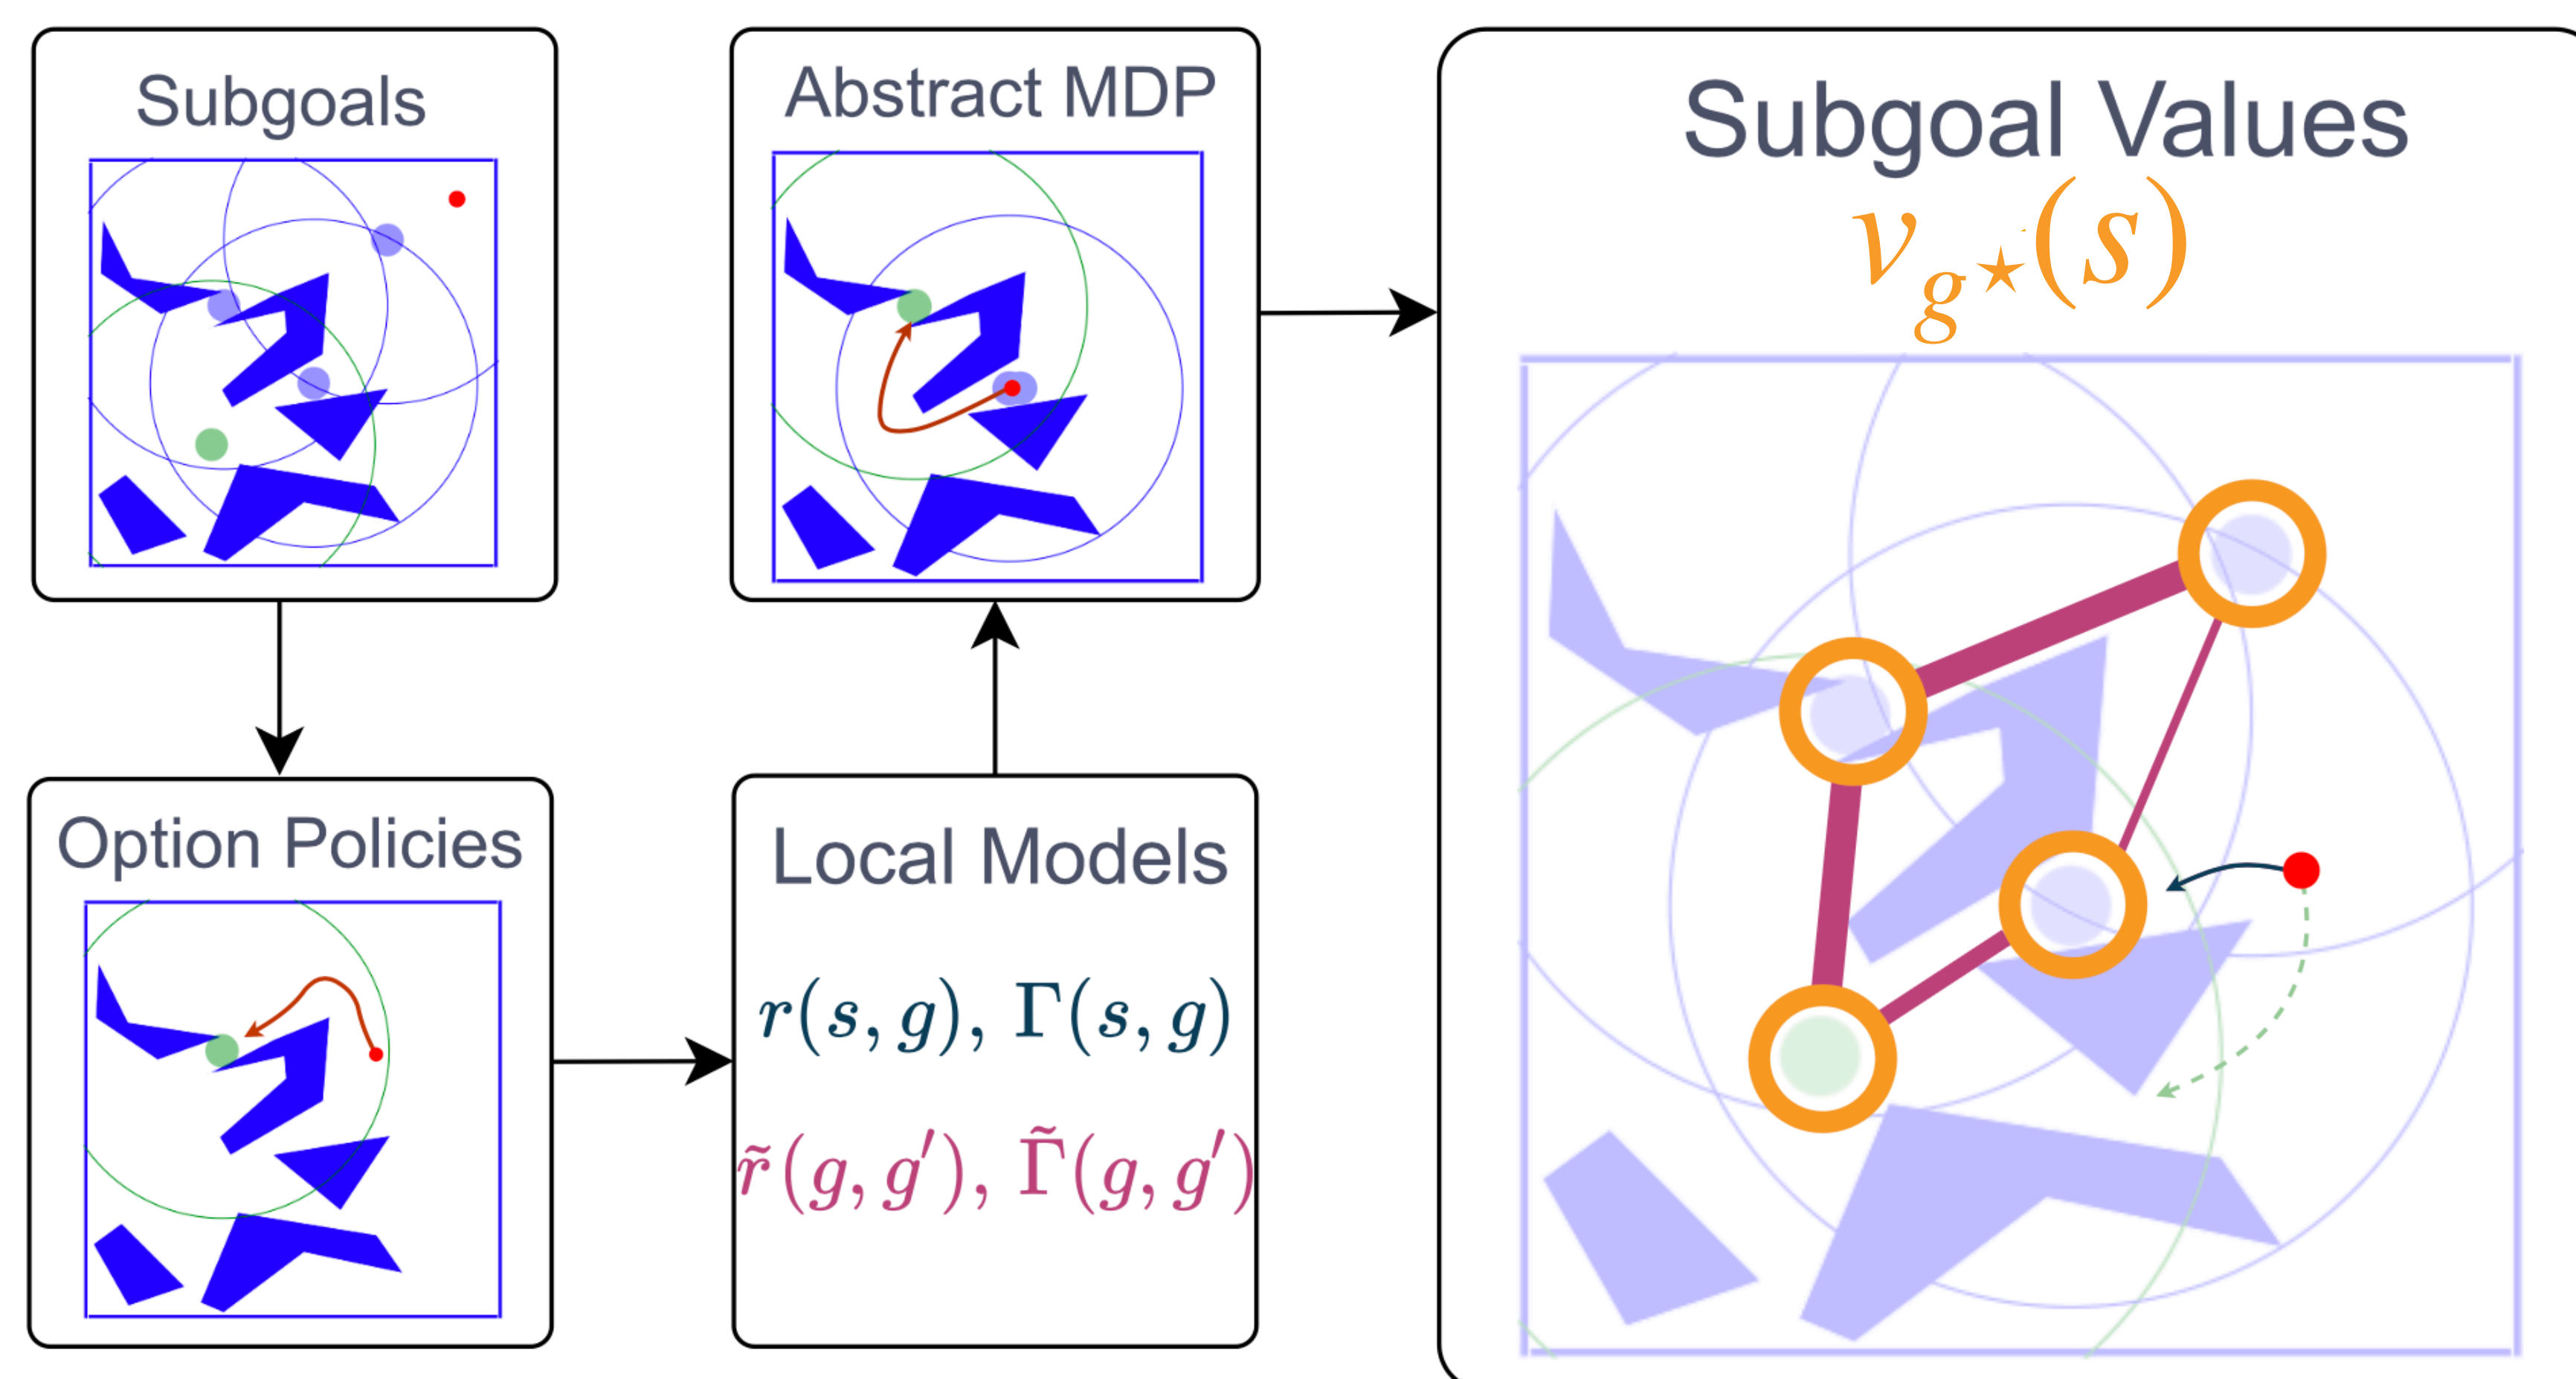
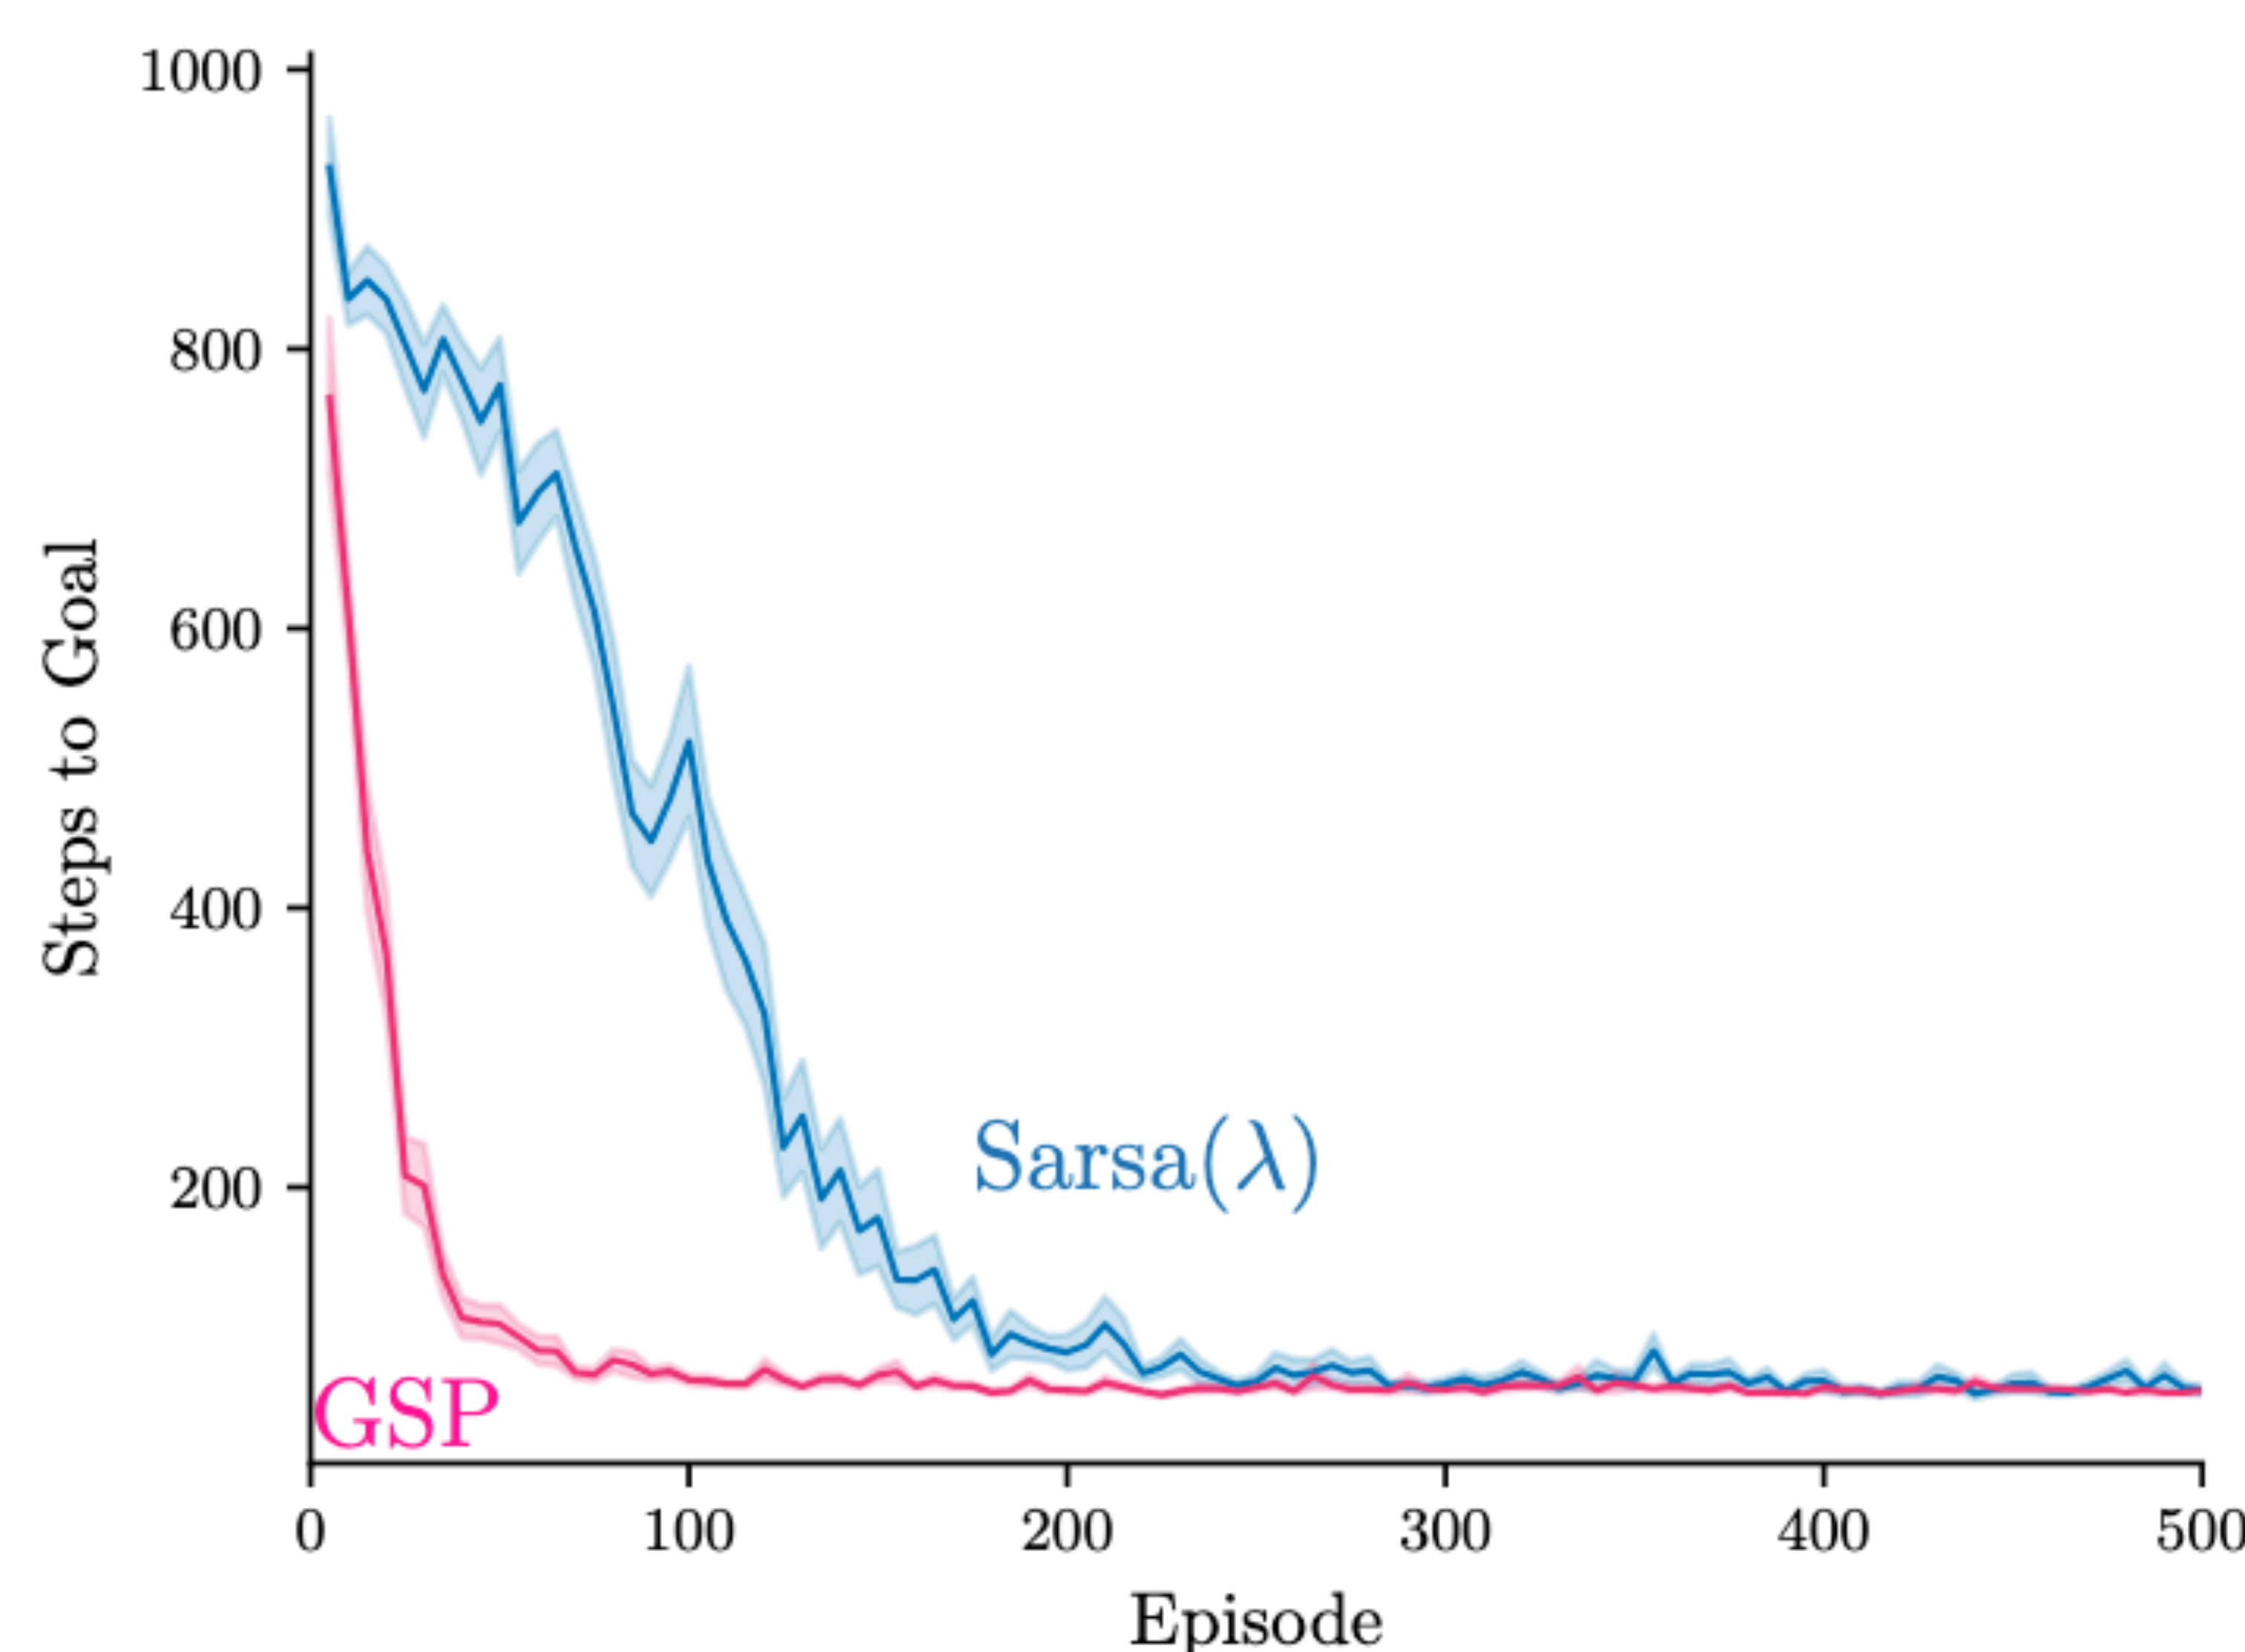


Planning in an Abstract Space



