

Skill-Driven Neurosymbolic State Abstractions

Alper Ahmetoglu, Steven James, Cameron Allen, Sam Lobel, David Abel, George Konidaris

1 When you have action abstractions ...

$\bar{a}_1 = (I_{\bar{a}_1}, \pi_{\bar{a}_1}, \beta_{\bar{a}_1})$

2 ... abstract states are distributions over ground states.

$$V(s) = \sum_{\bar{a}} \pi(\bar{a} | s) \mathbb{E}_{s'} [R(s, \bar{a}, s') + \gamma^T V(s')]$$

$$V(\bar{s}) = \mathbb{E}_{s \sim \bar{s}} [V(s)] \quad R(\bar{s}, \bar{a}, \bar{s}') = \mathbb{E}_{s \sim \bar{s}, s' \sim \bar{s}'} [R(s, \bar{a}, s')]$$

Abstract Bellman:

$$V(\bar{s}) \approx \sum_{\bar{a}} \bar{\pi}(\bar{a} | \bar{s}) [R(\bar{s}, \bar{a}, \bar{s}') + \gamma^T V(\bar{s}')]$$

3 Abstract states should be Markov ...

Transition is a function of $\bar{s}, \bar{a}, \bar{s}'$ ✓
 Reward is a function of $\bar{s}, \bar{a}, \bar{s}'$ ✓

Available actions are a function of \bar{s} . ✓

... and model-preserving.

$$\int_{s'} |T(s' | s, \bar{a}) - T(s' | \bar{s}, \bar{a})| ds' \leq \epsilon_T$$

$$|R(s, \bar{a}) - R(\bar{s}, \bar{a})| \leq \epsilon_R$$

$$\Rightarrow |Q^\pi(s, \bar{a}) - Q^\pi(\bar{s}, \bar{a})| \leq \frac{\epsilon_R + \gamma V_{MAX} \epsilon_T}{1 - \gamma}$$

4 Construct the initial abstract MDP ...

... and refine abstract states if necessary.

5 Abstract states recover the underlying model ...

neurosymbolic-mdps.github.io

6 ... and factored skills generate factored MDPs.

$$s = (s_1, s_2, s_3, s_4, s_5)$$

$$s' = (s_1, s_2, s_3, s'_4, s'_5)$$