of skill abstractions.

Associative model learning. A single skill may induce multiple distinct abstract state changes depending on the context of execution. To represent these *conditional effects*, SKILLWRAPPERclusters observed transitions based on their lifted effect sets, enabling it to learn a single operator across distinct instantiations of a skill. The preconditions of each operator are then computed as the intersection of all initial abstract states in the corresponding transitions, ensuring that each operator is both minimal and consistent.

Multi-type object-centrism. In realistic domains, objects do not fit neatly into single type categories, but rather belong to multiple overlapping categories (e.g., a Cup is fillable, while a Bottle is both fillable and openable). This complicates the process of generalizing grounded transitions into lifted operators, because it may be ambiguous which object attribute enables successful execution. We adopt a conservative strategy: SKILLWRAPPER assigns preconditions using the lowest level of the type hierarchy consistent with the data, preventing overgeneralization while retaining compositional structure.

Predicate re-evaluation. Predicates are generated sequentially, and early inventions may bias later stages if left unchecked. To mitigate this, SKILLWRAPPER re-applies the scoring function to the entire predicate set after each iteration of data collection. This allows spurious or redundant predicates to be discarded as more data is collected. In addition, tautological predicates—those that are always true or always false—are automatically filtered. As a result, the learned predicate set remains compact, informative, and aligned with the most recent transition data.

3.4 THEORETICAL ANALYSIS

We now provide theoretical guarantees for SKILLWRAPPER. These results show that the learned symbolic model is sound with respect to observed data and converges to a complete model with high probability. Full proofs are deferred to Appendix D.

Theorem 1 (Soundness of SKILLWRAPPER). Every operator $a \in \mathcal{A}_n$ in the model \mathcal{M}_n learned by SKILLWRAPPER is supported by at least one observed transition $\langle s, \omega, s' \rangle \in B_n$. That is, $\underline{s} \models \mathsf{PRE}_a$ and $\underline{s}' \models \mathsf{EFF}_a$.

Sketch. SKILLWRAPPER constructs each operator directly from sampled transitions; thus no unsupported operator can appear.

Lemma 1. For each $\omega \in \Omega$, the initiation set I_{ω} and termination set β_{ω} inferred by SKILLWRAPPER match exactly the corresponding predicate sets α_{ω} and ζ_{ω} derived from B_n , i.e. $I_{\omega} \triangle \alpha_{\omega} = \emptyset$ and $\beta_{\omega} \triangle \zeta_{\omega} = \emptyset$.

Sketch. Any mismatch would contradict the termination condition of the wrapper; thus initiation and termination sets are consistent with observed data.

Theorem 2 (Probabilistic-completeness of SKILLWRAPPER). Let \mathcal{M}^* be the true complete model. With probability at least $1 - |\mathcal{H}| \exp(-n\epsilon)$, the model $\widehat{\mathcal{M}}_n$ learned from n samples satisfies $d_{\text{compl}}(\widehat{\mathcal{M}}_n, \mathcal{M}^*) \leq \epsilon$, i.e. it misses fewer than an ϵ -fraction of feasible transitions.

Sketch. Since $\widehat{\operatorname{Err}}(\widehat{\mathcal{M}}_n) = 0$ by construction, a Chernoff bound and union bound over the finite hypothesis class imply that the true error is small with high probability.

Together, these results establish that SKILLWRAPPER learns symbolic operators that are sound with respect to observed transitions, consistent in their preconditions and effects, and probabilistically

complete relative to the true underlying model. These properties justify SKILLWRAPPER as a reliable model-learning procedure for task and motion planning. Next, we discuss empirical evaluation for SKILLWRAPPER.

4 EXPERIMENTS

For all experiments in this section, we consider images as fully observable state representations, assuming that an abstract state can be inferred from an image without uncertainty. Both the initial and goal states of each problem are specified using RGB images. These images may come from diverse sources, including a top-down view of an animated game, a third-person camera observing a robot, or the robot's own egocentric perspective. All quantitative results reported in this section are averaged over three independent runs.

4.1 IMPLEMENTATION OF PREDICATES

We employ foundation models for both predicate generation and evaluation:

- A foundation model gives us a string that can be used as a lifted predicate (generates interpretable relational classifiers with a good heuristic).
- After grounding with valid parameters, the predicate can be prompted to the foundation model again to acquire the truth value. In other words, a foundation model can be used as a relational classifier.

With these two properties, we can use the VLM's response in string form as a relational predicate, and a grounded version of the predicate can be used as a classifier. We use GPT-5 (OpenAI, 2025) for predicate generation and evaluation.

4.2 SIMULATION

We first conduct experiments in Robotouille (Gonzalez-Pumariega et al., 2025), which is a simulated grid world kitchen domain with an agent that has five high-level skills: *Pick*, *Place*, *Cut*, *Cook*, and *Stack*. In the environment, there are several objects: a patty, lettuce, a top bun, and a bottom bun; there is also a cutting board and a stove for cutting the lettuce and cooking the patty, respectively. We design and categorize 50 abstract planning problems: 20 easy problems, whose optimal solutions have no more than 7 steps; 20 hard problems, whose optimal solutions have no more than 15 steps; and 10 impossible tasks that cannot be realized in the environment.

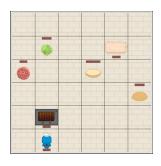


Figure 2: Screenshot of the Robotouille environment.

For a complete evaluation of SKILLWRAPPER, we compare SKILL-WRAPPER against four baselines:

- Expert Operators: A human expert who is familiar with PDDL is asked to interact with the environment and manually write predicates with semantics and PDDL operators.
- **System Predicates**: This baseline directly uses the built-in predicates defined by the simulator. In addition, instead of getting the truth values by classifying with foundation models, this baseline has access to the exact abstract state of the simulator. However, predicate invention is not allowed in this environment.
- ViLa (Hu et al., 2023): This baseline is a closed-loop VLM-based approach that iteratively prompts a foundation model for the next action until the goal state is reached, given an image observation and the agent's action history.
- Random Exploration: Instead of proposing skill sequences, this baseline randomly samples a skill and populates the arguments with valid objects. This baseline shares the same predicate invention and operator learning algorithms as SKILLWRAPPER.

For each baseline that performs operator learning, we run the learning algorithm for five iterations, with each iteration proposing and executing one skill sequence consisting of 15 steps. We then evaluate each method on the evaluation set and report the average results in Table 1, where Solved~%









(a) Franka: Initial

(b) Franka: Goal

(c) Bimanual Kuka: Initial

(d) Bimanual Kuka: Goal

Figure 3: Initial and Goal States for Real Robot Experiments.

is the percentage of the problem set that was successfully solved or where impossible tasks were correctly identified by returning an empty plan, and *PB* stands for planning budgets—the number of plans that were tried before solving the problem (adopted from Liang et al. (2025)). We set a planning budget cap of 10 across all problems; if the planning budget has been used for a problem, it is considered a failure. Theoretically, PB is an adequate metric that reflects the completeness of the learned model, and the *impossible* problems reflect the soundness of it. As shown in the table, SKILLWRAPPER significantly outperforms all baselines that have no access to privileged knowledge, and even surpasses the performance of the system predicates baseline (Sys Preds.) on hard problems. Here we present the key insights, while leaving case studies, example operators, and failure modes and analysis of SKILLWRAPPER in Appendix E.

Table 1: Baseline Comparison in Robotouille Environment

Method	Easy Solved % ↑	PB↓	Hard Solved % ↑	PB↓	Impossible
Expert Ops. Sys Preds.	80.0 80.0	2.0 2.6	56.7 16.7	4.8 8.5	100 43.3
ViLa Random Exp. Ours	33.3 5.0 73.3	9.6 2.9	1.7 0 38.3	10.0 6.1	16.7 100 100

Compared to the expert-constructed operators, which serve as an approximate upper bound for performance without privileged simulator access, SKILLWRAPPER demonstrates competitive accuracy while requiring only a small number of exploratory interactions. In particular, SKILLWRAPPER exhibits strong generalization from the easy to the hard set, indicating that the invented predicates and learned operators capture meaningful abstractions rather than overfitting to specific training traces. The gap between SKILLWRAPPER and ViLa highlights the benefit of explicitly learning a symbolic model instead of relying solely on open-loop prompting, while the poor performance of random exploration underscores the importance of guided data collection in learning effective operators. These findings suggest that SKILLWRAPPER achieves a favorable trade-off between data efficiency and model soundness and completeness, which is crucial for scaling to larger domains.

4.3 REAL ROBOTS

To demonstrate the applicability of SKILLWRAPPER for real-world agentic settings, we designed two sets of experiments with two robotic platforms: a Franka Emika Panda robot (Figures 3a and 3b) and a bimanual platform with two Kuka iiwa robots (Figures 3c and 3d). For both robot experiments, we assume that all skills are deterministic. The supplementary material contains video demonstrations of both experiments.

Generalization of SKILLWRAPPER In this setting, a tabletop Franka Emika Research 3 (Panda) manipulator has its skill set Ω consisting of five black-box skills: *Pick, Place, Stack, Pour*, and *Wipe*. The object set \mathcal{O} contains five objects: a mug, a teapot, a plate, a bowl, and a sponge. The robotic agent can pick and place all objects except the mug and plate, pour ingredients from the teapot into the mug, and use the sponge to wipe the plate if it is dirty. To validate the generalization ability of the learned skill representations, we design three smaller training environments, such that each environment only contains a subset of objects \mathcal{O} , and thus only a subset of skills Ω are executable. Each

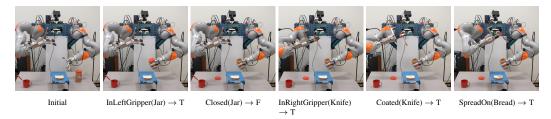


Figure 4: Sequence of Bimanual Robot Skill Rxecution with Predicate Value Changes

of the training environments contains fewer than 10 possible states. After running SKILLWRAPPER for exactly one iteration for each environment, we port the learned skill representations to the test environment that contains all objects in \mathcal{O} , which induces a state space of 34 possible abstract states. To quantify this generalization process, we similarly prepare an evaluation problem set that consists of five problems in the training environments, five problems in the test, and five *Impossible* problems across the two environments. The results are shown in Table 2.

Table 2: Results of Generalization Experiment

Method	In-domain Solved % ↑ PB ↓		Generalization Solved % ↑ PB ↓		Impossible
Expert Ops.	66.7	3.3	53.3	5.3	46.7
ViLa	46.7	-	6.7	-	6.7
Ours	76.7	2.7	60.0	4.0	66.7

Investigating Learning Curve of SKILLWRAPPER In this bimanual manipulation setting, a robot with two Kuka iiwa arms is equipped with a skill set Ω containing six black-box skills: LeftArmPick, RightArmPick, OpenJar, Scoop, Spread, and Drop. The object set \mathcal{O} consists of three objects: a peanut butter jar, a knife, and a slice of bread. The robotic agent can pick up the knife and jar, drop the knife, open the jar, scoop peanut butter with the knife, and spread it on the bread. Notably, this environment contains multiple dead ends, which would hinder the data gathering process. For example, the knife cannot be picked up again once dropped, the jar cannot be released once picked up, and the bread and knife cannot be cleaned once in contact with peanut butter. Moreover, the skills are heavily interdependent. We make this design to investigate the learning process of SKILLWRAPPER over iterations. Again, we compare the performance against two baseline methods, ViLa and Expert Operators. An example of predictive truth value changes induced by a sequence is shown as in Firgure 4. We observe that performance improves as SKILLWRAPPER progressively obtains more transition data and invents more predicates, finally surpassing the baselines. The performance improvement over iterations is shown in Figrue 5.

