

LEAGUE: Guided Skill Learning and Abstraction for Long-Horizon Manipulation

Shuo Cheng  and Danfei Xu 

Abstract—To assist with everyday human activities, robots must solve complex long-horizon tasks and generalize to new settings. Recent deep reinforcement learning (RL) methods show promise in fully autonomous learning, but they struggle to reach long-term goals in large environments. On the other hand, Task and Motion Planning (TAMP) approaches excel at solving and generalizing across long-horizon tasks, thanks to their powerful state and action abstractions. But they assume predefined skill sets, which limits their real-world applications. In this work, we combine the benefits of these two paradigms and propose an integrated task planning and skill learning framework named LEAGUE (Learning and Abstraction with Guidance). LEAGUE leverages the symbolic interface of a task planner to guide RL-based skill learning and creates abstract state space to enable skill reuse. More importantly, LEAGUE learns manipulation skills in-situ of the task planning system, continuously growing its capability and the set of tasks that it can solve. We evaluate LEAGUE on four challenging simulated task domains and show that LEAGUE outperforms baselines by large margins. We also show that the learned skills can be reused to accelerate learning in new tasks domains and transfer to a physical robot platform.

Index Terms—Reinforcement learning, task and motion planning, continual learning.

I. INTRODUCTION

DEVELOPING robots that can autonomously learn to work in everyday human environments, such as households, has been a long-standing challenge. Deep Reinforcement Learning (DRL) methods have shown promise in allowing robots to acquire skills with limited supervision [9], [16], but they are still far from enabling home robots on their own. Two significant challenges stand out: 1) real-world tasks are often long-horizon, requiring the learning agent to explore a vast space of possible action sequences that grows exponentially with the task duration, and 2) home robots must perform diverse tasks in varying environments, requiring the learner to either generalize or quickly adapt to new situations.

To better learn long-horizon tasks, many DRL methods propose to use domain knowledge and structural prior [2], [22], [28]. Automatic goal generation in curriculum learning guides a learning process using intermediate subgoals, enabling an agent

Manuscript received 23 June 2023; accepted 7 August 2023. Date of publication 23 August 2023; date of current version 29 August 2023. This letter was recommended for publication by Associate Editor R. Lioutikov and Editor J. Kober upon evaluation of the reviewers' comments. (Corresponding author: Shuo Cheng.)

The authors are with the Georgia Institute of Technology, Atlanta, GA 30308 USA (e-mail: shuocheng@gatech.edu; danfei@gatech.edu).

This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2023.3308061>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2023.3308061

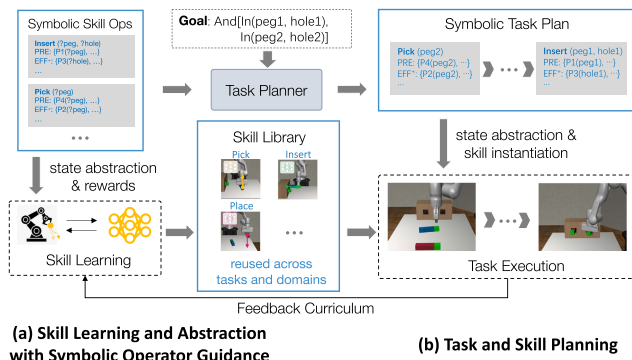


Fig. 1. Overview of the LEAGUE framework. We present an integrated task planning and skill learning framework. (a) The system uses the symbolic operator interface of a TAMP-like system as guidance to learn reusable manipulation skills (Algorithm 1). (b) A task planner composes the learned skills to solve long-horizon tasks (Algorithm 2). As an integrated system, the task planner acts as a feedback curriculum (bottom) to guide skill learning, and the RL-based skill learner continuously grows the set of tasks that the system can solve.

to explore and make incremental progress toward a long-horizon goal [22]. Other methods use skill primitives or learn hierarchical policies to enable temporally-extended decision-making [2], [19]. Although these approaches can outperform vanilla DRL, they still suffer from low sample efficiency, lack of interpretability, and fragile generalization [2], [28]. Most importantly, the learned policies are often task-specific and fall short in cross-task and cross-domain generalization.

In the meantime, more established paradigms in robotics have long sought to address these challenges. In particular, Task and Motion Planning (TAMP) [8], [14] leverages symbolic action abstractions to enable tractable planning and strong generalization. Specifically, the symbolic action operators divide a large planning problem into pieces that are each easier to solve. And the “lifted” action abstraction allows skill reuse across tasks and even domains. For example, a grasp skill operator and its underlying implementation can be easily adapted to solve a new task in a new domain. At the same time, most TAMP-style approaches assume access to a complete set of skills before deployment. This is impractical for two reasons. First, it is hard to prepare skills for all possible tasks. A robot must be able to grow its skill set on demand. Second, it is hard to hand-engineer manipulation skills for complex or contact-rich tasks (e.g., insertion). The challenges make TAMP methods difficult to deploy in real-world settings.

In this work, we introduce LEAGUE (LEarning and Abstraction with GUIDANCE), an integrated task planning and skill learning framework that learns to solve and generalize across long-horizon tasks (See Fig. 1). LEAGUE harnesses the merits of the two research paradigms discussed above. Starting

with a task planner that is equipped with skills that are easy to implement (e.g., *reaching*), LEAGUE continuously grows the skill set in-situ using a DRL-based learner. The intermediate goals in a task plan are prescribed as rewards for the learner to acquire and refine skills, and the mastered skills are used to reach the initial states of the new skills. Moreover, LEAGUE leverages the action operator definition, i.e., the preconditions and the effects, to determine a reduced state space for each learned skill, akin to the concept of information hiding in feudal learning [28]. The key idea is to *abstract away* task-irrelevant features to make the learned skills modular and reusable. Together, the result is a *virtuous cycle* where the task planner guides skill learning and abstraction, and the learner continuously expands the set of tasks that the system can solve.