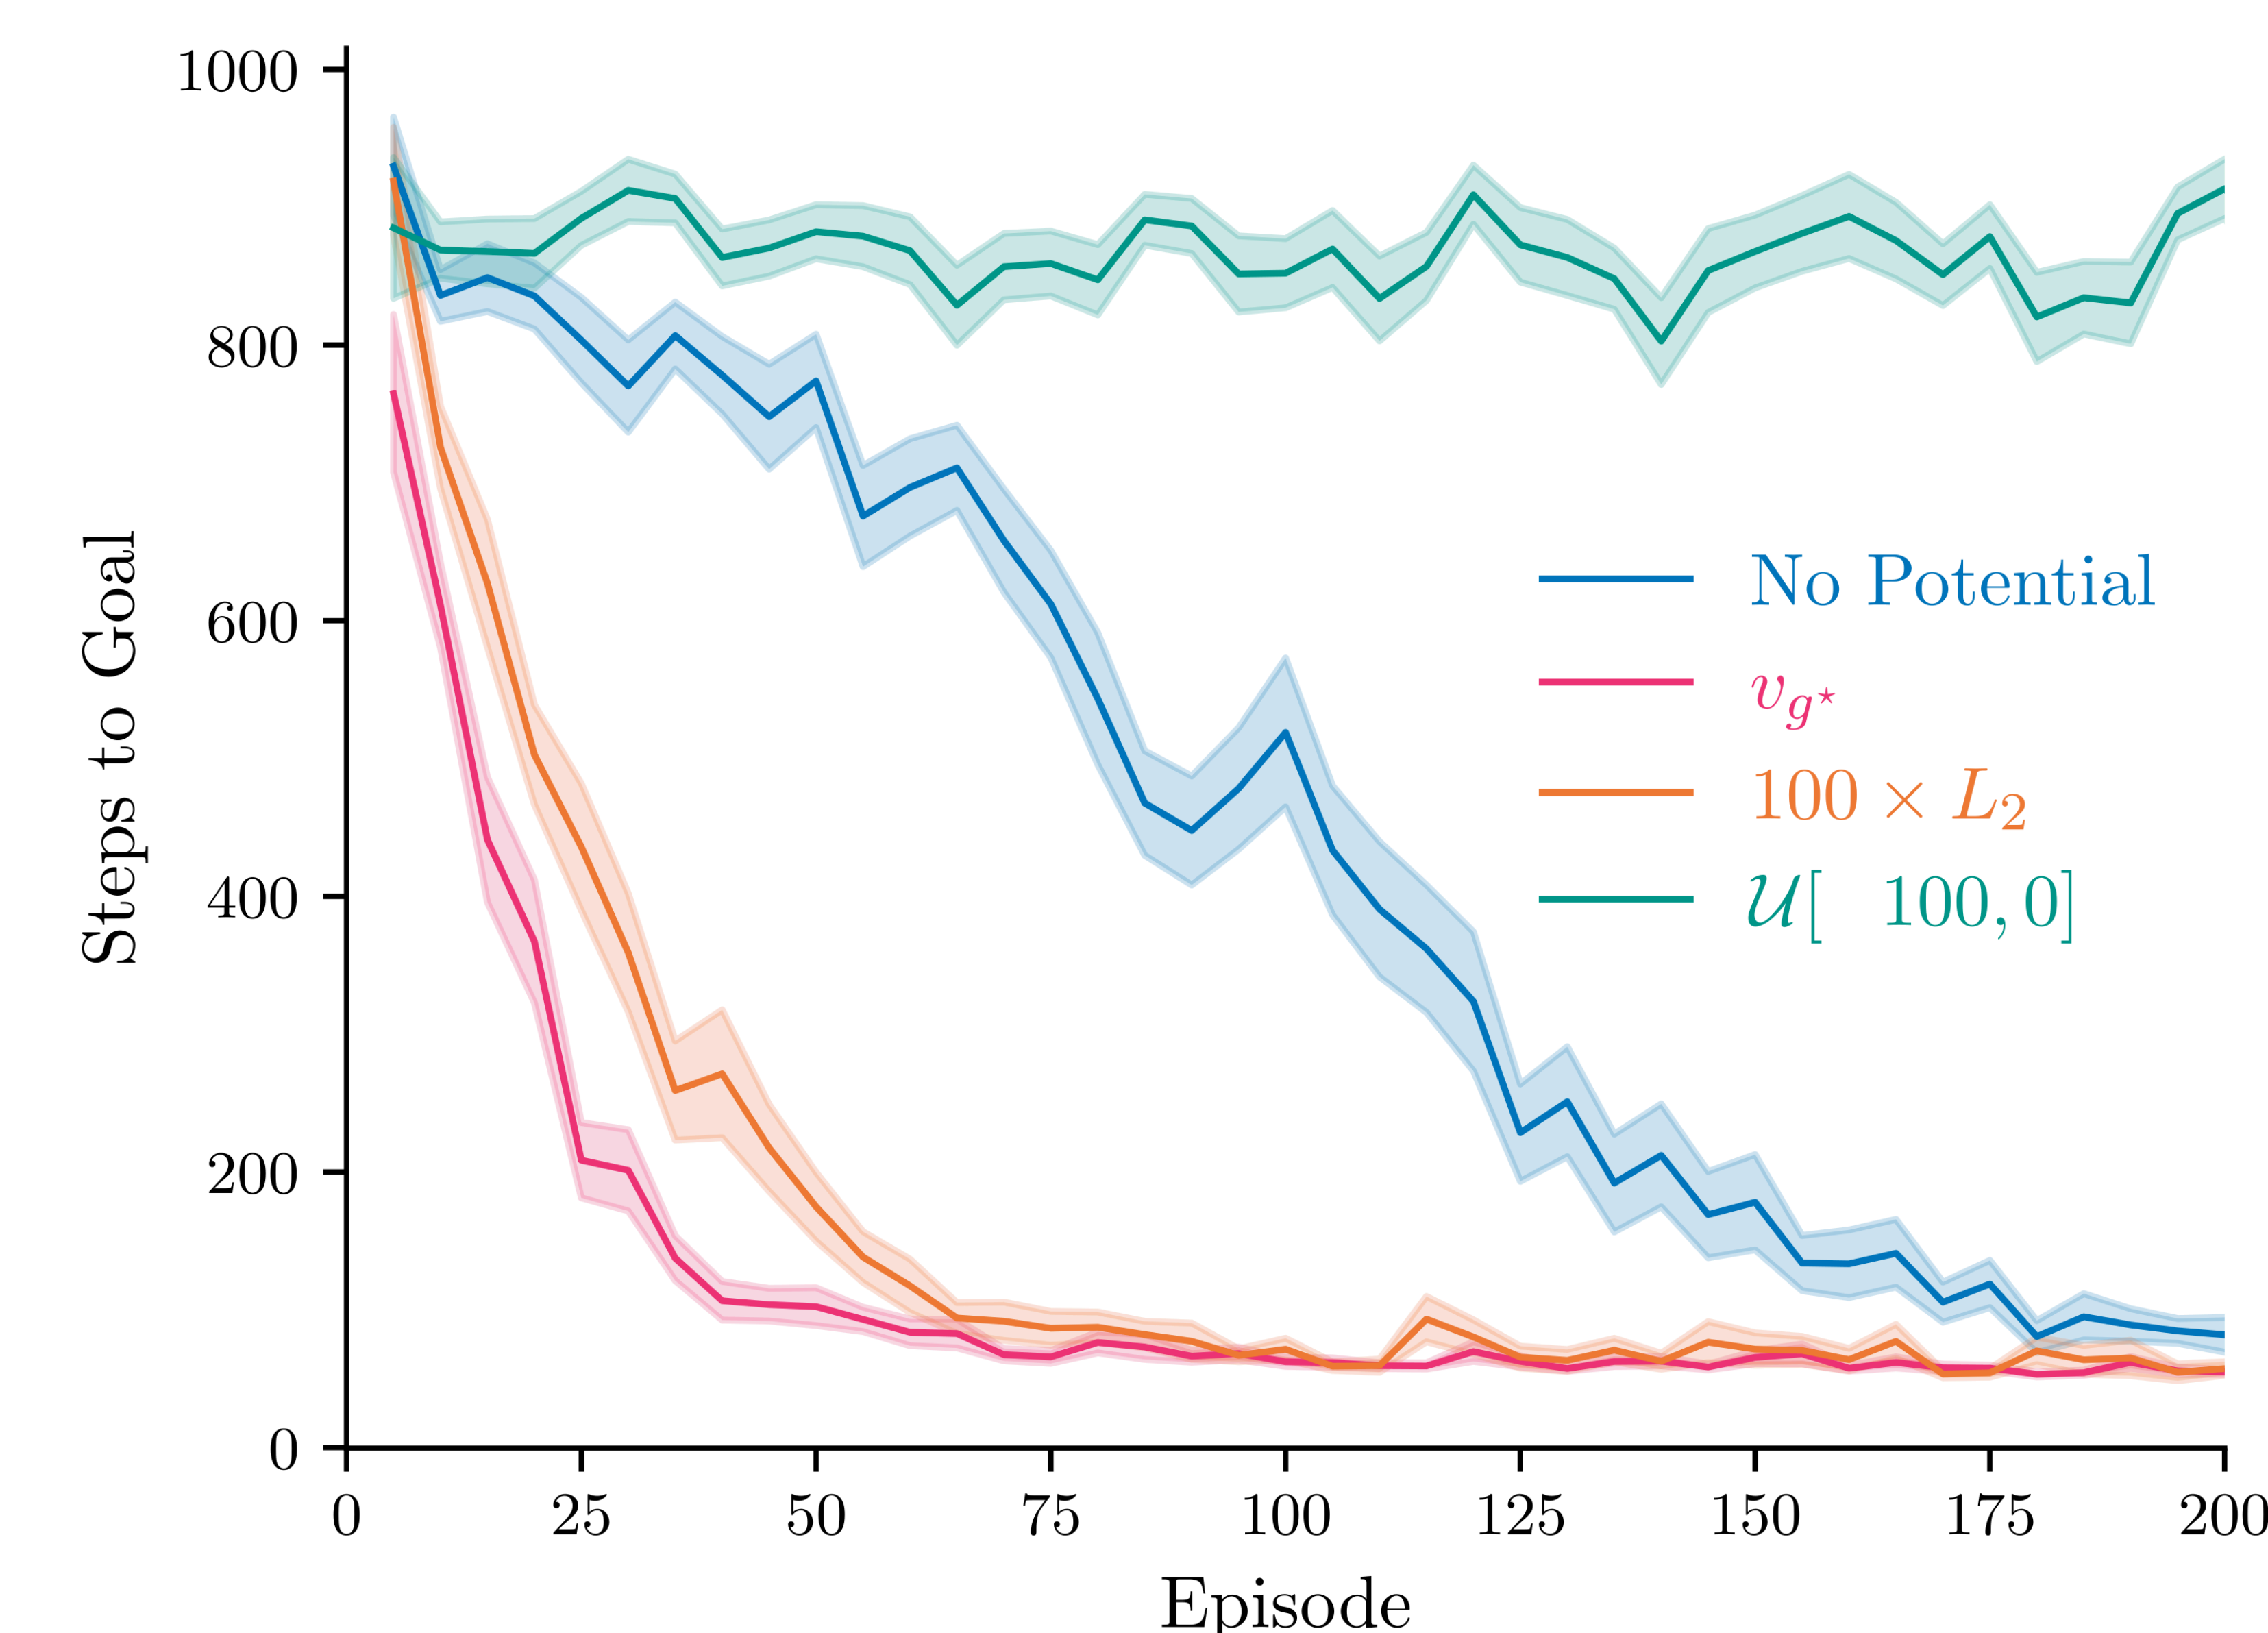
- Planning in large Markov Decision Processes (MDPs) is expensive.
- We construct an Abstract MDP of subgoals and options.
- We learn return and discounted probability models of these options.
- Perform value iteration in the abstract MDP and project values down to the original space for reward shaping.

Goal-Space Planning is Faster




Abstract models can be used for planning by reward shaping.

Shaping with World Models



GSP Propagates Value



Limitations

- Computational cost of model learning.
- Discovering subgoals for planning.
- Shaping is not as helpful in trajectory-based methods (REINFORCE or PPO)

Goal-Space Planning with Subgoal Models

Chunlok Lo¹², Kevin Roice¹²,
Parham M. Panahi¹², Scott M. Jordan¹²,
Adam White¹²³, Gabor Mihucz¹²,
Farzane Aminmansour¹², Martha White¹²³

¹ Department of Computing Science, University of Alberta

² Alberta Machine Intelligence Institute

³ Canada CIFAR AI Chair



Scan me!