These results demonstrate that SKILLWRAPPER is effective in real robot settings: it generalizes skill representations learned in restricted domains to richer environments, and it progressively improves in more challenging scenarios with irreversible actions and interdependent skills. By outperforming both expert-defined operators and baseline methods, SKILLWRAPPER highlights the importance of predicate invention and iterative learning for scaling symbolic representations to embodied tasks.

5 RELATED WORKS

Our work uses pre-trained foundation models for learning symbolic representations of black-box robot skills useful to planning and close to human language and understanding. This work draws ideas from different fields of research such as model learning, abstraction learning, and task and motion planning (TAMP). Several methods have used foundation models (mainly LLMs) as high-level planners (Ahn et al., 2022; Rana et al., 2023; Driess et al., 2023). Several approaches have used foundation models as robot action models (Brohan et al., 2023; Shridhar et al., 2023) or to generate reward functions for robot tasks (Wang et al., 2024b). Concurrent work has also explored how representations can be learned directly from pixels (Athalye et al., 2025). Although these approaches

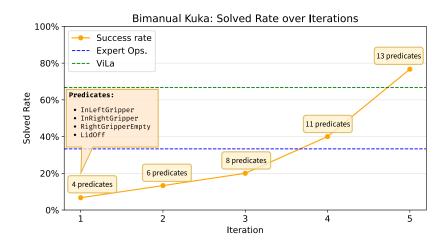


Figure 5: Bimanual Kuka Scenario Results over 5 Iterations with invented predicate count. As the number of predicates grows, SKILLWRAPPER improves over the baseline methods.

show promising results for short-horizon single-skill problems, they fail to scale to complex long-horizon problems (Kambhampati et al., 2024). Lastly, multiple approaches (Han et al., 2024; Liang et al., 2025) have used foundation models to learn symbolic representations of robot skills, but require extensive feedback or prior knowledge from human experts. To the best of our knowledge, our work is the first to use a foundation model to automatically learn the human-interpretable symbolic characterization of robot skills with *theoretical guarantees*.

TAMP has long been used to solve complex robot tasks (Dantam et al., 2018; Shah et al., 2020; Garrett et al., 2021). However, these approaches require symbolic models of the robot skills for task planning. Various approaches have been developed to learn such symbolic models compatible with TAMP solvers from high-dimensional inputs (Konidaris et al., 2018; Silver et al., 2023; Shah et al., 2024). Additionally, the abstract representations learned through these methods are not human-interpretable. On the other hand, we explicitly design our approach to work with high-dimensional inputs and generate human-interpretable abstractions using pre-trained foundation models. We consider abstractions as human-interpretable if they are semantically meaningful and use informative language descriptions. The proposed system also connects to other domains in robotics, and the full related work can be found in Appendix G.

6 Conclusion

We, for the first time, characterize important properties of a learned symbolic model and present the first known approach that employs off-the-shelf foundation models to invent symbolic representations for black-box skills of an agent while providing strong guarantees of soundness and completeness of the learned representations. By combining these theoretical guarantees with foundation model-driven data collection and predicate evaluation, SKILLWRAPPER produces interpretable operators directly usable by classical planners. Empirical results in a simulated burger domain and on real robots demonstrate that SKILLWRAPPER enables efficient long-horizon planning without hand-engineered abstractions, offering a principled path towards scalable skill reasoning.

7 REPRODUCIBILITY STATEMENT

To ensure reproducibility of our work, we provide source code and the prompts used in our experiments as supplementary materials. Although the reproducibility of real-world robot experiments is limited by hardware, simulation experiments run in Robotouille should be reliably reproduced, granted that the checkpoints of the foundation model (i.e., OpenAI's GPT-5 (OpenAI, 2025)) have not been moved.

REFERENCES

- Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, et al. Do As I Can, Not As I Say: Grounding Language in Robotic Affordances. In *Proceedings of the 6th Conference on Robot Learning (CoRL)*, pp. 287–318, 14–18 Dec 2022.
- Michael Ahn, Debidatta Dwibedi, Chelsea Finn, Montserrat Gonzalez Arenas, Keerthana Gopalakrishnan, Karol Hausman, Brian Ichter, et al. AutoRT: Embodied Foundation Models for Large Scale Orchestration of Robotic Agents. In *First Workshop on Vision-Language Models for Navigation and Manipulation (VLMNM) at ICRA 2024*, 2024.
- Ankuj Arora, Humbert Fiorino, Damien Pellier, Marc Métivier, and Sylvie Pesty. A Review of Learning Planning Action Models. *The Knowledge Engineering Review*, 33:e20, 2018.
- Ashay Athalye, Nishanth Kumar, Tom Silver, Yichao Liang, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. Predicate Invention from Pixels via Pretrained Vision-Language Models. In AAAI 2025 Workshop on Language Models for Planning (LM4Plan), 2025.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. In *Proceedings of the 7th Conference on Robot Learning*, pp. 2165–2183, 06–09 Nov 2023.
- Boyuan Chen, Zhuo Xu, Sean Kirmani, Brian Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. SpatialVLM: Endowing Vision-Language Models with Spatial Reasoning Capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14455–14465, 2024.
- William Chen, Oier Mees, Aviral Kumar, and Sergey Levine. Vision-Language Models Provide Promptable Representations for Reinforcement Learning. *Transactions on Machine Learning Research (TMLR)*, 2025. ISSN 2835-8856.
- Sijie Cheng, Zhicheng Guo, Jingwen Wu, Kechen Fang, Peng Li, Huaping Liu, and Yang Liu. EgoThink: Evaluating First-Person Perspective Thinking Capability of Vision-Language Models. In *Proceedings of the 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14291–14302, 2024.
- Cheng Chi, Zhenjia Xu, Chuer Pan, Eric Cousineau, Benjamin Burchfiel, Siyuan Feng, Russ Tedrake, and Shuran Song. Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots. In *Proceedings of Robotics: Science and Systems (RSS) XX*, 2024.
- Neil T Dantam, Zachary K Kingston, Swarat Chaudhuri, and Lydia E Kavraki. An incremental constraint-based framework for task and motion planning. *The International Journal of Robotics Research*, 37(10):1134–1151, 2018.
- S. Doncieux, D. Filliat, N. Díaz-Rodríguez, T. Hospedales, R. Duro, A. Coninx, D.M. Roijers, B. Girard, N. Perrin, and O. Sigaud. Open-ended learning: a conceptual framework based on representational redescription. *Frontiers in Neurorobotics*, 12:59, 2018.
- Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. PaLM-E: An Embodied Multimodal Language Model. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, pp. 8469–8488, 23–29 Jul 2023.
- Kuan Fang, Yuke Zhu, Silvio Savarese, and Li Fei-Fei. Adaptive Procedural Task Generation for Hard-Exploration Problems. In *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, 2021.
- Kuan Fang, Toki Migimatsu, Ajay Mandlekar, Li Fei-Fei, and Jeannette Bohg. Active Task Randomization: Learning Robust Skills via Unsupervised Generation of Diverse and Feasible Tasks. *Proceedings of the 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1–8, 2022.

- Kuan Fang, Fangchen Liu, Pieter Abbeel, and Sergey Levine. MOKA: Open-World Robotic Manipulation through Mark-Based Visual Prompting. *Proceedings of Robotics: Science and Systems (RSS) XX*, 2024.
 - Caelan Reed Garrett, Rohan Chitnis, Rachel Holladay, Beomjoon Kim, Tom Silver, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Integrated Task and Motion Planning. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:265–293, 2021.
 - Gonzalo Gonzalez-Pumariega, Leong Su Yean, Neha Sunkara, and Sanjiban Choudhury. Robotouille: An Asynchronous Planning Benchmark for LLM Agents. In *Proceedings of the 13th International Conference on Learning Representations (ICLR)*, 2025.
 - Jiayuan Gu, Devendra Singh Chaplot, Hao Su, and Jitendra Malik. Multi-skill Mobile Manipulation for Object Rearrangement. In *Proceedings of the 11th International Conference on Learning Representations (ICML)*, 2022.
 - Huy Ha, Pete Florence, and Shuran Song. Scaling Up and Distilling Down: Language-Guided Robot Skill Acquisition. In *Proceedings of the 7th Conference on Robot Learning (CoRL)*, pp. 3766–3777, 2023.
 - Muzhi Han, Yifeng Zhu, Song-Chun Zhu, Ying Nian Wu, and Yuke Zhu. InterPreT: Interactive Predicate Learning from Language Feedback for Generalizable Task Planning. In *Proceedings of Robotics: Science and Systems (RSS) XX*, 2024.
 - Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3D-LLM: Injecting the 3D World into Large Language Models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pp. 20482–20494, 2023.
 - Yingdong Hu, Fanqi Lin, Tong Zhang, Li Yi, and Yang Gao. Look Before You Leap: Unveiling the Power of GPT-4V in Robotic Vision-Language Planning. *arXiv preprint arXiv:2311.17842*, 2023.
 - Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language Models as Zero-Shot Planners: Extracting Actionable Knowledge for Embodied Agents. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, pp. 9118–9147, 2022.
 - Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, Pierre Sermanet, Tomas Jackson, Noah Brown, Linda Luu, Sergey Levine, Karol Hausman, and Brian Ichter. Inner Monologue: Embodied Reasoning through Planning with Language Models. In *Proceedings of the 6th Conference on Robot Learning (CoRL)*, pp. 1769–1782, 2023.
 - Hanxiao Jiang, Binghao Huang, Ruihai Wu, Zhuoran Li, Shubham Garg, Hooshang Nayyeri, Shenlong Wang, and Yunzhu Li. RoboEXP: Action-Conditioned Scene Graph via Interactive Exploration for Robotic Manipulation. In *Proceedings of the 8th Conference on Robot Learning*, pp. 3027–3052, 2025.
 - Minqi Jiang, Edward Grefenstette, and Tim Rocktäschel. Prioritized Level Replay. In *Proceedings* of the 38th International Conference on Machine Learning (ICML), pp. 4940–4950. PMLR, 2021.
 - Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recog*nition, pp. 2901–2910, 2017.
 - Brendan Juba, Hai S. Le, and Roni Stern. Safe Learning of Lifted Action Models. In *Proceedings* of the 18th International Conference on Principles of Knowledge Representation and Reasoning (KR), pp. 379–389, 11 2021.
 - Subbarao Kambhampati, Karthik Valmeekam, Lin Guan, Mudit Verma, Kaya Stechly, Siddhant Bhambri, Lucas Paul Saldyt, and Anil B Murthy. Position: LLMs Can't Plan, But Can Help Planning in LLM-Modulo Frameworks. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024.

