



Planning Goals for Exploration

Edward S. Hu, Richard Chang, Oleh Rybkin, Dinesh Jayaraman

Unsupervised Task Discovery

Dropped into an unknown environment, what should an agent do to discover what it can do and how to do it?



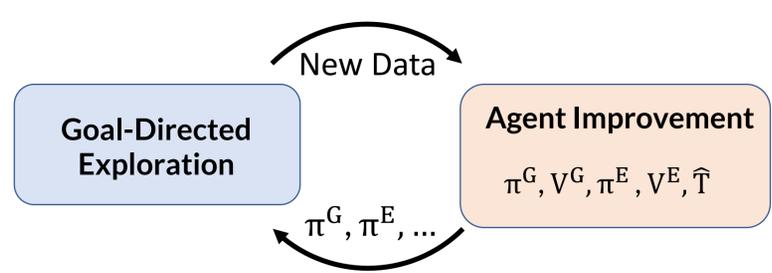
Agents first explore the environment, and then are asked to perform arbitrary tasks at test time.

More diverse exploration experience
⇒ more generalist agents

Task Discovery in Goal-Conditioned RL

Goal-conditioned RL: Train an agent $\pi^G(a|s, g)$ to reach any specified goal.

GCRL Training Loop



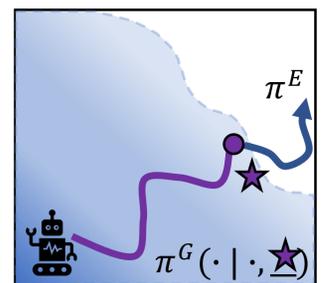
Goal Directed Exploration (Go-Explore style)

Go-phase:

Select a goal, execute π^G , and end up in s_T

Explore-phase:

Launch π^E from s_T to explore.

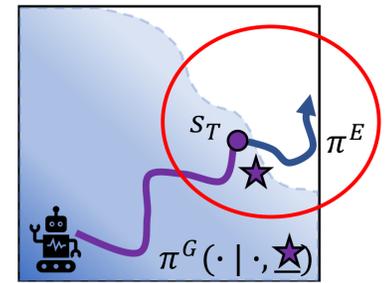


Open Question:

How to set goals for exploration?

Goals that maximize exploration

Optimize for the exploration rewards gathered by our exploration policy π^E ?



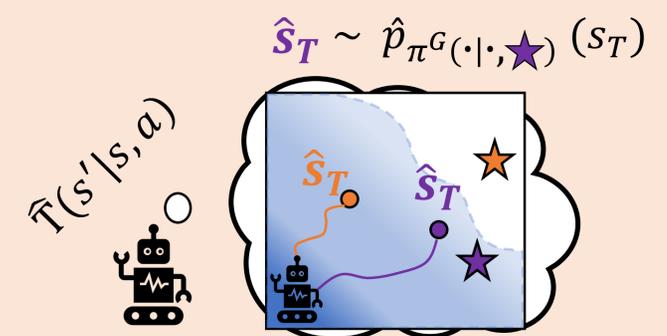
$$\max_g \mathbb{E}_{p_{\pi^G(\cdot|\cdot, g)}(s_T)} [V^E(s_T)]$$

Where π^G ends up following g How much will π^E explore

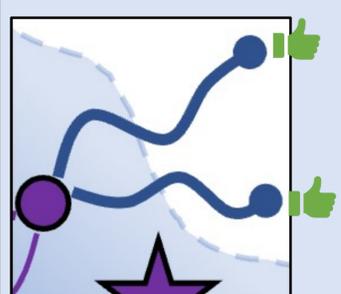
Choose goals that when followed by π^G , result in good launching points for π^E

Planning Exploratory Goals (PEG)