We conduct empirical studies on four challenging long-horizon manipulation tasks built on the Robosuite simulation framework [31]. We show that LEAGUE is able to outperform state-of-the-art hierarchical reinforcement learning methods [19] by a large margin. We also highlight that our method can achieve strong generalization to new task goals and even domains by reusing and adapting learned skills. As a result, LEAGUE can solve a challenging simulated coffee-making task where competitive baselines fall flat. We also demonstrate a LEAGUE system trained in simulation on a physical Franka Emika Panda robot. In summary, our primary contributions are: 1) we leverage the state and action abstractions readily available in a TAMP system to learn reusable skills, 2) we instantiate the synergies between the task planner and the skill learner as an integrated task planning and skill learning framework, and 3) we show that the framework can progressively learn skills to solve complex long-horizon tasks and generalize the learned skills to new task goals and domains.

II. RELATED WORK

TAMP and Learning for TAMP: Task and Motion Planning (TAMP) [8], [14] is a powerful paradigm to solve long-horizon manipulation tasks. The key idea is to break a challenging planning problem into a set of symbolic-continuous search problems that are individually easier to solve. However, TAMP methods require high-level skills and their kinematics or dynamics models *a priori*. The assumptions preclude domains for which hand-engineering manipulation skills is difficult, such as contact-rich tasks. Recent works proposed to learn dynamics models for TAMP by characterizing skill preconditions and effects [15], [17], [24]. For example, Konidaris et al. [15] learns compact symbolic models of an environment through trial-and-error. Liang et al. [17] uses graph neural networks to model skill effects. However, these works still require hand-engineering complete skill sets that can solve the target task, which may not be feasible in real-world applications. Our idea of learning skills to augment TAMP systems is closely related to Silver et al. [24], which proposed to learn neural-symbolic skills via imitation. But they require access to hard-coded demonstration policies that can readily solve the target tasks. Our work instead aims to progressively grow TAMP skill libraries via guided reinforcement learning to solve long-horizon contact-rich manipulation tasks.

Curriculum for RL: Our idea to guide skill learning with a task planner is connected to curriculum-based RL, which is to expose an agent to incrementally more difficult intermediate tasks before mastering a target task [18]. The intermediate tasks can take the form of environments [6] and subgoals [22], [27]. For example, VaPRL [22] starts with near-success initialization

and moves the initial states further away. While effective at accelerating task learning, existing curricula focus on teaching tasks or domain-specific policies. In contrast, our method leverages the symbolic abstraction of a task planner to learn a repertoire of modular and composable skills. We show that we can compose learned skills to achieve new goals and even transfer skills to new task domains.

State and Action Abstractions: State and action abstractions are crucial for learning tasks in a large environment [1]. State abstraction allows agents to focus on task-relevant features of the environment. Action abstraction enables temporally-extended decision-making for long-horizon tasks. There exists a large body of work on learning either or both types of abstractions [1], [3], [5], [13], [15], [30]. For example, Jonschkowski et al. [13] explores different representation learning objectives for effective state abstraction. Abel et al. [1] introduces a theory for value-preserving state-action abstraction. However, autonomously discovering suitable abstractions remains an open challenge. Our key insight is that a TAMP framework provides powerful state and action abstractions that can readily guide skill learning. Specifically, the symbolic interface of an action operator defines both the precondition and the effect (action abstraction) and the state subspace that is relevant to the action (state abstraction). The abstractions allow us to train skills that are compatible with the task planner and prevent the learned skills from being distracted by irrelevant objects, making skill reuse across tasks and domains possible.

Hierarchical Modeling in Robot Learning: Our method inherits the bi-level hierarchy of a TAMP framework. Hierarchical modeling has a rich history in robotics. In addition to TAMP, various general frameworks including hierarchical task networks [11], [20], [29], logical-geometric programming [26], and hierarchical reinforcement learning (HRL) [2], [28] have been proposed to exploit the hierarchical nature of common robotics tasks. In the context of HRL, a small number of works have explored symbolic planner-guided HRL [12]. However, these methods require tabular state representations and are thus limited to simple grid-world domains. In robotics domains, a closely related research thread is to use behavior primitives in RL [4], [19]. For example, MAPLE [19] trains a high-level policy that chooses hand-engineered behavior primitives and atomic actions. Our method instead leverages a symbolic planner to serve as the high-level controller to compose learned skills, allowing us to continuously extend the skill set while also leading to better generalization.

III. METHOD

We seek to enable robots to solve and generalize across long-horizon tasks. Our primary contribution is a novel integrated task planning and skill learning framework named LEAGUE. Here, we first provide the necessary background in Section III-A, and describe how LEAGUE (1) learns reusable skills guided by the symbolic operators of a task planner in Section III-B and (2) uses planner-generated task plans as an autonomous curriculum to continuously learn skills and expand the capability of the overall system in Section III-C.

A. Background

MDP: We consider a Markov Decision Process (MDP) $\langle \mathcal{X}, \mathcal{A}, \mathcal{R}(x, a), \mathcal{T}(x'|x, a), p(x^{(0)}), \gamma \rangle$, with continuous state space \mathcal{X} , continuous action space \mathcal{A} , reward function \mathcal{R} ,

and environment transition model \mathcal{T} . $p(x^{(0)})$ denotes the distribution of the initial states, $x^{(H)}$ denotes terminal state, and γ is the discount factor. The objective for RL training is to maximize the expected total reward of the policy $\pi(a|x)$ that the agent uses to interact with the environment: $J = \mathbb{E}_{x^{(0)}, a^{(0)}, \dots, x^{(H)} \sim \pi, p(x^{(0)})} [\sum_t \gamma^t \mathcal{R}(x^{(t)}, a^{(t)})]$.

