Difference-Aware Retrieval Policies for Imitation Learning

Quinn Pfeifer[1], Ethan Pronovost[1], Paarth Shah[2], Khimya Khetarpal[3,4], Siddhartha Srinivasa[1], Abhishek Gupta[1,2]

[1]Paul G. Allen School of Computer Science & Engineering, University of Washington, [2]Toyota Research Institute, [3]Google DeepMind, [4]Mila

Behavior Cloning (BC)

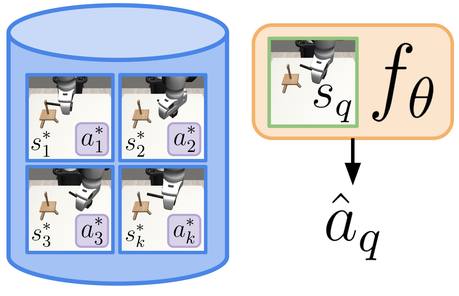

Problem: Accumulation of error leads agent to OOD states

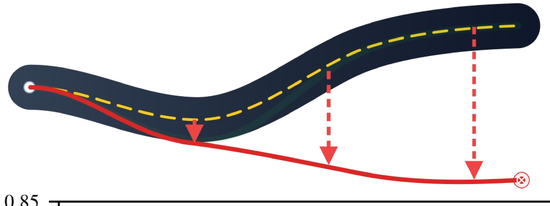

Difference-Aware Retrieval Policies

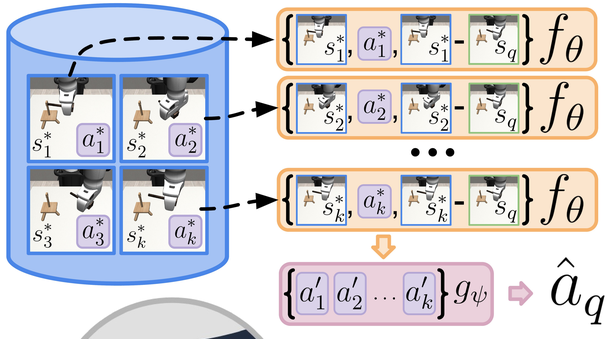

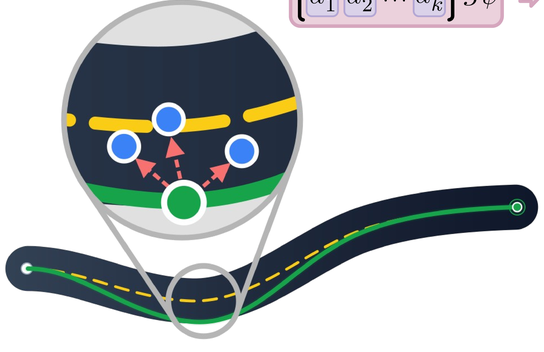

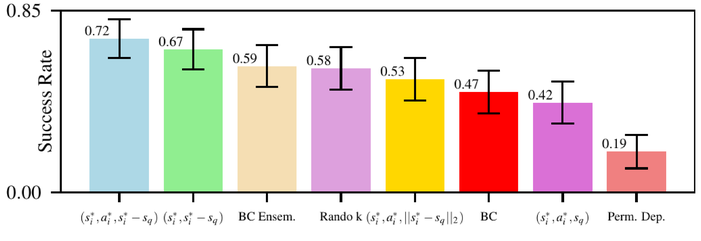

What Matters?

Relative distance vectors rather than scalar distance or the query point directly.

A **permutation invariant aggregator.**

Results

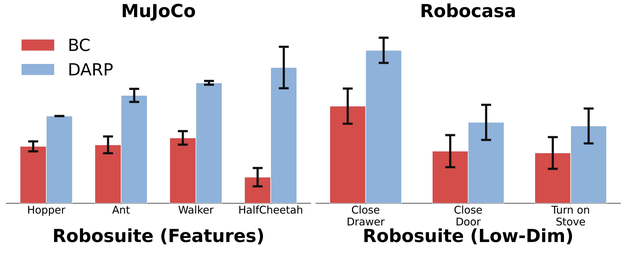

Enhances state of the art diffusion models

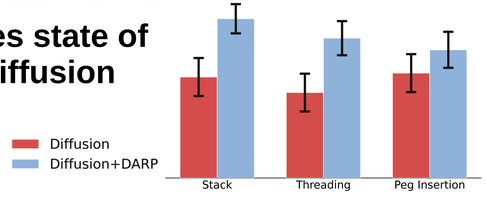

Why Does This Work?

Theoretical Understanding

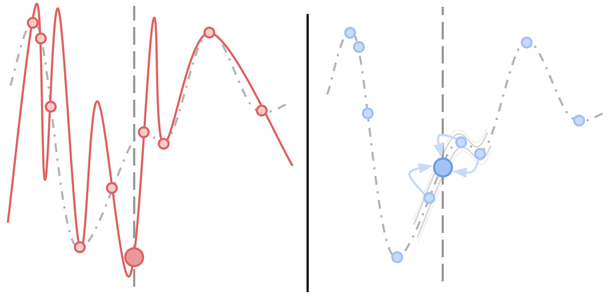

BC: Arbitrarily jagged

DARP: Local neighborhood induces smoothness
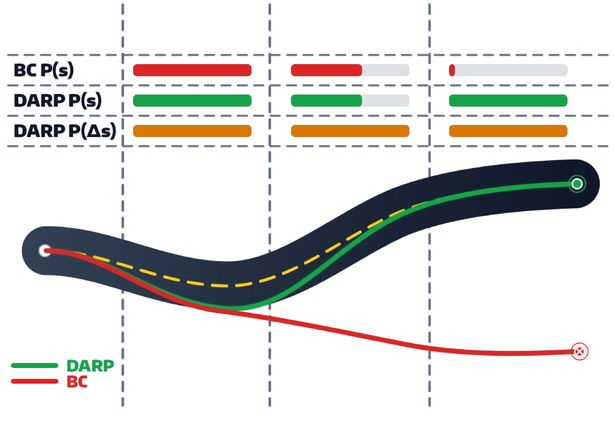
Empirical Understanding