1. Forecast $\pi^G(\cdot|\cdot, g)$ trajectories



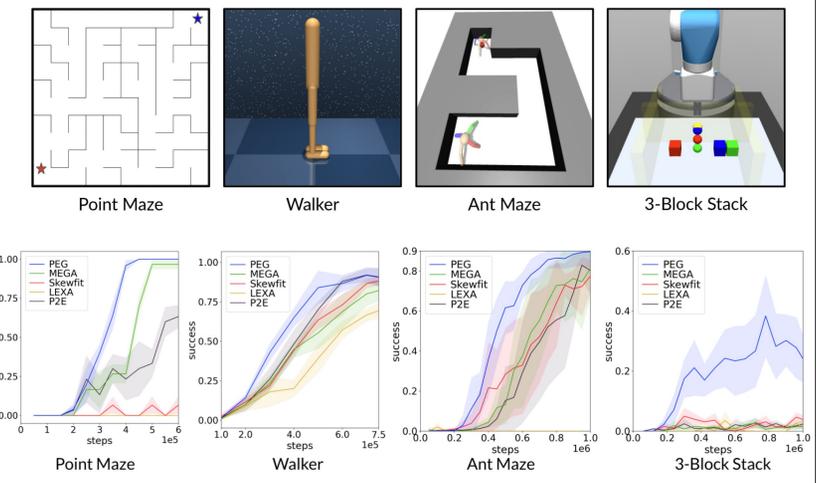
2. Score goals by expl. value of terminal state



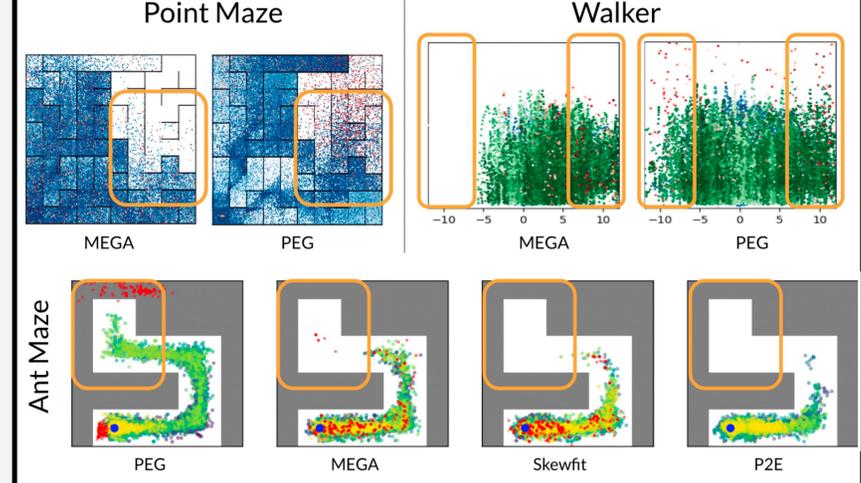
$$\begin{aligned} \text{Score}(\star) &= V^E(\hat{s}_T) \\ &= \mathbb{E}_{\pi^E} [\sum r_t^e] \\ &= (\text{thumbs up} + \text{thumbs up}) / 2 \\ &= \text{thumbs up} \end{aligned}$$

3. Use best goal for Go-Explore exploration

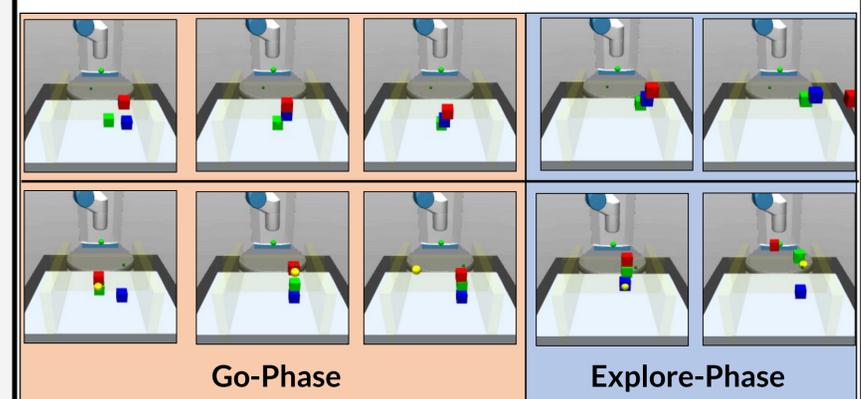
Experiments



How does PEG set goals for exploration?



Block Stacking Exploration



PEG commands the Go-policy π^G to stack blocks and the Explore-policy π^E throws them out.

Check out our website for more info!
<https://sites.google.com/view/exploratory-goals>