SkiLD: Unsupervised Skill Discovery Guided by Local Dependencies

Zizhao Wang*, Jiaheng Hu*, Caleb Chuck*, Amy Zhang, Roberto Martin-Martin, Scott Niekum, Peter Stone

Code: Contact authors

Contact: calebc@cs.utexas.edu

Paper:

Overview

Core Takeaway: Diverse skills that induce **object interactions** lead to better downstream performance by exploiting **dynamic bottlenecks**

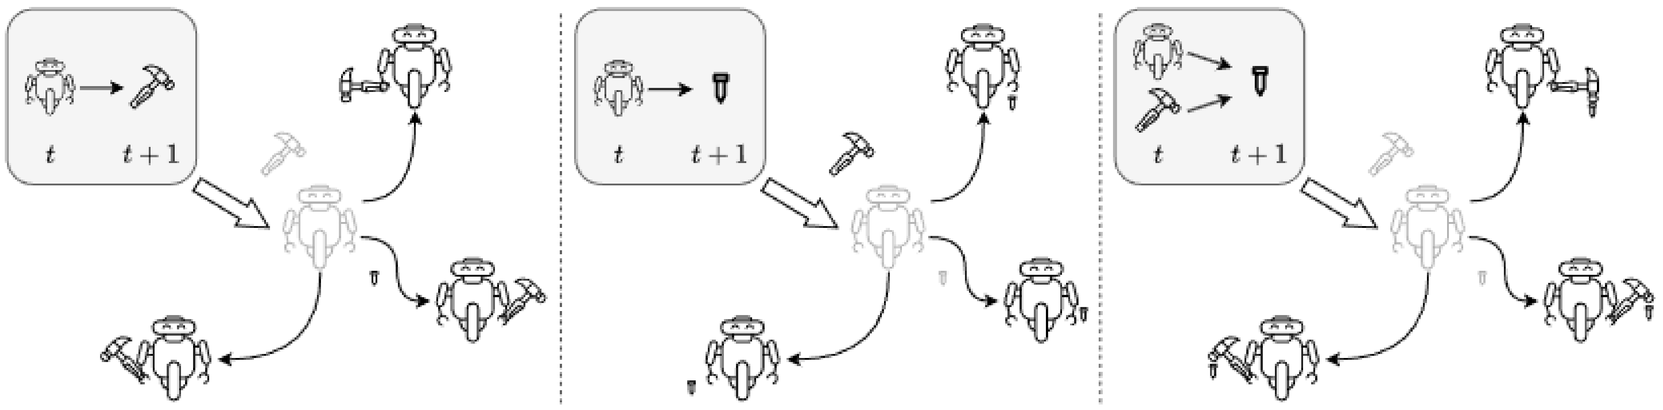

1. Identify possible interactions with **learned dynamics models**
2. Learn policies that explore **diversity** in the space of a desired interaction
3. Learn dynamics bottlenecks by **exploring** rare interactions
4. Learn a task-specific policy to **exploit** interaction-inducing skills

Skill Learning

Interaction Inference via pointwise Conditional Mutual Information

$$\mathrm{pCMI}(y; x^i | \{x/x^i\}) = \log \frac{p(y|x)}{p^{\{X/X^i\}}(y|\{x/x^i\})} > 0.$$

$$\mathcal{G}^{ij} := \mathrm{pCMI}((x^i)'; x^j | \{x/x^j\})$$

Skill Learning Reward $\quad \mathcal{R}_{\text{skill}} = \mathbb{1}[g_{\text{induced}} = g] \cdot (1 + \lambda \mathcal{R}_{\text{diversity}}),$

Graph-selection Reward $\quad \mathcal{R}_{\mathcal{G}} = \frac{1}{\sqrt{C(g_{\text{visited}})}},$