Task planning space: To support task planning, we assume the environment is augmented with a symbolic interface $\langle \mathcal{O}, \Lambda, \bar{\Psi}, \bar{\Omega}, \mathcal{G} \rangle$, where \mathcal{O} denotes the object set and Λ denotes a finite set of object types. Each object entity $o \in \mathcal{O}$ (e.g., `peg1`) has a specific type $\lambda \in \Lambda$ (e.g., `peg`) and a tuple of $\dim(\lambda)$ -dimensional feature containing information such as poses and joint angles, and the environment state $x \in \mathcal{X}$ is a mapping from object entities to features: $x(o) \in \mathbb{R}^{\dim(\text{type}(o))}$. Predicates $\bar{\Psi}$ describe the relationships among multiple objects. Each predicate $\bar{\psi}$ (e.g., `holding(?object:peg)`) is characterized by a tuple of object types $(\lambda_1, \dots, \lambda_m)$ and a binary classifier that determines whether the relationship holds: $c_{\bar{\psi}} : \mathcal{X} \times \mathcal{O}^m \rightarrow \{True, False\}$, where each substitute entity $o_i \in \mathcal{O}$ is restricted to have type $\lambda_i \in \Lambda$. Evaluating a predicate on the state by substituting corresponding object entities will result in a ground atom (e.g., `holding(peg1)`). A task goal $g \in \mathcal{G}$ is represented as a set of ground atoms, where a symbolic state $x_{\bar{\Psi}}$ can be obtained by evaluating a set of predicates $\bar{\Psi}$ and keeping all positive ground atoms:

$$x_{\bar{\Psi}} = \text{PARSE}(x, \mathcal{O}, \bar{\Psi}) \\ \triangleq \{\underline{\psi} : c_{\bar{\psi}}(x, \mathcal{O}^{\bar{\psi}}) = True, \forall \mathcal{O}^{\bar{\psi}} \subseteq \mathcal{O}, \forall \bar{\psi} \in \bar{\Psi}\} \quad (1)$$

where $\mathcal{O}^{\bar{\psi}}$ is a subset of object entities that each entity o_i has the same object type λ_i specified by the predicate $\bar{\psi}$.

Symbolic skill operators: Following prior works [8], we characterize lifted skill operator $\bar{\omega} \in \bar{\Omega}$ by a tuple $\langle \text{PAR}, \text{PRE}, \text{EFF}^+, \text{EFF}^- \rangle$, where `PRE` denotes the precondition of the operator, which is a set of lifted atoms defining the condition that the operator is executable. `EFF+` and `EFF-` are lifted atoms that describe the expected effects (changes in conditions) upon successful skill execution. `PAR` is an ordered parameter list that defines all object types used in `PRE`, `EFF+`, and `EFF-`. A ground skill operator $\underline{\omega}$ substitutes lifted atoms with object instances: $\underline{\omega} = \langle \bar{\omega}, \delta \rangle \triangleq \langle \text{PRE}, \text{EFF}^+, \text{EFF}^- \rangle$, where $\delta : \Lambda \rightarrow \mathcal{O}$. Given a task goal, a symbolic task plan is a list of ground operators that, when the instantiated skills are executed successfully, lead to an environment state that satisfies the goal condition.

As a running example, consider a short task of inserting a peg (`peg1`) into the target hole (`hole1`). The applicable operators for this task are defined as:

```
Pick(?object)
  PAR: [?object:peg]
  PRE: {HandEmpty(), OnTable(?object)}
  EFF-: {HandEmpty(), OnTable(?object)}
Insert(?object, ?hole)
  PAR: [?object:peg, ?hole:hole]
  PRE: {Holding(?object), IsClear(?hole)}
  EFF-: {Holding(?object), IsClear(?hole)}
  EFF+: {HandEmpty(), In(?object, ?hole)}
```

The environment starts with `peg1` on the table. Evaluating the `PARSE` function (1) yields a symbolic state $\{\text{HandEmpty}(), \text{IsClear}(\text{hole1}), \text{OnTable}(\text{peg1})\}$, a set of grounded atoms that satisfies the preconditions of the grounded operator `Pick(peg1)`. This grounded operator, if executed successfully, should reach the symbolic state of $\{\text{Holding}(\text{peg1}), \text{IsClear}(\text{hole1})\}$, which is an intermediate subgoal for the final task goal that is characterized by the grounded atoms $\{\text{HandEmpty}(), \text{In}(\text{peg1}, \text{hole1})\}$. The symbolic task plan is therefore $\mathcal{P} = [\text{Pick}(\text{peg1}), \text{Insert}(\text{peg1}, \text{hole1})]$.

We are interested in learning primitive manipulation skills for accomplishing individual subgoals induced by the expected effects of the corresponding operators – the building blocks that constitute a symbolic task plan. In our setting, each lifted operator $\bar{\omega}$ will have a corresponding skill policy π to be learned, while during execution the ground operators belong to the same lifted operator $\bar{\omega}$ share the same skill policy. We assume access to the predicates $\bar{\Psi}$ and the lifted operators $\bar{\Omega}$ of the environments and focus on efficiently learning the skills for achieving the effects. Note that it is possible to invent and learn predicates and operators [23], [25], but the topics are beyond the scope of this work.

B. Skill Learning and Abstraction With Operator Guidance

Action and state abstractions are fundamental to TAMP's abilities to solve and generalize across long-horizon tasks [8]. Our key insight is that these abstractions, in the form of symbolic action operators, can readily guide RL-trained policies to gain similar abilities. Specifically, for action abstraction, we train temporally-extended skills to reach desired effects of a skill operator by prescribing the effect condition as shaped reward. For state abstraction, we take inspiration from the idea of *information hiding* in feudal learning [28] and use the precondition and effect signature of an operator to determine a *skill-relevant* state space for its corresponding learned policy. This allows the policy to be robust against domains shift and achieve generalization, especially in large environments where most elements are impertinent to a given skill. To further accelerate skill learning, we leverage the existing motion planning capability of a TAMP system to augment the learned skill with a transition primitive. Below we describe each component in detail.