- George Konidaris. On the Necessity of Abstraction. *Current Opinion in Behavioral Sciences*, 29:
 1–7, 2019. ISSN 2352-1546.
 - George Konidaris and Andrew Barto. Skill Discovery in Continuous Reinforcement Learning Domains using Skill Chaining. In *Advances in Neural Information Processing Systems (NIPS)*, volume 22, 2009.
 - George Konidaris, Leslie Pack Kaelbling, and Tomas Lozano-Pérez. From Skills to Symbols: Learning Symbolic Representations for Abstract High-Level Planning. *Journal of Artificial Intelligence Research*, 61:215–289, 2018.
 - Leonardo Lamanna, Luciano Serafini, Mohamadreza Faridghasemnia, Alessandro Saffiotti, Alessandro Saetti, Alfonso Gerevini, and Paolo Traverso. Planning for Learning Object Properties. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(10):12005–12013, Jun. 2023.
 - Amber Li and Tom Silver. Embodied Active Learning of Relational State Abstractions for Bilevel Planning. In *Proceedings of The 2nd Conference on Lifelong Learning Agents (CoLLAs)*, pp. 358–375, 2023.
 - Zhaoyi Li, Kelin Yu, Shuo Cheng, and Danfei Xu. LEAGUE++: Empowering Continual Robot Learning via Guided Skill Acquisition with Large Language Models. In *ICLR 2024 Workshop on Large Language Model (LLM) Agents*, 2024.
 - Yichao Liang, Nishanth Kumar, Hao Tang, Adrian Weller, Joshua B. Tenenbaum, Tom Silver, Joao F. Henriques, and Kevin Ellis. VisualPredicator: Learning Abstract World Models with Neuro-Symbolic Predicates for Robot Planning. In *Proceedings of the 13th International Conference on Learning Representations (ICLR)*, 2025.
 - Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mcvay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent Berges, Shiqi Zhang, Pulkit Agrawal, Yonatan Bisk, Dhruv Batra, Mrinal Kalakrishnan, Franziska Meier, Chris Paxton, Alexander Sax, and Aravind Rajeswaran. OpenEQA: Embodied Question Answering in the Era of Foundation Models. In *Proceedings of the 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 16488–16498, 2024.
 - Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B Tenenbaum, and Jiajun Wu. The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences From Natural Supervision. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*, 2019.
 - D. McDermott, M. Ghallab, A. Howe, C. Knoblock, A. Ram, M. Veloso, D. Weld, and D. Wilkins.
 PDDL The Planning Domain Definition Language. Technical report, CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control, 1998.
 - Toki Migimatsu and Jeannette Bohg. Grounding Predicates through Actions. In *Proceedings of the 2022 International Conference on Robotics and Automation (ICRA)*, pp. 3498–3504, 2022.
 - Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. EmbodiedGPT: Vision-Language Pre-Training via Embodied Chain of Thought. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 2024.
 - Soroush Nasiriany, Fei Xia, Wenhao Yu, Ted Xiao, Jacky Liang, Ishita Dasgupta, Annie Xie, Danny Driess, Ayzaan Wahid, Zhuo Xu, Quan Vuong, Tingnan Zhang, Tsang-Wei Edward Lee, Kuang-Huei Lee, Peng Xu, Sean Kirmani, Yuke Zhu, Andy Zeng, Karol Hausman, Nicolas Heess, Chelsea Finn, Sergey Levine, and Brian Ichter. PIVOT: Iterative Visual Prompting Elicits Actionable Knowledge for VLMs. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024.
 - OpenAI. Introducing GPT-5, 2025. URL https://openai.com/index/introducing-gpt-5/. Accessed: September 25, 2025.

- Shreyas Sundara Raman, Vanya Cohen, Ifrah Idrees, Eric Rosen, Ray Mooney, Stefanie Tellex, and David Paulius. CAPE: Corrective Actions from Precondition Errors using Large Language Models. In *Proceedings of the 2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 14070–14077, 2024.
 - Krishan Rana, Jesse Haviland, Sourav Garg, Jad Abou-Chakra, Ian Reid, and Niko Suenderhauf. SayPlan: Grounding Large Language Models using 3D Scene Graphs for Scalable Robot Task Planning. In *Proceedings of the 7th Conference on Robot Learning (CoRL)*, volume 229, pp. 23–72, 06–09 Nov 2023.
 - Allen Z. Ren, Jaden Clark, Anushri Dixit, Masha Itkina, Anirudha Majumdar, and Dorsa Sadigh. Explore until Confident: Efficient Exploration for Embodied Question Answering. In *Proceedings of Robotics: Science and Systems (RSS) XX*, 2024.
 - Pierre Sermanet, Tianli Ding, Jeffrey Zhao, Fei Xia, Debidatta Dwibedi, Keerthana Gopalakrishnan, Christine Chan, Gabriel Dulac-Arnold, Sharath Maddineni, Nikhil J Joshi, et al. RoboVQA: Multimodal Long-Horizon Reasoning for Robotics. In *Proceedings of the 2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 645–652. IEEE, 2024.
 - Naman Shah, Deepak Kala Vasudevan, Kislay Kumar, Pranav Kamojjhala, and Siddharth Srivastava. Anytime Integrated Task and Motion Policies for Stochastic Environments. In *Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9285–9291. IEEE, 2020.
 - Naman Shah, Jayesh Nagpal, Pulkit Verma, and Siddharth Srivastava. From Reals to Logic and Back: Inventing Symbolic Vocabularies, Actions and Models for Planning from Raw Data. *arXiv* preprint arXiv:2402.11871, 2024.
 - C. E. Shannon. A mathematical theory of communication. The Bell System Technical Journal, 27 (3):379–423, 1948.
 - Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-Actor: A Multi-Task Transformer for Robotic Manipulation. In *Proceedings of the 6th Conference on Robot Learning*, volume 205, pp. 785–799, 14–18 Dec 2023.
 - Tom Silver, Rohan Chitnis, Nishanth Kumar, Willie McClinton, Tomás Lozano-Pérez, Leslie Kaelbling, and Joshua B. Tenenbaum. Predicate Invention for Bilevel Planning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(10):12120–12129, Jun. 2023.
 - Theodore Sumers, Kenneth Marino, Arun Ahuja, Rob Fergus, and Ishita Dasgupta. Distilling Internet-Scale Vision-Language Models into Embodied Agents. In *Proceedings of the Fortieth International Conference on Machine Learning (ICML)*, pp. 32797–32818, 2023.
 - Richard S. Sutton, Doina Precup, and Satinder Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1):181–211, 1999.
 - Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Singh Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimír Vondruš, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel Chang, Zsolt Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra. Habitat 2.0: Training Home Assistants to Rearrange their Habitat. In *Advances in Neural Information Processing Systems* (NeurIPS), volume 34, pp. 251–266, 2021.
 - Karthik Valmeekam, Matthew Marquez, Sarath Sreedharan, and Subbarao Kambhampati. On the Planning Abilities of Large Language Models A Critical Investigation. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pp. 75993–76005, 2023.
 - Pulkit Verma, Shashank Rao Marpally, and Siddharth Srivastava. Discovering User-Interpretable Capabilities of Black-Box Planning Agents. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, volume 19, pp. 362–372, 2022.

- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An Open-Ended Embodied Agent with Large Language Models. *Transactions on Machine Learning Research (TMLR)*, 2024a. ISSN 2835-8856.
- Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O Stanley. Paired Open-Ended Trailblazer (POET): Endlessly Generating Increasingly Complex and Diverse Learning Environments and Their Solutions. *arXiv preprint arXiv:1901.01753*, 2019.
- Yufei Wang, Zhanyi Sun, Jesse Zhang, Zhou Xian, Erdem Biyik, David Held, and Zackory Erickson. RL-VLM-F: Reinforcement Learning from Vision Language Foundation Model Feedback. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pp. 51484–51501, 21–27 Jul 2024b.
- Yufei Wang, Zhou Xian, Feng Chen, Tsun-Hsuan Wang, Yian Wang, Katerina Fragkiadaki, Zackory Erickson, David Held, and Chuang Gan. RoboGen: Towards Unleashing Infinite Data for Automated Robot Learning via Generative Simulation. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024c.
- Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. FoundationPose: Unified 6D Pose Estimation and Tracking of Novel Objects. In *Proceedings of the 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17868–17879, 2024.
- Kai Xi, Stephen Gould, and Sylvie Thiébaux. Neuro-Symbolic Learning of Lifted Action Models from Visual Traces. *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, 34(1):653–662, May 2024.
- Jingkang Yang, Yuhao Dong, Shuai Liu, Bo Li, Ziyue Wang, Haoran Tan, Chencheng Jiang, Jiamu Kang, Yuanhan Zhang, Kaiyang Zhou, et al. Octopus: Embodied Vision-Language Programmer from Environmental Feedback. In *Proceedings of the 2024 European Conference on Computer Vision (ECCV)*, pp. 20–38, 2024.
- Naoki Yokoyama, Alex Clegg, Joanne Truong, Eric Undersander, Tsung-Yen Yang, Sergio Arnaud, Sehoon Ha, Dhruv Batra, and Akshara Rai. ASC: Adaptive Skill Coordination for Robotic Mobile Manipulation. *IEEE Robotics and Automation Letters*, 9(1):779–786, 2024.

APPENDIX

A FORMAL FRAMEWORK

We consider a problem setting in which an agent is equipped with a set of predefined, "black-box" skills. The agent can evaluate whether a skill is executable in the current state, but it does not have a complete transition model of the skills a priori, and therefore cannot compose its skills to solve long-horizon problems without considering low-level details. However, if the agent were to learn symbolic models of its skills, it could use classical planning to efficiently compose them to solve new tasks. In this section, we formalize this setting as the problem of *Skill Model Learning*.

A.1 PRELIMINARIES

Environment Model We define an *environment* as a tuple $(\mathcal{S}, \mathcal{T}, \Omega, T)$, where the *state space* \mathcal{S} is assumed to be high-dimensional, continuous, and fully observable; the *type set* \mathcal{T} enumerates the possible object types; and the *skills* Ω are object-centric, such that each $\omega \in \Omega$ is parameterized by object types drawn from the *type set* \mathcal{T} . The *transition function* $T: \mathcal{S} \times \Omega \to \mathcal{S}$ characterizes the environment dynamics but is unknown to the agent. For example, Pour (?teapot, ?mug) may be used to pour tea from a teapot into a mug. However, certain environmental aspects (e.g., which mugs are available to pour into) may differ between settings during learning. We therefore define a specific *setting* by the tuple (s_0, \mathcal{O}) , specifying an initial state $s_0 \in \mathcal{S}$ and a set of typed objects \mathcal{O} , where the type(s) of object $o \in \mathcal{O}$ are denoted $t_o \subseteq \mathcal{T}$.

Black-box Skills We model skills as object-centric options (Sutton et al., 1999) with discrete, object-typed parameters. Formally, a *skill* $\omega \in \Omega$ is defined by a tuple $(\mathcal{I}_{\omega}, \pi_{\omega}, \beta_{\omega}, \Theta_{\omega})$, where the *initiation set* $\mathcal{I}_{\omega} \subseteq \mathcal{S}$ contains states from which the skill may be executed; the *policy* π_{ω} controls the agent during the skill; the *termination set* $\beta_{\omega} \subseteq \mathcal{S}$ is the set of states at which the skill immediately terminates; and the *skill parameters* $\Theta_{\omega} = (\theta_{\omega}^1, \dots, \theta_{\omega}^k)$ specify type constraints $t_{\theta_{\omega}^i} \subseteq \mathcal{T}$ for valid skill arguments. Specifically, a skill $\omega \in \Omega$ may be instantiated to create a *skill instance* ω using objects $o_{1:k} \in \mathcal{O}^k$ if and only if for $1 \le i \le k$, $t_{\theta_{\omega}^i} \subseteq t_{o_i}$. We assume that executing a skill only affects the state of objects passed as arguments.