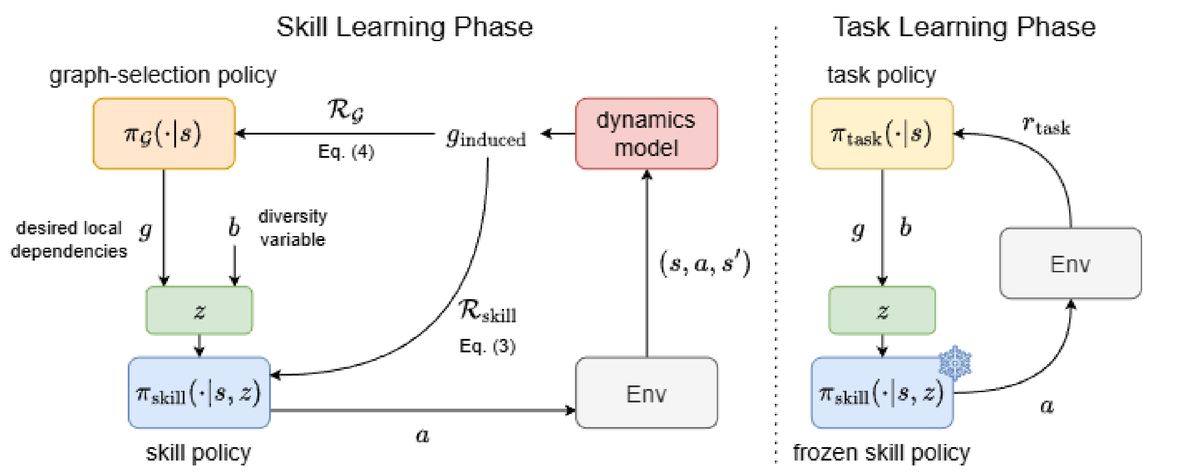

Evaluation

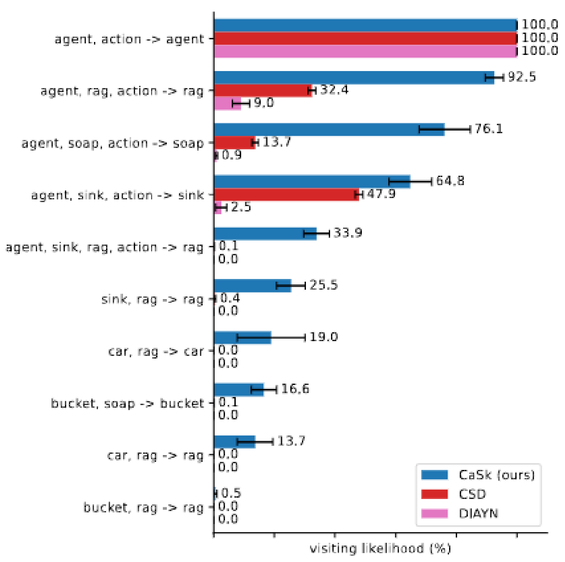

Induces relevant interactions **more frequently** than other methods

Skills improve **downstream task performance**, especially in high dimensional, challenging tasks

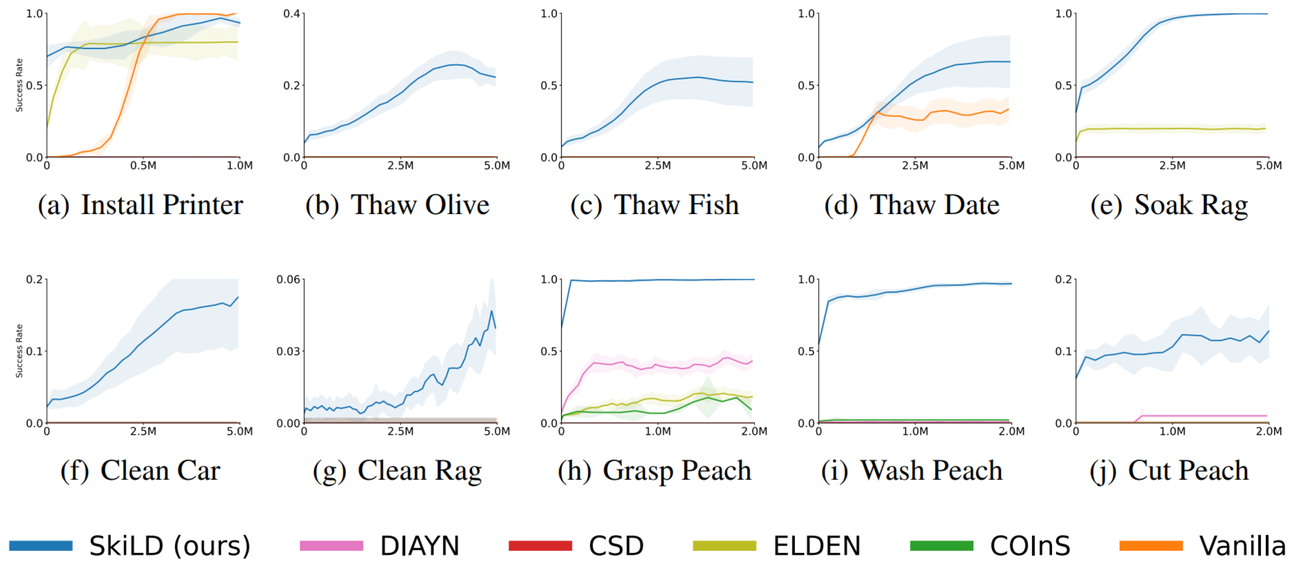

SCALAR LAB Safe, Confident, and Aligned Learning • Robotics

MIDI lab

RobIn ROBOT INTERACTIVE INTELLIGENCE LAB

LARG Learning Agents Research Group The University of Texas at Austin

TEXAS The University of Texas at Austin

TEXAS Robotics

University of Massachusetts Amherst

Sony AI

* Equal Contribution