Symbolic operators as reward guidance: Our skill learner leverages the existing RL method that supports continuous action space. In this work, we use Soft Actor-Critic (SAC) [10] as the basis for skill learning. SAC leverages entropy regularization to enhance exploration. Given the ground operator $\underline{\omega}$ of a skill, we can define an operator-guided reward $\mathcal{R}_{\bar{\Psi}}$ for each individual skill based on continuous environment state x and the action a produced by the corresponding policy π that takes in skill-related state \hat{x} (which will be described later), the objective for our skill learning is therefore rewritten as:

$$J = \mathbb{E}_{x^{(0)}, a^{(0)}, \dots, x^{(H)} \sim \pi, p(x^{(0)})} \left[\sum_t \gamma^t (\mathcal{R}_{\bar{\Psi}}(x^{(t)}, a^{(t)}, \underline{\omega}) + \alpha \mathcal{H}(\pi(\cdot | \hat{x}^{(t)}))) \right] \quad (2)$$

where $R_{\bar{\Psi}}(\cdot) \mapsto [0, 1]$, $\mathcal{H}(\cdot)$ is the entropy term introduced by SAC. While it is possible to learn directly from sparse reward defined by the symbolic state, in practice we associate each operator-guided reward with a dense reward function implemented in the Robosuite [31] benchmark for better learning

efficiency. Continuing our running example, the shared reward for `Pick` is defined as $1 - \tanh(10.0 * d)$, where d is the distance between the gripper center and target object center, and the target object is identified by the task planner.

Enhance skill reuse with feudal state abstraction: With the precondition and effect signature of a ground operator ω , we can determine a skill-relevant state space to further prevent the learned policy from being distracted by task-irrelevant objects:

$$\hat{x} = \text{EXTRACT}(x, \underline{\omega}, \mathcal{O}) \triangleq \{x(o) : o \in \underline{\text{PAR}}, \forall o \in \mathcal{O}\} \quad (3)$$

where $\underline{\text{PAR}}$ is the parameter list of the ground operator. In our running example, the skill-related state \hat{x} for `Pick(peg1)` includes the 6D pose of `peg1` and the state of the robot. This design echoes previous works that learn to impose constraints on states [3], except that here the constraints are directly informed by the task planner.

Accelerate learning with transition motion primitives: A key to our method is learning modular manipulation skills that can be composed to solve long tasks. However, for complex manipulation problems, even learning such short skills can be challenging. On the other hand, although TAMP systems fall short when facing contact-rich manipulation, they excel at finding collision-free paths. To this end, we propose to augment our policy with motion planner-based transition primitives. The key idea is to first approach the skill-relevant object (per the skill operator) using an off-the-shelf motion planner, before convening RL-based skill learning. For the target of motion planning, we simply set the goal position to be 0.04 m higher than the object or placement position of interest that was identified by the task planning. The component significantly speeds up the exploration while still allowing the system to learn closed-loop contact-rich manipulation skills.

C. Integrated Task Planning and Skill Learning

So far, we have described a recipe for learning reusable skills using symbolic skill operators as guidance. But these skills are not learned in silos. A key to LEAGUE's success is to learn skills in-situ of a task planning system. The integrated planning and learning scheme ensures that the learned skills are compatible with the planner, and the skill learner can continuously extend the capability of the overall system to solve more tasks. Here we first describe how LEAGUE performs task planning and execution at inference time, and then we introduce an algorithm that uses task plans as an autonomous curriculum to schedule skill learning.

Task planning and skill execution: To plan for task goal g , we first `PARSE` (1) the continuous environment state x for obtaining symbolic state x_Ψ , which affords symbolic search with ground operators. We then ground each lifted operator $\bar{\omega} \in \bar{\Omega}$ on the object set \mathcal{O} by substituting object entities in preconditions and effects, leading to ground operators $\underline{\omega} = \langle \underline{\text{PRE}}, \underline{\text{EFF}}^+, \underline{\text{EFF}}^- \rangle$ that support operating with symbolic states. A ground operator is considered executable only when its preconditions are satisfied: $\underline{\text{PRE}} \subseteq x_\Psi$. The operators induce an abstract transition model $F(x_\Psi, \underline{\omega})$ that allows planning in symbolic space:

$$x'_\Psi = F(x_\Psi, \underline{\omega}) \triangleq (x_\Psi \setminus \underline{\text{EFF}}^-) \cup \underline{\text{EFF}}^+ \quad (4)$$

We use PDDL [7] to build the symbolic planner and use A* search for generating high-level plans.

With the generated task plan, we sequentially invoke the corresponding skill π^* to reach the subgoal that complies with

Algorithm 1: SKILLCURRICULUM.

hyperparameters:

Number of training iterations K

input:

env ▷ task environment
 g ▷ symbolic task goal
 $\bar{\Psi}$ ▷ state predicates
 $\bar{\Omega}$ ▷ lifted operators

start

$\Pi \leftarrow [\pi_1^{(0)}, \dots, \pi_{|\bar{\Omega}|}^{(0)}]$ ▷ initialize all skill policies

$t \leftarrow 0$

while Not Converged do

$\mathcal{D} \leftarrow \emptyset$ ▷ buffer for failed skills

for $i \leftarrow [1, \dots, N]$ **do**

$\mathcal{D} \leftarrow \mathcal{D} \cup \text{TRYSOLVETASK}(\text{env}, g, \bar{\Psi}, \bar{\Omega}, \Pi)$

end for

for $i, s, \underline{\omega} \leftarrow \mathcal{D}$ **do**

$\pi_i^{(t)} \leftarrow \Pi[i]$

for $k \leftarrow [1, \dots, K]$ **do**

$\pi_i^{(t+k)} \leftarrow \text{SAC}(\text{env}, s, \pi_i^{(t+k-1)}, \underline{\omega})$ ▷ RL training

end for

$\Pi[i] \leftarrow \pi_i^{(t+K)}$

end for

$t \leftarrow t + K$

end while

return Π

the effects of each ground operator $\underline{\omega}$ in the plan. We rollout each skill controller until it fulfills the effects of the operator or a maximum skill horizon H is reached. To verify whether the l -th skill is executed successfully, we first obtain the corresponding symbolic state x'_Ψ by parsing the ending environment state x^* . The execution is considered successful only when the environment state x^* conforms to the expected effects: $F(x'_\Psi, \underline{\omega}_l) \subseteq x'_\Psi$. We keep track of the failed skills and the starting simulator info s^* to inform the learning curriculum.