Symbolic Abstractions Because the agent must evaluate \mathcal{I}_{ω} on individual states, it cannot distinguish essential skill information from irrelevant details, making long-horizon planning combinatorially difficult. However, each skill affects only a few objects at once, leaving the rest of the world unchanged. The agent can exploit this property by instead learning a factored state representation—formalized here as first-order logic models in PDDL (McDermott et al., 1998)—providing an abstract transition model of its skills. Such abstractions enable the agent to use classical planning to compose its skills and accomplish unseen, long-horizon goals.

We use symbolic $\operatorname{predicates} \mathcal{P}$ to express abstract relations between objects. Each predicate $\sigma \in \mathcal{P}$ is a tuple $(C_{\sigma}, \Theta_{\sigma})$, where the $\operatorname{predicate}$ classifier C_{σ} tests whether the predicate holds in a state, given a binding of objects to the $\operatorname{predicate}$ parameters $\Theta_{\sigma} = (\theta_{\sigma}^1, \dots, \theta_{\sigma}^n)$. Each parameter θ_{σ}^i specifies a type constraint $t_{\theta_{\sigma}^i} \subseteq \mathcal{T}$ on corresponding object arguments. Formally, the partial function $C_{\sigma}: \mathcal{O}^n \rightharpoonup (\mathcal{S} \to \{0,1\})$ is defined on objects $o_{1:n} \in \mathcal{O}^n$ if and only if $t_{\theta_{\sigma}^i} \subseteq t_{o_i}$ for $1 \leq i \leq n$. In such cases, we say that the objects $o_{1:n}$ are valid arguments for σ .

By observing different effects, the agent must construct a set of *operators* \mathcal{A} that define abstract transition models for the agent's skills. We define each operator $a \in \mathcal{A}$ by $(\omega_a, \Theta_a, \mathsf{PRE}, \mathsf{EFF}^+, \mathsf{EFF}^-)$, where $\omega_a \in \Omega$ is the corresponding skill; the *operator parameters* $\Theta_a = (\theta_a^1, \dots, \theta_a^m)$ impose type constraints $\mathsf{t}_{\theta_a^i} \subseteq \mathcal{T}$ on operator arguments; the *preconditions* PRE are a conjunction of literals over \mathcal{P} defining the conditions necessary to apply the operator; and the *add* and *delete effects*, EFF^+ and EFF^- , are the subsets of \mathcal{P} that become true and false, respectively, after the operator is applied.

Grounded Abstractions To apply an abstract transition model to a real-world setting, the agent must map the low-level state to an abstract state, which requires *grounding* the known predicates

¹We do not, however, assume that a skill necessarily affects the state of *all* objects passed in as arguments.

²We denote a partial function using $f: \mathcal{X} \rightharpoonup \mathcal{Y}$.

using concrete objects so that their truth value may be determined. A predicate $\sigma \in \mathcal{P}$ may only be grounded using valid arguments, inducing a grounded predicate $\underline{\sigma} = \sigma(o_{1:n})$. Given a low-level state, the classifier $C_{\underline{\sigma}}: \mathcal{S} \to \{0,1\}$ tests whether the predicate holds for the object arguments. We define the set of grounded predicates for a setting (s_0,\mathcal{O}) under predicates \mathcal{P} as $\underline{\mathcal{P}} = \{\sigma(o_{1:n}): \sigma \in \mathcal{P}, o_{1:n} \in \mathcal{O}^n, \bigwedge_{i=1}^n t_{\theta_{\sigma}^i} \subseteq t_{o_i}\}$, inducing an abstract state space $\underline{\mathcal{S}} = 2^{\underline{\mathcal{P}}}$ where each abstract state corresponds to a specific combination of grounded relations.

Given a set of operators \mathcal{A} in a setting (s_0, \mathcal{O}) , the abstract action space $\underline{\mathcal{A}}$ is the set of all valid groundings of the operators using objects from \mathcal{O} . Each abstract action (i.e., grounded operator) is defined as $\underline{a} = a(o_{1:m})$, where $a \in \mathcal{A}$, $o_{1:m} \in \mathcal{O}^m$, and for $1 \leq i \leq m$, $t_{\theta_a^i} \subseteq t_{o_i}$. Grounding an operator induces ground preconditions \underline{PRE} , which are a conjunction of ground literals over $\underline{\mathcal{P}}$; ground add effects $\underline{EFF}^+ \subseteq \mathcal{P}$; and ground delete effects $\underline{EFF}^- \subseteq \mathcal{P}$.

The grounding function $\mathcal{G}:\underline{\mathcal{P}}\to 2^{\mathcal{S}}$ maps each grounded predicate to its grounding set $\mathcal{G}(\underline{\sigma})\subseteq \mathcal{S}$, defined as $\mathcal{G}(\underline{\sigma})=\{s\in\mathcal{S}:C_{\underline{\sigma}}(s)=1\}$. We overload this notation for abstract states $\underline{s}\in\underline{\mathcal{S}}$ so that $\mathcal{G}(\underline{s})=\bigcap_{\underline{\sigma}_i\in\underline{s}}\mathcal{G}(\underline{\sigma}_i)$. Conversely, the abstraction function ABSTRACT: $\mathcal{S}\to\underline{\mathcal{S}}$ maps each low-level state $\underline{s}\in\mathcal{S}$ to the abstract state $\underline{s}\in\underline{\mathcal{S}}$ defined by ABSTRACT $(s)=\{\underline{\sigma}\in\underline{\mathcal{P}}:C_{\sigma}(s)=1\}$.

A.2 PROBLEM DEFINITION

Definition 1. Given an environment (S, T, Ω, T) containing settings $\{(s_0, \mathcal{O})\}_{i=1}^N$, we define a Skill Model Learning problem as learning an abstract transition model $\mathcal{M} = (\mathcal{P}, \mathcal{A})$ for the skills Ω .

After a period of continual learning in one or more settings, an agent may be evaluated on an *skill* planning problem $\mathbf{p}=(s_0,\mathcal{O},\mathcal{S}_g)$, where $\mathcal{S}_g\subseteq\mathcal{S}$ is the set of goal states to be reached. Given a model $\mathcal{M}=(\mathcal{P},\mathcal{A})$, a classical planner can be used to search for an abstract plan $[\underline{a}_1,\ldots,\underline{a}_n]$ that solves the abstract planning problem $\mathbf{p}=(\underline{s}_0,\mathcal{O},\underline{\mathcal{S}}_g)$.

Definition 2. An abstract plan $[\underline{a}_1, \dots, \underline{a}_n]$ is called a solution for skill planning problem $\mathbf{p} = (s_0, \mathcal{O}, \mathcal{S}_g)$ iff for $1 \leq i \leq n$, $s_i = T(s_{i-1}, \underline{\omega}_i)$, $s_{i-1} \in \mathcal{I}_{\underline{\omega}_i}$, and $s_n \in \mathcal{S}_g$, where $\underline{\omega}_i = \underline{\omega}_{\underline{a}_i}$.

In Sec. 3, we describe how SKILLWRAPPER constructs \mathcal{M} from raw skill executions.

B ALGORITHMS

B.1 SKILL SEQUENCE PROPOSAL

Each skill sequence $\sigma = [\omega_1(o_1), \dots, \omega_m(o_m)]$ (Section 3.1) proposed by the foundation model is scored using two heuristics: coverage (C) and chainability (Ch). This section provides more details on how these heuristics are computed and algorithmically used to assign scores to each sequence.

Coverage (C) Coverage (Algorithm 4) evaluates the information gain on all possible pairs of consecutively executed skills over existing transitions after executing a new skill sequence. Specifically, the information gain is measured by the increase in Shannon entropy (Shannon, 1948) over the distribution of all consecutive skill pairs resulting from executing the proposed skill sequence

$$C = \frac{\mathcal{Q}'}{\Sigma \mathcal{Q}'} \times log(\frac{\mathcal{Q}'}{\Sigma \mathcal{Q}'}) - \frac{\mathcal{Q}}{\Sigma \mathcal{Q}} \times log(\frac{\mathcal{Q}}{\Sigma \mathcal{Q}})$$
 (1)

where Q and Q' are matrices that tabulate the number of pairs of consecutively executed skills that occur before and after executing a new skill sequence, respectively. Maximizing coverage encourages the generation of proposed skill sequences that contain the least explored skill pairs. More importantly, this would allow our method to uncover a larger set of interdependencies across the preconditions and effects of all skills.

Chainability (Ch) Chainability predicts the ratio of successful to failed pairs of consecutively executed skills. By computing chainability, we estimate the degree to which the preconditions of operators learned in each iteration are satisfied, and executability can be inferred from the estimated symbolic states and the operators. With an appropriate chainability score, the collected dataset of skill execution traces maintains a balance between number of successful executions and failure executions, which is ideal for identifying possible mismatched pairs and thus inventing predicates.

Skill Sequence Proposal We prompt a foundation model is provided with the agent's skill set Ω and the current set of abstract predicates \mathcal{P} to generate and propose sequences of skills σ , with which the agent collects a dataset of transitions \mathcal{D} . The skill sequence proposal procedure (Algorithm 3) assigns a score tuple (C,Ch) to all sequences and maintains a subset of pareto-front sequences that cannot strictly dominate another sequence, i.e., $(C_i < C_j) \lor (Ch_i < Ch_j) \lor (C_i \le C_j \land Ch_i \le Ch_j)$ where $i \ne j$. An output skill sequence is finally chosen from this pareto-front subset.

Algorithm 3 Skill Sequence Proposal

```
1: Input: Operator set \mathcal{A}, Dataset of skill execution traces \mathcal{D}, batch size of candidate sequences n
2: Output: Proposed skill sequence \sigma
3: \operatorname{seq\_batch} \leftarrow \operatorname{GENERATEGROUNDEDSEQUENCES}(\mathcal{A}, n)
4: \operatorname{Score} \leftarrow \{\}
5: \operatorname{for} \sigma \operatorname{in} \operatorname{seq\_batch} \operatorname{do}
6: \operatorname{cov} \leftarrow \operatorname{COVERAGE}(\mathcal{D}, \sigma)
7: \operatorname{chain}, \operatorname{stateseq} \leftarrow \operatorname{CHAINABILITY}(\mathcal{A}, \mathcal{P}, \sigma)
8: \operatorname{suff} \leftarrow \operatorname{CONSISTENCY}(\mathcal{D}, \operatorname{stateseq})
9: \operatorname{Score}[\sigma] \leftarrow (\operatorname{cov}, \operatorname{chain}, \operatorname{suff}, \operatorname{stateseq})
10: \operatorname{end} \operatorname{for}
11: \sigma^* \leftarrow \operatorname{PARETOOPTIMALITY}(\operatorname{Score}, \mathcal{D}, \mathcal{A}, \mathcal{P})
12: \operatorname{return} \sigma^*
```

