

Leveraging Experience in Lazy Search Mohak Bhardwaj, Sanjiban Choudhury, Byron Boots, Siddhartha Srinivasa



Motion Planning on Roadmaps



Piano Mover's Problem



Nodes - Robot Configurations **Edges** - Robot Motions

Plan collision-free paths on a graph of robot motions in real-time



Collision checking edges is the computational bottleneck [1]

Imitation of Clairvoyant Oracles

Key Insight 2: Leverage IL to learn MDP policy from data (Sample efficiency by guidance from (near-optimal) experts [3])

Optimal set of edge-checks if we are clairvoyant (world known)?



Lazy search defers edge evaluation until absolutely necessary [2]

Several checks of robot mesh against world

Lazy Search as a Markov Decision Process

LazySP minimizes edge evaluations[1]











STROLL - Search through Oracle Learning and Laziness Leveraging no-regret online learning for IL [4]









SelectAlter

Sample a world from database.

Execute learner(evaluate edges) to reach a state

Query clairvoyant oracle for Continue aggregating data edge (action) to evaluate. till the end, update learner.

Experimental Results

Linear policy on graph features such as Posterior, Centrality, Location

STROLL focuses on edges with high invalid likelihood+centrality









Optimal Selector for LazySP is the optimal policy for the MDP , i.e. minimizes edge evaluations



Challenges in Solving the MDP

Convergence of Large state space approximate VI **Sparse Rewards** Potentially unstable Exponential in Tabula rasa RL

Reference

WALI

Better Posterior+Centrality

TWOWALL

FOREST

In 7D STROLL suffers due to poor posterior + centrality estimates

Future Improvements

RL+IL

BAFFLE

Oracle is not realizable - RL can improve policy further. [5]