Task planner as an automated curriculum: To efficiently acquire all necessary skills for a given multi-step task, we leverage the task planner as an automated curriculum to learn skills in a progressive manner. The key idea is to use more proficient skills to reach the preconditions of skills that require additional learning (See Fig. 1). The algorithm is sketched in Algorithms 1 and 2. On a high level, we repeat task planning and skill learning until convergence. We keep track of failed skills during N task executions and adopt strict scheduling criteria, where a skill is scheduled for learning (Section III-B) if it ever fails during the N episodes. Notably, we share the replay buffers for different skill instances (e.g., `Pick(peg1)` and `Pick(peg2)`) that belong to the same lifted operator, so that the relevant experience can be reused to further improve the learning efficiency and generalization.

IV. EXPERIMENTS

Our experiments aim to show that 1) LEAGUE can progressively learn and refine skills to solve long-horizon tasks and 2) our novel operator-guided skill learning and abstraction algorithm produces composable and reusable skills, enabling quick adaptation to new tasks and domains. Finally, we demonstrate transferring a trained LEAGUE system to a physical robot.

Algorithm 2: TRY SOLVE TASK.**hyperparameters:**Maximal skill horizon H **input:**

env \triangleright task environment
 g \triangleright symbolic task goal
 $\bar{\Psi}$ \triangleright state predicates
 $\bar{\Omega}$ \triangleright lifted operators
 Π \triangleright skill policies

start $\mathcal{O}, x^{(0)} \leftarrow env.get_state()$ $x_{\Psi}^{(0)} \leftarrow PARSE(x^{(0)}, \mathcal{O}, \bar{\Psi})$ \triangleright continuous state to symbolic state $\underline{\Omega} \leftarrow GROUND(\mathcal{O}, \bar{\Omega})$ \triangleright get grounded operators $[\underline{\omega}_1, \dots, \underline{\omega}_L] \leftarrow SEARCH(x_{\Psi}^{(0)}, g, \underline{\Omega})$ \triangleright found plan with length L $\mathcal{D}, l \leftarrow [], 0$ **while** $l < L$ **do** $i \leftarrow LOOKUPSKILL(\underline{\omega}_l)$ $\pi^* \leftarrow \Pi[i]$ $s^* \leftarrow env.get_sim()$ \triangleright get simulator state $x^* \leftarrow ROLLOUT(env, \pi^*, H)$ **if** $ISUCCESS(x^*, \underline{\omega}_l)$ **then** $l \leftarrow l + 1$ \triangleright advance to the next skill**continue****else** $\mathcal{D} \leftarrow \mathcal{D} \cup (i, s^*, \underline{\omega}_l)$ \triangleright collect failed skills and states**end if****end while****return** \mathcal{D} **A. Experimental Setup**

We conduct evaluations in four simulated domains, in which we devise tasks that require multi-step reasoning, contact-rich manipulation, and long-horizon interactions (See Fig. 3).

StackAtTarget is to stack two cubes on a tight target region with a specific order. The applicable skill operators are `Pick` and `Place`. Since the cubes are randomly placed in the scene, the top cube may occupy the target region, in which situation the robot must first remove the top cube before stacking.

StowHammer requires the robot to stow two hammers into different closed cabinets. It involves four skills: `Pick`, `Place`, `Pull`, `Push`. Since the workspace is tight, the robot needs to close an opened cabinet before being able to open the other one, which requires multi-step reasoning.

PegInHole is to pick up and insert two pegs into two horizontal holes. The applicable operators are `Pick` and `Insert`. This task challenges the robot with contact-rich manipulations and multi-step planning.

MakeCoffee is to pick up a coffee pod from a closed cabinet, insert it into the holder of the coffee machine, and finally close both the lid and the cabinet. The applicable operators are `Pick`, `Pull`, `Push`, `CloseLid`, and `InsertHolder`.

The environments are built on Robosuite [31] simulator. We use a Franka Emika Panda robot arm that is controlled at 20 Hz with an operational space controller (OSC), which has 5 degrees of freedom: end-effector position and the yaw angle and the position of the gripper. See Fig. 3 for an illustration.

B. Visualize the Progressive Skill Learning Process

Before discussing quantitative comparisons, we seek to gain intuitive understanding of our progressive skill learning scheme (Section III-C), where the learning curriculum adjusts based on the proficiencies of the skills. In Fig. 2, we visualize the proficiency level of each skill throughout the process of learning a simplified StowHammer task, where the goal is to stow away one hammer instead of two. The y axis is the average normalized reward a skill receives. Note that we only visualize a subset of skills scheduled for training at an iteration. The corresponding behavior of each skill at a certain stage is visualized in the snapshots on top of the plot.

At the beginning of the training, the system can only reach the precondition for executing the `Pull(?cabinet)` skill but not other skills, thus the experience of `Pull(?cabinet)` skill is collected and it is repeatedly selected for training. Until the agent is able to open one of the cabinets, the second planned skill `Pick(?object)` is then instantiated for learning and execution. Finally at the end of the training, all skills become proficient to be used to execute the entire task. The result qualitatively shows that LEAGUE's automated curriculum is effective at progressively learning skills to achieve long-horizon task goals.

C. Quantitative Evaluation

Here, we seek to highlight various aspects of our solution paradigm through quantitatively comparing LEAGUE with a number of strong baseline methods. Below we describe the baselines and discuss the results.