Algorithm 4 Coverage Scoring Function

```
1: Input: Dataset of skill execution traces \mathcal{D}, Proposed skill sequence \sigma
 2: Output: Coverage score C
3: Q \leftarrow 0
                                                                                \triangleright construct a matrix \mathcal{Q} \in \mathcal{M}_{|\Omega| \times |\Omega|} of skill-pair counts
 4: for \tau_i, \tau_{i+1} in \mathcal{D} do
                                                                                           \triangleright iterate over all pairs of transitions in dataset \mathcal D
           \tau_i = (s_i, \omega(\mathbf{o})_i, s_i)
           \tau_{i+1} = (s_{i+1}, \omega(\mathbf{o})_{i+1}, s_{i+1})
 6:
 7:
           \mathcal{Q}[\omega(\mathbf{o})_i][\omega(\mathbf{o})_{i+1}] + = 1
 8: end for
9: Q \leftarrow Q
                                                                                                                  ⊳ new skill-pair count initialized
10: ent_{curr} = -1 \times (Q/\sum Q) \times log(Q/\sum Q)
                                                                                                                           \triangleright compute entropy with Q
                                                                           \triangleright iterate over all pairs of consecutively executed skills in \sigma
11: for (\omega(\mathbf{o})_i, \omega(\mathbf{o})_{i+1}) in \sigma do
            Q'[\omega(\mathbf{o})_i, \omega(\mathbf{o})_{i+1}] + = 1
14: ent_{new} = -1 \times (\mathcal{Q}' / \sum \mathcal{Q}') \times log(\mathcal{Q}' / \sum \mathcal{Q}')
                                                                                                                          \triangleright compute entropy with Q'
15: C = ent_{new} - ent_{curr}
                                                                                                                            16: return C, Q
```

Algorithm 5 Chainability Scoring Function

918 919

920

921

922

923

924

925

926

927

928

929

930

931

932

933

934

935

936

937 938 939

30: end if

```
1: Input: Operator set A, Symbolic predicate set P, Proposed skill sequence \sigma
 2: Output: Chainability score Ch, sequence of states seq
 3: len_{tot} \leftarrow len(t)
                                                                                                               ⊳ store total number of skills
 4: len_{exec} \leftarrow 0
                                                                                               > store total number of executable skills
 5: s_{curr} = \{ p : \mathcal{F}_p(s) \text{ for } p \text{ in } \mathcal{P} \}

    b track the current state

 6: seq = [s_{curr}]
                                                                                                        > store sequence of abstract states
 7: for \omega in t do
                                                                                                                  \triangleright iterate over all skills in t
          if \forall p in \text{PRE}_{\mathcal{A}[\omega]} \wedge \mathcal{F}_p(s_{curr}) = 1 then
9:
               len_{exec} = len_{exec} + 1
                for p_{\mathrm{Eff}^+} in \mathrm{Eff}^+_{\mathcal{A}[\omega]} do
10:
11:
                     s_{curr}[p_{\rm Eff^+}] = 1
12:
                end for
                for p_{\mathrm{EfF}^-} in \mathrm{EfF}_{\mathcal{A}[\omega]}^- do
13:
                     s_{curr}[p_{\rm Eff^-}] = 0
14:
15:
                end for
16:
          end if
17:

    b add current state to sequence

          seq \leftarrow seq \cup s_{curr}
18: end for
19: Ch = |(len_{exec}/len_{tot}) - 0.5|
                                                                                                               ⊳ compute chainability score
20: return Ch, seq
```

```
940
            Algorithm 6 Scoring Functions for Invented Predicates
941
             1: Input: Operator set A, skill \omega, dataset of skill execution traces \mathcal{D}, abstraction function \Gamma
942
             2: Parameters: Threshold h
943
944
             3: ScorePrecond:
945
             4: total \leftarrow 0
946
             5: valid \leftarrow 0
947
             6: for \langle s, \underline{\omega}, s' \rangle \in \mathcal{D} do
948
                       if \exists a \in \mathcal{A}, \mathbf{o} \subseteq \mathcal{O}, s.t. \mathbf{s} \in \mathsf{PRE}_{a_\mathbf{o}}, and \mathbf{s} \in I_\omega then
949
             8:
                            valid + +
950
                       end if
             9:
951
            10:
                       total + +
952
            11: end for
            12: if valid/total > h then
953
                       return True
            13:
954
            14: else
955
            15:
                       return False
956
            16: end if
957
958
            17: ScoreEff:
959
            18: total \leftarrow 0
960
            19: valid \leftarrow 0
961
            20: for \langle s, \underline{\omega}, s' \rangle \in \mathcal{D} do
962
                       if \exists a \in \mathcal{A}, \mathbf{o} \subseteq \mathcal{O}, s.t. \Gamma(s') \setminus \Gamma(s) = \mathrm{EfF}_{a_0}, and s \in I_\omega then
            21:
963
            22:
                            valid + +
            23:
                       end if
964
            24:
                       total + +
965
            25: end for
966
            26: if valid/total > h then
967
            27:
                       return True
968
            28: else
969
                       return False
            29:
970
```

B.2 OPERATOR LEARNING

972

973 974

975

976

977

978

979

980

981

982

983

984

985

986

987 988 989

990 991

992

993

994

995

996 997

998

999

1000

1001

1002 1003

1004

1005

1008

1009 1010 1011

1012 1013 1014

1015 1016

1017

1018

1019

1020

1021

1022

1023 1024

1025

Algorithm 7 Learn Operators

```
1: Input: Skill execution traces \mathcal{D}(\underline{\omega}) = \{\langle s, \underline{\omega}, s' \rangle\}_{\omega}, abstraction function \Gamma
 2: Output: Operator set A
 3: eff_dict \leftarrow defaultdict()

    Store clustered effects

 4: for \langle s, \underline{\omega}, s' \rangle \in \mathcal{D} do
             eff \leftarrow \Gamma(s')/\Gamma(s)
             eff\_dict[eff] \leftarrow eff\_dict[eff] \cup ((s, s'))
 6:
 7: end for
 8: \mathcal{A} \leftarrow []
 9: for eff \in eff_dict do
             execution\_list \leftarrow eff\_dict[eff]
10:
11:
             precond \leftarrow \Pi_{\langle s,\underline{\omega},s'\rangle \in \text{execution\_list}}\Gamma(s)
             \mathcal{A} \leftarrow \mathcal{A} \cup [\text{precond}, \text{eff}]
13: end for
14: return A
```

PROPERTIES OF LEARNED SYMBOLIC MODELS

Relational predicates are the basic units of the abstract representation of the low-level state space. In this section, we characterize the conditions of the learned representations of a finite set of skills using relational predicates to support high-level planning, in the context of model learning. From here, chaining the skills is enabled by applying predicates from the representation of each low-level skill to others for clustering (described in Section ??).

Effective skill planning requires an accurate abstract model and grounding function. All forms of abstractions are typically lossy, i.e., while learning an abstract transition model, certain low-level environment details may not be captured. Conversely, the learned model must accurately retain the information needed to produce sound and complete plans. In this section, we characterize the conditions under which an abstract model facilitates exact, sound, and/or complete planning, for the purpose of constraining how we can *construct* such a model.

To begin formalizing the relationship between a skill $\omega \in \Omega$ and some abstract model $\mathcal{M} = (\mathcal{P}, \mathcal{A})$, we define the alpha set α_{ω} and zeta set ζ_{ω} , representing the states in which the model predicts that the skill may either be initiated (when $s \in \alpha_{\omega}$) or terminated (when $s \in \zeta_{\omega}$), respectively.

Definition 3. Given a skill instance $\underline{\omega} \in \underline{\Omega}$ and abstract model $\mathcal{M} = (\mathcal{P}, \mathcal{A})$, we define the **instance** alpha set $\alpha_{\underline{\omega}}$ and instance zeta set $\zeta_{\underline{\omega}}$ as follows:

$$\alpha_{\underline{\omega}} = \bigcup_{\underline{a} \in \underline{A}_{\underline{\omega}}} \mathcal{G}(\underline{PRE}_{\underline{a}})$$

$$\zeta_{\underline{\omega}} = \bigcup_{\underline{a} \in \underline{A}_{\omega}} \mathcal{G}((\underline{PRE}_{\underline{a}} \setminus \underline{EFF}_{\underline{a}}^{-}) \cup \underline{EFF}_{\underline{a}}^{+})$$
(2)

$$\zeta_{\underline{\omega}} = \bigcup_{\underline{a} \in \underline{A}_{\underline{\omega}}} \mathcal{G}((\underline{\operatorname{Pre}}_{\underline{a}} \setminus \underline{\operatorname{Eff}}_{\underline{a}}^{-}) \cup \underline{\operatorname{Eff}}_{\underline{a}}^{+})$$
(3)

For a non-instantiated skill $\omega \in \Omega$, we define the alpha set $\alpha_{\omega} = \bigcup_{\underline{\omega}} \alpha_{\underline{\omega}}$ and zeta set $\zeta_{\omega} = \bigcup_{\underline{\omega}} \zeta_{\underline{\omega}}$.

We define an exact abstract model as one that perfectly captures the initiation set and termination set of all skills. Although such a representation is infeasible to learn in practice, its properties provide an "ideal case" from which other definitions can weaken assumptions.

Definition 4. (Exact Model). Let $\mathcal{M} = (\mathcal{P}, \mathcal{A})$ be a model for environment $(\mathcal{S}, \mathcal{T}, \Omega, T)$ and the set of skills Ω , and let α_{ω} and ζ_{ω} be the approximate initiation and termination sets (Def. 3). The model \mathcal{M} is **an exact model** iff:

$$\forall \omega \in \Omega, s \in \mathcal{S} \colon s \in \mathcal{I}_{\omega} \iff s \in \alpha_{\omega} \text{ and } \beta_{\omega}(s) = 1 \iff s \in \zeta_{\omega}. \tag{4}$$

An exact planning model supports accurate planning as it precisely characterizes skills' initiation and termination sets. However, this accuracy comes at the cost of practical feasibility, as any exact

model achieves very little in terms of *abstraction*: it must express the full initiation and termination sets of each skill! Therefore, in many settings, alternative model properties that approximate the exactness of the learned model, namely *soundness*, *suitability*, and *correctness*, may be preferred as objectives for model learning. We now define these properties.

A symbolic model is *sound* if it correctly predicts the effects of a plan: whenever a complete and sound planner predicts that a sequence of skills will reach some abstract state, executing the corresponding skills in the environment truly leads there. Soundness rules out spurious symbolic transitions that do not correspond to realizable outcomes. Formally, we define soundness as follows: Let $\mathcal{M} = (\mathcal{P}, \mathcal{A})$ denote a symbolic planning model, where \mathcal{P} is a finite set of predicates defining an *abstract state space* $\bar{\mathcal{S}}$, and \mathcal{A} is a set of *abstract actions (skills)* with preconditions and effects expressed in terms of \mathcal{P} . Each abstract state $\bar{s} \in \bar{\mathcal{S}}$ is obtained by a learned *grounding function* $\Gamma: \mathcal{S} \to \bar{\mathcal{S}}$ that maps low-level agent states $s \in \mathcal{S}$ to truth assignments over \mathcal{P} .

Definition 5 (Soundness). The model \mathcal{M} is sound iff, for any valid plan π produced by a complete symbolic planner over \mathcal{M} and for all task instances $\mathbf{p}_i \in \mathbf{P}$,

$$\Gamma(\mathcal{T}(\pi, s_0)) = \bar{\mathcal{T}}_{\mathcal{M}}(\pi, \Gamma(s_0)),$$

where s_0 is the initial state, $\mathcal{T}(\pi, s_0)$ is the set of states reachable by executing π from s_0 , and $\bar{\mathcal{T}}_{\mathcal{M}}(\pi, \Gamma(s_0))$ is the abstract state predicted by \mathcal{M} after executing π from $\Gamma(s_0)$.

A symbolic model is *complete* if it never omits real solutions: whenever the environment admits a way to solve a task, the planner can find a corresponding abstract plan in the model. Completeness rules out gaps in symbolic coverage that would make feasible problems appear unsolvable.

Definition 6 (Completeness). The model $\mathcal{M} = (\mathcal{P}, \mathcal{A})$ is complete if, for any task instance $\mathbf{p}_i \in \mathbf{P}$ and any sequence of low-level actions that achieves the goal from an initial state s_0 , there exists a symbolic plan π over \mathcal{M} such that

$$\Gamma(\mathcal{T}(\pi, s_0)) \models G_i,$$

where G_i is the goal condition of \mathbf{p}_i .

A symbolic model is *suitable* if it correctly characterizes *when* a skill can be applied, meaning the symbolic preconditions predicted by the model align with the skill's real initiation conditions. A skill is applicable in an abstract state *iff* it is applicable in the corresponding grounded state.

Definition 7 (Suitability). The model \mathcal{M} is suitable if, for any valid plan π produced by a complete symbolic planner and for all task instances $\mathbf{p}_i \in \mathbf{P}$,

$$\Gamma(\mathcal{T}(\underline{a_i}, s_0)) \in \bar{I}_{a_{i+1}} \iff \mathcal{T}(\omega_i, s_0) \in I_{\omega_{i+1}}, \quad \forall a \in \mathcal{A}, \forall \omega \in \Omega$$

where s_0 is an initial agent state, $\mathcal{T}(\pi, s_0)$ is the state reached by executing π from s_0 , and \bar{I}_a is the abstract initiation set of skill a.

The minimum requirement for a model \mathcal{M} to solve abstract planning problems is that it is *suitable* and *complete*, such that an abstract plan can always be found, and that it is always executable, if there exists a skill plan as a solution. Although methods for constructing such models exist and have been investigated in previous work, they do not provide these guarantees.

D Proofs

Lemma 2 (Predicate invention). Let \mathcal{D} be the set of transition tuples $\langle s, \omega, s' \rangle$ with $s, s' \in \mathcal{S}$ and $\omega \in \Omega$. Let B_n be a buffer of n i.i.d. samples from \mathcal{D} , each labeled with outcome and containing at least one successful and one failed transition. Let $\mathcal{M}_n = (\mathcal{P}_n, \mathcal{A}_n)$ be the model learned by SKILLWRAPPER from B_n , where each operator in \mathcal{A}_n corresponds to some skill $\omega \in \Omega$. Then, for every $\omega \in \Omega$ appearing in B_n ,

$$I_{\omega} \triangle \alpha_{\omega} = \emptyset, \tag{5}$$

$$\beta_{\omega} \triangle \zeta_{\omega} = \emptyset,$$
 (6)

where $\alpha_{\omega} = \bigcup_{a \in \mathcal{A}_n^{\omega}} \mathcal{G}(\mathsf{PRE}_a)$, $\zeta_{\omega} = \bigcup_{a \in \mathcal{A}_n^{\omega}} \left(\mathcal{G}(\mathsf{PRE}_a) \setminus \mathcal{G}(\mathsf{Eff}_a^-) \right) \cup \mathcal{G}(\mathsf{Eff}_a^+)$, and \triangle is the symmetric difference.

Proof sketch. Suppose for contradiction that $\exists \langle s, \omega, s' \rangle \in B_n$ with $s \in I_\omega \triangle \alpha_\omega$. Two cases arise:

- (i) False positive. $s \notin I_{\omega}$ but $s \in \alpha_{\omega}$. By construction of α_{ω} , there must also exist $s_j \in I_{\omega} \cap \alpha_{\omega}$. Thus $1_{\alpha_{\omega}}(s) = 1_{\alpha_{\omega}}(s_j)$ while $I_{\omega}(s) \neq I_{\omega}(s_j)$, contradicting the update rule.
- (ii) False negative. $s \in I_{\omega}$ but $s \notin \alpha_{\omega}$. Then some $s_j \notin I_{\omega} \cap \alpha_{\omega}$ must also exist, yielding the same contradiction.

Hence no such s exists, and Eq. (6) holds. The proof for Eq. (7) is identical, replacing α_{ω} with ζ_{ω} defined from operator effects.

Theorem 3 (Probabilistic-completeness of SKILLWRAPPER). Let \mathcal{M}^* be a Complete Model for a set of skills Ω , where each $\omega \in \Omega$ has initiation set $I_{\omega} \subseteq \mathcal{S}$ and termination set $\beta_{\omega} \subseteq \mathcal{S}$. Let μ be a probability distribution over $\mathcal{S} \times \mathcal{A} \times \mathcal{S}$. Consider a finite hypothesis class \mathcal{H} , where each $\mathcal{M} \in \mathcal{H}$ assigns a learned initiation set $\widehat{I}_{\omega} = \bigcup_{a \in \mathcal{A}_{\omega}} \mathcal{G}(\mathsf{PRE}_a)$ and termination set $\widehat{\beta}_{\omega} = \bigcup_{a \in \mathcal{A}_{\omega}} \left(\mathcal{G}(\mathsf{PRE}_a) \setminus \mathcal{G}(\mathsf{EfF}_a^-)\right) \cup \mathcal{G}(\mathcal{G}_a^+)$ to each $\omega \in \Omega$.

For any $M \in \mathcal{H}$, define

$$d_{\text{compl}}(\mathcal{M}, \mathcal{M}^*) = \Pr_{(s, \omega, s') \sim \mu} [(s \in \widehat{I}_{\omega} \triangle I_{\omega}) \lor (s' \in \widehat{\beta}_{\omega} \triangle \beta_{\omega})].$$

Let n i.i.d. samples $\{(s_i, \omega_i, s_i')\}_{i=1}^n$ be drawn from μ . Then, for every $\epsilon > 0$,

$$\Pr\left[d_{\text{compl}}(\widehat{\mathcal{M}}_n, \mathcal{M}^*) \le \epsilon\right] \ge 1 - |\mathcal{H}| e^{-n\epsilon},$$

i.e., with high probability, $\widehat{\mathcal{M}}_n$ misses fewer than an ϵ -fraction of feasible transitions under μ .

Proof sketch. For $\mathcal{M} \in \mathcal{H}$, define the *true error*

$$\operatorname{Err}(\mathcal{M}) = d_{\operatorname{compl}}(\mathcal{M}, \mathcal{M}^*),$$

and the empirical error

$$\widehat{\operatorname{Err}}(\mathcal{M}) = \frac{1}{n} \sum_{i=1}^{n} \mathbf{1} \left[(s_i \in \widehat{I}_{\omega} \triangle I_{\omega}) \vee (s_i' \in \widehat{\beta}_{\omega} \triangle \beta_{\omega}) \right].$$

By Lemma 2, $\widehat{\operatorname{Err}}(\widehat{\mathcal{M}}_n)=0$. Suppose some $\mathcal{M}\in\mathcal{H}$ satisfies $\operatorname{Err}(\mathcal{M})\geq\epsilon$. Then each sample has probability at least ϵ of revealing an error. The chance of seeing none in n i.i.d. draws is at most $e^{-n\epsilon}$ (by Hoeffding/Chernoff).

Applying a union bound over all $\mathcal{M} \in \mathcal{H}$,

$$\Pr[\exists \mathcal{M} \in \mathcal{H} : \operatorname{Err}(\mathcal{M}) \ge \epsilon \wedge \widehat{\operatorname{Err}}(\mathcal{M}) = 0] \le |\mathcal{H}| e^{-n\epsilon}.$$

Since $\widehat{\mathcal{M}}_n$ has $\widehat{\operatorname{Err}} = 0$, the event $\operatorname{Err}(\widehat{\mathcal{M}}_n) \geq \epsilon$ is contained in this bound. Thus, with probability at least $1 - |\mathcal{H}| e^{-n\epsilon}$, $d_{\operatorname{compl}}(\widehat{\mathcal{M}}_n, \mathcal{M}^*) \leq \epsilon$.

Theorem 4 (Soundness of SKILLWRAPPER). Let \mathcal{T} be the set of transition tuples $\langle s, \omega, s' \rangle$ with $s \in \mathcal{S}$ and $\omega \in \Omega$. Let B_n be an experience buffer of n samples drawn from \mathcal{T} , each labeled with outcome. Suppose SKILLWRAPPER learns a model

$$\mathcal{M}_n = (\mathcal{P}_n, \mathcal{A}_n),$$

where each $a \in A_n$ corresponds to some skill $\omega \in \Omega$. Then, for every operator $a \in A_n$ associated with ω , there exists a real transition $\langle s, \omega, s' \rangle \in B_n$ such that

$$\underline{s} \models \mathsf{PRE}_a \quad and \quad \underline{s}' \models \mathsf{EFF}_a.$$

Proof sketch. By construction, SKILLWRAPPER derives each operator $a \in \mathcal{A}_n$ from transitions in B_n through its wrapper procedure. If a is associated with skill ω , then its preconditions PRE_a are obtained from the abstract representation of some observed s, and its effects EFF_a from the corresponding s'. Hence there must exist $\langle s, \omega, s' \rangle \in B_n$ such that $\underline{s} \models \text{PRE}_a$ and $\underline{s'} \models \text{EFF}_a$.

If no such transition existed, then a would be unsupported by data and would not have been generated. Thus every operator in A_n corresponds to at least one valid observed transition, proving soundness.

Proof (by Contradiction). Assume, for contradiction, that there is an operator $a \in \mathcal{A}_n$ for which no transition $\langle s, \omega, s' \rangle \in B_n$ supports it. That would mean:

(1) a is in the learned model,

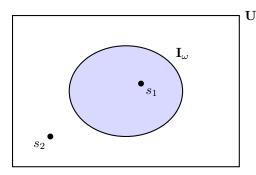
(2) a has no real sample $\langle s, \omega, s' \rangle \in B_n$ s.t. $\underline{s} \models \mathsf{PRE}_a \land \underline{s}' \models \mathsf{EFF}_a$.

However, SKILLWRAPPER introduces or refines operators *only* in response to observed transitions $\langle s, \omega, s' \rangle$ that cannot be explained by any existing operator in \mathcal{A}_n . Therefore, if a exists in the final model, it must have been created when the system encountered a transition $\langle s, \omega, s' \rangle$ with \underline{s} and \underline{s}' not accounted for by any previously existing operator. That transition becomes the "anchor" for a's preconditions and effects.

Hence, there *must* be at least one real transition $\langle s, \omega, s' \rangle \in B_n$ matching the preconditions/effects of a, contradicting assumption (2). Consequently, our assumption is false, and each operator indeed has a supporting transition in B_n . This completes the proof.

D.1 ILLUSTRATION OF PREDICATE INVENTION IN LOW-LEVEL SPACE

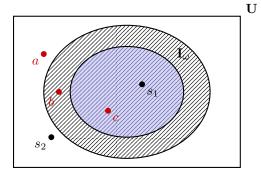
Following the last section, we illustrate how SKILLWRAPPER invents new predicates under all circumstances with guarantees using the condition. We discuss the case of precondition here, and guarantees of effect follow the same logic.



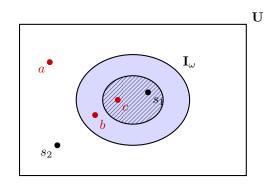
By assumption, there exist two transitions: $\langle s_1, \omega, s_1' \rangle, \langle s_2, \omega, s_2' \rangle$ such that $s_1 \in I_\omega, s_2 \notin I_\omega$ initially.