- *RL (SAC)*: We adopt SAC [10] as a strong RL baseline. To facilitate a fair comparison, we extend the vanilla task reward function to staged rewards using an oracle task plan, where the reward at each step is the summation of achieved rewards for each completed subgoal and the reward for the current subgoal.
- *Curriculum RL (CRL)*: We follow the main idea of state-of-the-art curriculum RL approaches [22], [27], which starts the training with near-success initializations and gradually move the initial states back to the true environment initial states. To facilitate a fair comparison, we sample the curriculum's initial states based on the subgoals of an oracle task plan (in reverse) and adopt the same staged reward described above.
- *Hierarchical RL (HRL)*: This baseline adapts the recent primitive-based HRL frameworks [4], [19] for our tasks. The key idea is to train a high-level meta controller to compose parameterized skill primitives and atomic actions. We base our implementation on MAPLE [19] and use the oracle task plan to identify the target objects for defining the affordance to guide the exploration.
- *Symb+MP*: An open-loop baseline that resembles a vanilla TAMP framework, which greedily generates a motion plan for each skill in a task plan. The robot then executes the plan through a trajectory controller.
- *Symb+RL*: An ablation baseline of LEAGUE that removes the state abstraction (III-B) and retains all other features including the symbolic plan-based curriculum.

The multi-stage nature of our evaluation tasks makes designing smooth task-level metrics difficult. Thus we adopt **task progress** as our metric, which is defined as the summed reward

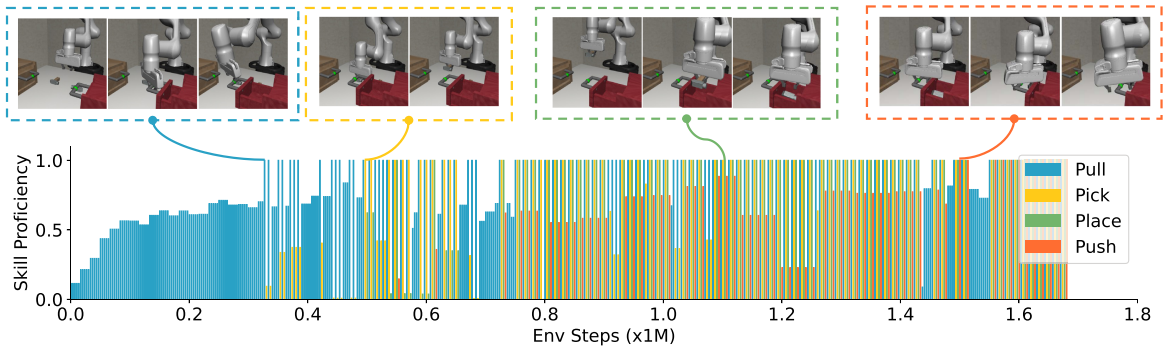


Fig. 2. Visualizing skill learning progress. The plot shows the proficiency level of each skill throughout the process of learning a simplified StowHammer task. The skill proficiency is the average normalized reward a skill receives at an iteration.

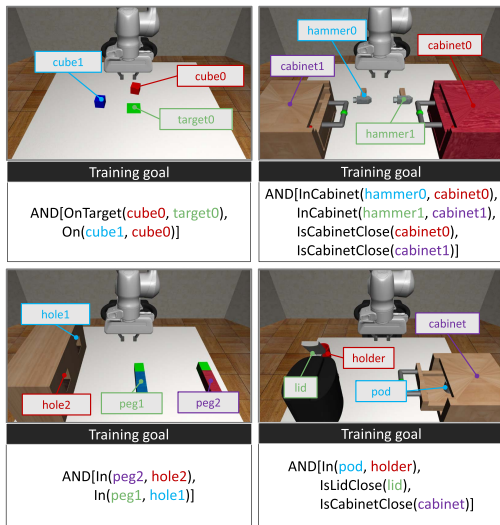


Fig. 3. Simulation task setup. We show the initialization of the simulation setups for the four tasks: StackAtTarget, StowHammer, PegInHole, and Make-Coffee.

of all task stages normalized to $[0, 1]$. Below we discuss the main findings based on Fig. 4.

High-level reasoning is critical for solving long-horizon tasks: We observe that in StackAtTarget, a long-horizon task with relatively simple manipulation steps, methods equipped with a task planner (LEAGUE, Symb+MP, and Symb+RL) significantly outperforms all other baselines. The most competitive HRL baseline occasionally learns to move the bottom cube to the target region. This shows the value of explicit high-level reasoning, in particular as a plan-informed automated curriculum in LEAGUE. Notably, the open loop Symb+MP performs on par with LEAGUE because simple picking and placing can readily be solved by open-loop trajectories.

LEAGUE can solve long-horizon, contact-rich manipulation tasks: LEAGUE significantly outperforms all other baselines in StowHammer and PegInHole, which are both long-horizon and require contact-rich manipulation. Notably, most baselines cannot advance beyond opening the cabinet in StowHammer and picking up the first peg in PegInHole.

Skill reuse is critical to learning structured tasks: Common multi-step tasks have repeating structures, which can be leveraged by methods that explicitly reuse learned skills. We

note that both LEAGUE and Symb+RL perform competitively in StackAtTarget that involve repeating steps (i.e., stack two cubes). On the other hand, HammerPlace and PegInHole involve more objects, most of which are not relevant to a given skill. This prevents naïve skill reuse — a policy may learn spurious correlation to these irrelevant features — and necessitates state abstraction, which we will discuss next.

State abstraction facilitates skill reuse in complex environments: We observe that LEAGUE outperforms Symb+RL in both HammerPlace and PegInHole. This shows that state abstraction can further improve skill reuse in complex environments by ignoring features that are irrelevant to a skill. We will also show in Section IV-D that skill reuse enables our method to generalize to novel task goals and domains.

Other observations: We observe that without explicit prior structures such as motion primitive, SAC baseline is able to exploit environment artifacts and learn shortcut behaviors. For example, in the StowHammer task, SAC agent learns to grip the head of the hammer to prevent slipping, but the grasping pose precludes it from fitting the hammer to the drawer. Moreover, our analysis found that the CRL agent often failed to reach the final goal from some intermediate states due to the strong sequential dependency of our evaluation tasks: the robot must succeed in one stage to reach the pre-condition of the next. And because the environment steps budget is distributed to multiple stages, CRL often underperforms other baselines (e.g., SAC) in completing the initial stages of a task.