There are three possible circumstances of the resulting learning model from $\langle s_1, \omega, s_1' \rangle$ and $\langle s_2, w, s_2' \rangle$. For each of them, we discuss all possible cases of more transitions with the starting state falling into each section of the state space.

(1). The learned model has $I_{\omega} \subset \alpha$. Then, for each section that initial states of new transitions can fall in:



- (a): $a \notin \alpha, a \notin I_{\omega}$. Thus, $\nexists \langle s, \omega, s' \rangle$ such that $\mathbf{1}_{\alpha}(a) = \mathbf{1}_{\alpha}(s)$ while $I_{\omega}(a) \neq I_{\omega}(s)$. No additional predicate need to be invented.
- (b): $b \in \alpha, a \notin I_{\omega}$. Thus, $\mathbf{1}_{\alpha}(b) = \mathbf{1}_{\alpha}(s_1)$ while $I_{\omega}(a) \neq I_{\omega}(s)$. New predicate will be invented.
- (c): $a \in \alpha, a \in I_{\omega}$. Thus, $\nexists \langle s, \omega, s' \rangle$ such that $\mathbf{1}_{\alpha}(a) = \mathbf{1}_{\alpha}(s)$ while $I_{\omega}(a) \neq I_{\omega}(s)$. No additional predicate need to be invented.
- (2). The learned model has $\alpha \subset I_{\omega}$. Then, for each section that initial states of new transitions can fall in:



1189 1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1203

1204

1205

1206

1208

1209

1210

1211

1212

1213

1214

1215 1216

1217

1218 1219

1220

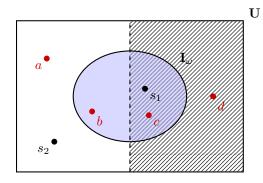
1222

• (a): $a \notin \alpha, a \notin I_{\omega}$. Thus, $\nexists \langle s, \omega, s' \rangle$ such that $\mathbf{1}_{\alpha}(a) = \mathbf{1}_{\alpha}(s)$ while $I_{\omega}(a) \neq I_{\omega}(s)$. No additional predicate need to be invented.

• **(b)**: $b \notin \alpha, a \in I_{\omega}$. Thus, $\mathbf{1}_{\alpha}(b) = \mathbf{1}_{\alpha}(s_1)$ while $I_{\omega}(a) \neq I_{\omega}(s)$. New predicate will be invented.

• (c): $a \in \alpha, a \in I_{\omega}$. Thus, $\nexists \langle s, \omega, s' \rangle$ such that $\mathbf{1}_{\alpha}(a) = \mathbf{1}_{\alpha}(s)$ while $I_{\omega}(a) \neq I_{\omega}(s)$. No additional predicate need to be invented.

(3). The learned model has $\alpha \cap I_{\omega} \neq \emptyset$, $\alpha \nsubseteq I_{\omega}$, $I_{\omega} \nsubseteq \alpha$. Then, for each section that initial states of new transitions can fall in:



• (a): $a \notin \alpha, a \notin I_{\omega}$. Thus, $\nexists \langle s, \omega, s' \rangle$ such that $\mathbf{1}_{\alpha}(a) = \mathbf{1}_{\alpha}(s)$ while $I_{\omega}(a) \neq I_{\omega}(s)$. No additional predicate need to be invented.

• (b): $b \notin \alpha, a \in I_{\omega}$. Thus, $\mathbf{1}_{\alpha}(b) = \mathbf{1}_{\alpha}(s_1)$ while $I_{\omega}(a) \neq I_{\omega}(s)$. New predicate will be invented.

• (c): $a \in \alpha, a \in I_{\omega}$. Thus, $\nexists \langle s, \omega, s' \rangle$ such that $\mathbf{1}_{\alpha}(a) = \mathbf{1}_{\alpha}(s)$ while $I_{\omega}(a) \neq I_{\omega}(s)$. No additional predicate need to be invented.

• (d): $d \in \alpha, a \notin I_{\omega}$. Thus, $\mathbf{1}_{\alpha}(d) = \mathbf{1}_{\alpha}(s_1)$ while $I_{\omega}(d) \neq I_{\omega}(s)$. New predicate will be invented.

So far, we have discussed all possible cases where initial states of new transitions can fall in, and the predicate invention condition is proven to handle all cases.

E LEARNED OPERATORS AND CASE STUDIES

E.1 LEARNED OPERATORS

```
1223
        # Example Learned Predicate and Operator of Burger domain
1224
1225
        # Predicate
          name: cut into pieces
1226
          types:
1227
          - cuttable
          semantic: "the_cuttable_appears_as_at_least_two_non-touching_visible_pieces_(has_multiple_
1228
          disconnected_regions),_indicating_it_is_cut_rather_than_a_single_intact_piece.
1229
1230
        # Operator
        (:action Cut 7
1231
         :parameters (?cuttable_p0 - cuttable
                   ?pickupable_p1 - pickupable
1232
                   ?pickupable_p2 - pickupable
1233
                   ?pickupable_p6 - pickupable
1234
                   ?robot\_p5 - robot
                   ?station_p4 - station)
1235
         :precondition (and
            (not (= ?pickupable_p1 ?pickupable_p2))
1236
            (not (= ?pickupable_p1 ?pickupable_p6))
            (not (= ?pickupable_p2 ?pickupable_p6))
            (gripper_empty ?robot_p5)
1238
            (on_cutting_board ?cuttable_p0)
1239
            (on_station ?cuttable_p0 ?station_p4)
            (top_most ?cuttable_p0)
1240
            (not (cut_into_pieces ?cuttable_p0))
1241
            (not (holding ?robot_p5 ?cuttable_p0))
            (not (holding ?robot_p5 ?pickupable_p1))
```

```
1246
        # Example Learned Predicate and Operator of Franka domain
1247
        # Predicate
1248
          name: plate_is_dirty
1249
          types:
          - plate
1250
          semantic: the specified plate's upper surface contains visible food traces whose appearance
1251
          differs from the plate's base surfac (visible residue to be wiped).
1252
        # Operator
1253
         (:action Stack_5
         :parameters (?pickupable_p0 - pickupable
1254
                   ?pickupable_p1 - pickupable
?pickupable_p3 - pickupable
1255
                   ?plate_p4 - plate
1256
                   ?robot_p2 - robot)
1257
         :precondition (and
            (not (= ?pickupable_p0 ?pickupable_p1))
1258
            (not (= ?pickupable_p0 ?pickupable_p3))
1259
            (not (= ?pickupable_p1 ?pickupable_p3))
1260
            (holding ?robot_p2 ?pickupable_p1)
            (plate_top_unoccupied ?plate_p4)
1261
            (not (gripper_empty ?robot_p2))
            (not (holding ?robot_p2 ?pickupable_p0))
1262
            (not (holding ?robot_p2 ?pickupable_p3))
1263
            (not (plate_is_dirty ?plate_p4))
            (not (stacked_on ?pickupable_p1 ?plate_p4)))
1264
         :effect (and
1265
            (gripper_empty ?robot_p2)
            (stacked_on ?pickupable_p1 ?plate_p4)
1266
            (not (holding ?robot_p2 ?pickupable_p1))
1267
            (not (plate_top_unoccupied ?plate_p4))))
1268
```

```
1269
         # Example Learned Predicate and Operator of Bi-manual Kuka domain
1270
         # Predicate
1271
          name: Coated
1272
          types:
          - utensil
1273
          semantic: a visible layer or clump of material adheres to the utensil's working end (e.g.,
          the blade shows a smear that was absent before).
1274
1275
         # Operator
1276
         (:action Scoop_156
1277
         :parameters (?openable_p0 - openable
                   ?robot_p1 - robot
1278
                    ?utensil_p2 - utensil)
1279
         :precondition (and
            (HeldByRobot ?robot_p1 ?openable_p0)
1280
            (InLeftGripper ?robot_p1 ?openable_p0)
1281
            (InRightGripper ?robot_p1 ?utensil_p2)
            (LidOff ?openable_p0)
1282
            (not (Closed ?openable_p0))
1283
            (not (Coated ?utensil_p2))
            (not (InContainer ?utensil_p2))
1284
            (not (LeftGripperEmpty ?robot_p1))
(not (OpenableOnTable ?openable_p0))
1285
            (not (RightGripperEmpty ?robot_p1))
1286
            (not (UtensilOnTable ?utensil_p2)))
1287
         :effect (and
            (Coated ?utensil_p2)))
1288
```

Additionally, we present here a predicate and an operator written by the PDDL expert:

1289

```
# Predicate and Operator Written by PDDL Expert

# Predicate
name: is_on_station
types:
- pickupable
- station
semantic: "A_'pickupable'_object_is_on_top_of_a_'station'."
```

```
1296
1297
        # Operator
        (:action Cut
1298
         :parameters (?robot - robot
1299
                   ?cuttable - cuttable
                   ?board - cuttingboard)
1300
         :precondition (and
1301
            (hand_empty)
            (obj free ?cuttable)
1302
            (is_on_station ?cuttable ?board)
1303
            (not (is cut ?cuttable)))
         :effect (and
1304
            (is cut ?cuttable)))
1305
```

E.2 CASE STUDIES IN BURGER

 (1) failure mode: over partitioning (2) sys preds not having topmost() or clear()

System predicates are designed in a way such that every possible state can be specified with a compact set of eight predicates. However, this is proven insufficient to describe the model dynamics. For example, the predicate clear_above or its semantic equivalent is necessary for describing the precondition of all operators learned from the *Pick* skill, such that the model can express "Picking an item with another item on top will fail". Since predicate invention is disabled in Sys Preds., it often returns invalid plans that try to pick objects that are underneath a stack. As a result, the lack of soundness induces a low solved rate for hard and impossible problems of System Predicates.

Random exploration is heavily limited by inefficient data gathering. However, if we take a close look at the operators it learns, they are of good quality, which is the benefit of having the same predicate invention algorithm of SKILLWRAPPER. In fact, one frequent failure mode is never to execute certain skills with complex preconditions, e.g., *Cook* requires the item on the stove and the agent's hand to be empty. Since the same predicate invention algorithm with SKILLWRAPPER is being used, and so can only learn operators for a subset of skills, did not find all operators for *Pick*, and thus can only solve the simplest pick and stack tasks.

Failure mode of SKILLWRAPPER Comparing the poorly performing PDDL operators learned by SKILLWRAPPER to Expert Operator, SKILLWRAPPER only learns one extra predicate, *on_cutting_board(item)*, which divides the previous cluster that shares the same effect into even smaller clusters. Then, the transition data in these smaller cluster cannot support the model learning algorithm to effectively eliminate spurious preconditions. This case study points out that the balance between predicate invention and data gathering is a critical factor in the learning process: if too many predicates are invented without adequate transition data, the resulting operators could possibly contain spurious preconditions, such that it cannot generalize at all. In SKILLWRAPPER, the balance is controlled by empirically tuning the threshold of the scoring function and the length of each skill sequence proposed.

F IMPLEMENTATION DETAILS

E 1 Don

F.1 ROBOT EXPERIMENTS

Single-Arm Manipulation We employ a Franka Emika Research 3 robotic arm equipped with a UMI gripper (Chi et al., 2024). The workspace is observed by a single Intel RealSense D455 exocentric RGB-D camera, oriented to capture both the tabletop scene and the robot. The RGB data from this camera are used for learning symbolic models, while the depth information supports object pose estimation. Object poses are estimated using FoundationPose (Wen et al., 2024), which leverages high-fidelity 3D scanned models of the target objects. System-level communication and coordination are implemented in ROS 2 (Humble), which interfaces with motion planning, perception, and control modules.