D. Generalization to New Tasks and Domains

To validate that our method can effectively generalize to new task goals and even new task domains by reusing learned skills, we present the following experiments.

Generalize to new task goals: Besides evaluating the training goals (shown in Fig. 3), we directly test our models on new task goals for the StowHammer and the PegInHole domains. For StowHammer domain, the first test goal is to swap the hammer-cabinet mapping. The second test goal is to place hammer1 into cabinet0 and keep cabinet1 open. For PegInHole, the first test goal is to swap the peg-hole mapping. The second goal is to only insert peg1 into hole2. The results are in Table I. We observe that LEAGUE experiences little performance drop when generalizing to new task goals without additional training, demonstrating strong compositional generalization capability and skill modularity.




Fig. 4. Baseline comparison. We compare relevant methods on three task domains. The plot shows the corresponding average task progress during evaluation throughout training, which is measured as the summation of achieved rewards of each successfully executed skill in the task plan and normalized to 1. The results are reported using five random seeds, with the standard deviation shown as the shaded area.

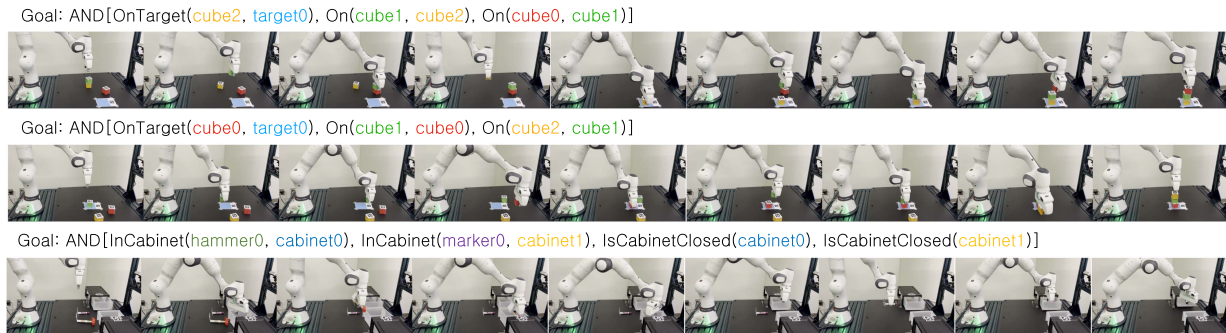


Fig. 5. Real robot demonstration. Key frames of three task execution processes (bottom) and their final task goals (top).

TABLE I

WE REPORT THE PERFORMANCE OF APPLYING OUR METHOD TO NEW TASK GOALS IN THE STOWHAMMER AND THE PEGINHOLE DOMAINS WITHOUT ADDITIONAL LEARNING

	Training Goal	Test Goal1	Test Goal2
StowHammer	0.94 ± 0.21	0.90 ± 0.12	0.73 ± 0.31
PegInHole	0.87 ± 0.23	0.53 ± 0.05	1.00 ± 0.00




Fig. 6. Generalization to new domain. For MakeCoffee task, we compare learning the task from scratch and learning by adapting the skills (Pick, Pull, and Push) learned from the StowHammer domain.

Quick adaptation to new domains: Another exciting possibility of LEAGUE is to transfer skills learned from one domain to another. We design an experiment to validate this feature. The target domain is MakeCoffee, which is the most challenging task of the four. We adapt skills Pick, Pull, and Push learned in the StowHammer domain for learning the MakeCoffee task. As shown in Fig. 6, compared to learning from scratch, transferring learned skills can significantly accelerate learning (the x -axis is shorter than in Fig. 4) and enables the robot to solve the entire

task. This highlights LEAGUE's strong potential for continual learning.

E. Real World Demonstration

We demonstrate transferring simulation-trained LEAGUE system to two real-world task domains: **StackThreeAtTarget** and **StowObject**. For the StackThreeAtTarget task, we randomly place three cubes and a target region on the table. The task is to stack the cubes at the target region. We directly reuse the skills trained in StackAtTarget in simulation to demonstrate generalization to different number of objects and initial conditions. The StowObject is to stow two objects into two cabins. Similar to StowHammer, the task also requires the robot to operate the cabinets. We reuse skills trained in the simulated StowHammer domain.

Our system uses a Franka Emika Panda robot. We take RGBD images from an Intel RealSense D435 camera and use AprilTag [21] to detect the 6D poses of task-relevant objects. Our system performs state estimation prior to each skill execution, synchronizes the states to a simulated environment, and executes each skill generated by LEAGUE from the simulated environment through open-loop control.

Fig. 5 shows the key frames of three task execution processes and the corresponding task goals. Our system achieves an 8/10 success rate for the StackThreeAtTarget task, and a 6/10 success rate for the StowObject task. The failure mode for the StackThreeAtTarget task is that the AprilTags getting occluded from the camera in some initial configurations. The failure mode for StowObject task is that sometimes the learned policy is not able to generate a valid motion for operating the drawer, and the objects slipping from the gripper.

V. CONCLUSIONS, LIMITATIONS, AND FUTURE WORKS

We introduced LEAGUE, an integrated task planning and skill learning framework that represents a virtuous-cycle system: It leverages the high-level reasoning ability and abstraction of a TAMP framework to facilitate the exploration and generalization of an RL skill learner, which in turn expands the capability of the overall system. Through challenging manipulation tasks in both simulation and the real world, we demonstrated that LEAGUE is effective at solving long-horizon tasks and generalizing to new tasks and domains.

While empirically effective, our method does have a number of limitations. As we discussed in Section III-A, we assume access to a library of skill operators that serve as the basis for skill learning. Relatedly, our assumptions for skill-relevant state abstraction, although effective, may not hold in certain cases (e.g. unintended consequences during exploration). A possible path to address both challenges is to learn skill operators with sparse transition models from unstructured experiences [23], [25]. Second, our skill learning process relies on the environment-provided dense reward function. RL algorithms that can better learn from sparse reward would allow LEAGUE to build a tighter connection with the symbolic space. Finally, in the real-world setting, LEAGUE is limited by the capability of the off-the-shelf perception algorithms. We plan to explore learning visuomotor control policies to make LEAGUE easier to deploy in the real world.