This setup supports five parameterized skills: *Pick*, *Place*, *Stack*, *Pour*, and *Wipe*. The first four skills (*Pick*, *Place*, *Stack*, and *Pour*) are executed through motion planning with the MoveIt framework, conditioned on both the end-effector and object poses. The *Wipe* skill is implemented by replaying a teleoperated trajectory.

Bimanual Manipulation We use a robot with two horizontally mounted KUKA LBR iiwa 7 R800 manipulators, one with a BarrettHand BH8-282 gripper, and the other with a Schunk Dextrous Hand 2.0 gripper. The robot collects RGB data used for learning symbolic models with a MultiSense S7 camera mounted on a Pan-Tilt unit, while using an Intel RealSense D455 camera for RGB-D data used in pose estimation (Wen et al., 2024) of the objects in the scene. We use ROS 1 and KUKA FRI to communicate with the robot and utilize the built-in joint impedance control with position target as the low-level controller. At the high level, we create collision models of all objects in the scene and use a task and motion planner to generate motion plans for each skill. The *Pick* skills (compatible withknife and peanut butter jar) are implemented using motion planning. The *OpenJar*, *Scoop*, *Spread* skills are implemented using a combination of motion planning and pre-defined trajectory playback.

F.2 LANGUAGE MODEL PROMPTS

In this section, we provide the prompts used for the core components of SKILLWRAPPER (specifically skill sequence proposal, predicate invention, and predicate evaluation) as well as the ViLA (Hu et al., 2023) baseline. For predicate evaluation (Appendix F.5), we empirically observed that it is more accurate when evaluation is done in batches, where the truth values of multiple predicates are evaluated at once rather than one at a time. In addition, when asking for a fixed and structured output, the accuracy is significantly lower than a free-form output. Therefore, we adopt a two-stage evaluation process: in the first stage, the foundation model generates a response in any format, and in the second stage, it provides a summary of the output from the previous step.

F.3 SKILL SEQUENCE PROPOSAL

System Prompt

<AGENT_DESCRIPTION> is attempting to learn the preconditions and
effects for a finite set of skills by executing exploratory skill
sequences and exploring the environment.

Skill Sequence Proposal Prompt Propose a set of skill sequences for a robot to execute. The robot is attempting to learn the preconditions and effects for a finite set of operators. The robot can navigate the environment freely but only has one gripper. The robot has access to the following skills with their associated arguments: [SKILL_PROMPT] The list of objects the robot has previously encountered in the environment are: [OBJECT_IN_SCENE] [ENV_DESCRIPTION] The pairs of consecutive skills (skill1, skill2) that have been least explored are: [[LEAST_EXPLORED_SKILLS]]. Certain skills have similar names and arguments, but different preconditions and effects. Using the list of objects and the skill preconditions / effects learned, generate 5 skill sequences and their sequence of skills such that: (1) the skill sequences should violate their preconditions occasionally. (2) at least 1 unexplored skill pair is used in each skill sequence. (3) all skill sequences have at least 15 skills in sequence. (4) there are no same skills with same arguments consecutively in the sequence. Output only the sequence of skills to execute, ensuring to follow the naming/syntax/arguments for skills provided. Output 1 skill every new line, following the format below: Skill Sequence 1: GoTo(CounterTop) PickUp(Apple, CounterTop) Skill Sequence 2:

F.4 PREDICATE INVENTION

Predicate Invention Prompt

[AGENT_DESCRIPTION]

The robot has been programmed with the skill [LIFTED_SKILL] two times. In the first execution, the grounded skill [GROUNDED_SKILL_1] [SUCCESS_1], and in the second execution, [GROUNDED_SKILL_2] [SUCCESS_2]. The difference in outcomes suggests that the existing predicate set is insufficient to fully capture the preconditions for successful execution of this skill.

Your task is to propose a single new high-level predicate and its semantic meaning based on the visual comparison of the two input images taken before each execution.

Predicates should meet these criteria:

- The predicate must be grounded in visual state only (e.g., "gripper is open," "object is above table," "arm is holding object").
- Describe object state or spatial relations relevant to task success (e.g., gripper open/closed, object on left/right of gripper, object touching/supporting another object, etc.)
- Do not infer properties like affordances (is_graspable), alignment with grippers, or success likelihood that are vaguely defined and cannot be clearly determined visually.
- ${\mathord{\text{--}}}$ Avoid using concept like grasping zone or robot's reachability to define the predicate since they are not defined by common sense.
- Use at most 2 parameters (e.g., predicate(x), predicate(x, y), predicate()), where robot arm must be included for any robot-environment relation.
- Avoid predicates that assume internal properties like is_graspable, is_properly_aligned, or any accessibility/reachability reasoning that cannot be determined visually.
- The semantic meaning should be a grounded and objective description of the predicate in terms of the physical scene (e.g., "the object is fully enclosed by the robot's gripper"), not about execution success or skill dynamics.
- The parameters of the predicate must be subset of the parameters of the skill.

Format your output as follows:

'predicate_name (parameters) ': semantic_meaning.

for example:

'CloseTo(arm, location)': the robot arm is close to the location.

Current predicates: [PRED_LIST]

Previously proposed but rejected predicates: [TRIED_PRED]

Avoid duplicates or near-duplicates of existing predicates and rejected predicates. Reason over using a paragraph and generate the predicate and the semantic meaning in the given format in a separate line.

One new predicate candidate for improving the representation of the precondition for [LIFTED_SKILL] (Don't use any parameter other than [PARAMETERS]):

F.5 PREDICATE EVALUATION

Predicate Evaluation Prompt: Step 1 Given the current observation of the simulated kitchen domain, the object types, and the list of predicates, what are the true grounded predicates? [ENVIRONMENTAL_DESCRIPTION] Objects: [OBJECTS] Predicates: [PREDICATES]

Predicate Evaluation Prompt: Step 2

Summarize what the true grounded predicates are from this response, and list them in the format of predicate_name(arg1, arg2, ...) in separate lines with no any formatting. If the response contains typos of object names or redundant indices, you should correct them. Correct object names are: [OBJECT_NAMES]. If the response include redundant predicates that are not in this list, you should filter them. Correct predicates are: [PRED_NAMES]. The response is:

[RESPONSE]

F.6 VILA (HU ET AL., 2023)

ViLA Prompt

```
You are [AGENT_DESCRIPTION]. As a robot, you are able to execute the following skills: [SKILLS]
```

Here are the objects and their types that are compatible with your skills:

[OBJECTS]

You are given two images: The first one captures your current observation, and the second one specifies your goal. Given both images, your job is to generate a plan starting from the *current state* to the goal state. You should first reason about the goal of the task and how the skills can be chained to solve it in the first paragraph. After the reasoning, return the plan from the current state in a new paragraph by listing skills in separate lines with no additional explanation, header, or numbering. Use "Done" in the skill list to indicate the task is complete, and report if the task is impossible to solve by simply returning "Impossible".

G RELATED WORKS

Skill Abstraction There has been a long track of works focusing on building hierarchies that abstract away high-dimensional details with low-dimensional abstractions for planning (Konidaris & Barto, 2009; Konidaris et al., 2018; Shah et al., 2024), and those applied to robotics are usually

 connected to task and motion planning (TAMP) (Shah et al., 2020; Garrett et al., 2021). These approaches, however, are incapable of handling high-dimensional sensory-motor signals (such as images) as input. Research on action model learning (Xi et al., 2024; Juba et al., 2021) learn symbolic action models for input skills. However, unlike our method, these approaches require symbols to be provided as input. Similar to our system's integration of self-play and focus on uncovering skill conditions, Verma et al. (2022) focus on assessing capabilities of black-box agents for grid world-like tasks while assuming that the agent is an oracle. A tangential research effort on chaining various skills in novel environments involves training extra models (Yokoyama et al., 2024) and STRIPS task planner with action primitives (Gu et al., 2022; Szot et al., 2021).

Predicate Learning for Robotic Tasks Predicates provide a convenient way to abstract away low-level details of the environment and build efficient and compact representations. Prior to foundation models, previous attempts to build classifiers for predicates from raw image inputs originated from the neuro-symbolic domain (Johnson et al., 2017; Mao et al., 2019), and their initial application for robotics took a similar supervised learning approach with labeled demonstrations (Migimatsu & Bohg, 2022) or generated tasks (Lamanna et al., 2023). After the emergence of foundation models, recent works guide skill learning with predicates generated by LLM or together with human interaction. Li et al. (2024) invents symbolic skills for reward functions used for RL training but cannot generalize to skills learned through latent objectives, which is more commonly seen in imitation learning. Li & Silver (2023) and Han et al. (2024) leverage human experts to provide feedback to the LLM to help it improve the learned predicates and skills.

Task Generation for Robotics The approach of automatically proposing tasks has been studied for active learning and curriculum learning in grid worlds and games (Wang et al., 2019; Jiang et al., 2021) to robotic domains (Fang et al., 2021; 2022). Lamanna et al. (2023) generates tasks in PDDL as training sets to learn classifiers for object properties in predicates format, while they assume the action operators are given. With the commonsense reasoning ability of foundation models, recent works have applied the idea of automatic task proposing and self-playing for exploration (Nasiriany et al., 2024; Ren et al., 2024), data collection (Wang et al., 2024c; Yang et al., 2024; Ahn et al., 2024), boosting skills learning (Ha et al., 2023; Wang et al., 2024a), and scene understanding (Jiang et al., 2025). These works indicate a promising direction for generating robotic data and scaling up. Following the idea, we equipped our system with a task-proposing module for generating skill sequences specific to skills and predicates, which serves the idea of both data collection and exploration.

Embodied Reasoning with Foundation Models There has been a track of work on leveraging large language models (LLMs) for embodied decision-making (Huang et al., 2022; Raman et al., 2024) and reasoning (Huang et al., 2023), while vision-language models (VLMs) are often considered to have limited embodied reasoning ability due to their pre-training corpora that focus primarily on language generation (Valmeekam et al., 2023). Common ways of addressing this issue include fine-tuning on datasets from a specific domains (Hong et al., 2023; Mu et al., 2024; Chen et al., 2024) or knowledge distillation (Sumers et al., 2023; Yang et al., 2024). Meanwhile, many works manage to leverage preexisting models without further training from direct visual observation (Fang et al., 2024; Nasiriany et al., 2024) to complete robotic tasks (Jiang et al., 2025). In these works, the embodied reasoning ability of the foundation models serves as the central part of the systems. However, most benchmarking works evaluate the embodied reasoning ability of the models in a question-answering fashion (Sermanet et al., 2024; Majumdar et al., 2024; Cheng et al., 2024; Chen et al., 2025), where it remains unclear whether they are capable of solving robotic tasks.

H USE OF LARGE LANGUAGE MODELS

Our work incorporates language models as part of SKILLWRAPPER, particularly for the three important components of our system discussed in Section 3. We utilize OpenAI's GPT-5 (OpenAI, 2025) as our foundation model of choice. We acknowledge that all the content of the manuscript has been generated by the authors. However, we have used LLMs for basic editing, polishing, and grammar checking.