REFERENCES

- [1] D. Abel, N. Umbanhowar, K. Khetarpal, D. Arumugam, D. Precup, and M. Littman, "Value preserving state-action abstractions," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2020, pp. 1639–1650.
- [2] P.-L. Bacon, J. Harb, and D. Precup, "The option-critic architecture," in *Proc. AAAI Conf. Artif. Intell.*, 2017, pp. 1726–1734.
- [3] R. Chitnis, T. Silver, B. Kim, L. Kaelbling, and T. Lozano-Perez, "CAMPS: Learning context-specific abstractions for efficient planning in factored MDPs," in *Proc. Conf. Robot Learn.*, 2021, pp. 64–79.
- [4] M. Dalal, D. Pathak, and R. R. Salakhutdinov, "Accelerating robotic reinforcement learning via parameterized action primitives," in *Proc. Neural Inf. Process. Syst.*, 2021, pp. 21847–21859.
- [5] S. Emmons, A. Jain, M. Laskin, T. Kurutach, P. Abbeel, and D. Pathak, "Sparse graphical memory for robust planning," in *Proc. Neural Inf. Process. Syst.*, 2020, pp. 5251–5262.
- [6] K. Fang, Y. Zhu, S. Savarese, and L. Fei-Fei, "Adaptive procedural task generation for hard-exploration problems," in *Proc. Int. Conf. Learn. Representations*, 2021.
- [7] M. Fox and D. Long, "PDDL2.1: An extension to PDDL for expressing temporal planning domains," *J. Artif. Intell. Res.*, vol. 20, pp. 61–124, 2003.
- [8] C. R. Garrett et al., "Integrated task and motion planning," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 4, pp. 265–293, 2021.
- [9] S. Gu, E. Holly, T. Lillicrap, and S. Levine, "Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 3389–3396.
- [10] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1861–1870.
- [11] B. Hayes and B. Scassellati, "Autonomously constructing hierarchical task networks for planning and human-robot collaboration," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016, pp. 5469–5476.
- [12] L. Illanes, X. Yan, R. T. Icarte, and S. A. McIlraith, "Symbolic plans as high-level instructions for reinforcement learning," in *Proc. Int. Conf. Automated Plan. Scheduling*, 2020, pp. 540–550.
- [13] R. Jonschkowski and O. Brock, "Learning state representations with robotic priors," *Auton. Robots*, vol. 39, pp. 407–428, 2015.
- [14] L. P. Kaelbling and T. Lozano-Pérez, "Hierarchical task and motion planning in the now," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2011, pp. 1470–1477.
- [15] G. Konidaris, L. P. Kaelbling, and T. Lozano-Perez, "From skills to symbols: Learning symbolic representations for abstract high-level planning," *J. Artif. Intell. Res.*, vol. 61, pp. 215–289, 2018.
- [16] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *J. Mach. Learn. Res.*, vol. 17, pp. 1334–1373, 2016.
- [17] J. Liang, M. Sharma, A. LaGrassa, S. Vats, S. Saxena, and O. Kroemer, "Search-based task planning with learned skill effect models for lifelong robotic manipulation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2022, pp. 6351–6357.
- [18] S. Narvekar, B. Peng, M. Leonetti, J. Sinapov, M. E. Taylor, and P. Stone, "Curriculum learning for reinforcement learning domains: A framework and survey," *J. Mach. Learn. Res.*, vol. 21, pp. 7382–7431, 2020.
- [19] S. Nasiriany, H. Liu, and Y. Zhu, "Augmenting reinforcement learning with behavior primitives for diverse manipulation tasks," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2022, pp. 7477–7484.
- [20] N. Nejati, P. Langley, and T. Konik, "Learning hierarchical task networks by observation," in *Proc. Int. Conf. Mach. Learn.*, 2006, pp. 665–672.
- [21] E. Olson, "AprilTag: A robust and flexible visual fiducial system," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2011, pp. 3400–3407.
- [22] A. Sharma, A. Gupta, S. Levine, K. Hausman, and C. Finn, "Autonomous reinforcement learning via subgoal curricula," in *Proc. Neural Inf. Process. Syst.*, 2021, pp. 18474–18486.
- [23] T. Silver, R. Chitnis, J. Tenenbaum, L. P. Kaelbling, and T. Lozano-Pérez, "Learning symbolic operators for task and motion planning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2021, pp. 3182–3189.
- [24] T. Silver, A. Athalye, J. B. Tenenbaum, T. Lozano-Pérez, and L. P. Kaelbling, "Learning neuro-symbolic skills for bilevel planning," in *Proc. Conf. Robot Learn.*, 2022, pp. 701–714.
- [25] T. Silver et al., "Predicate invention for bilevel planning," in *Proc. AAAI Conf. Artif. Intell.*, 2023, pp. 12120–12129.
- [26] M. Toussaint, "Logic-geometric programming: An optimization-based approach to combined task and motion planning," in *Proc. Int. Joint Conf. Artif. Intell.*, 2015, pp. 1930–1936.
- [27] I. Uchendu et al., "Jump-start reinforcement learning," in *Proc. Int. Conf. Mach. Learn.*, 2022, pp. 34556–34583.
- [28] A. S. Vezhnevets et al., "Feudal networks for hierarchical reinforcement learning," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 3540–3549.
- [29] D. Xu et al., "Neural task programming: Learning to generalize across hierarchical tasks," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 3795–3802.
- [30] D. Xu, A. Mandlekar, R. Martín-Martín, Y. Zhu, S. Savarese, and L. Fei-Fei, "Deep affordance foresight: Planning through what can be done in the future," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 6206–6213.
- [31] Y. Zhu, J. Wong, A. Mandlekar, and R. Martín-Martín, "Robosuite: A modular simulation framework and benchmark for robot learning," 2020, *arXiv:2009.12293*.