






ABLE: Anticipating Belief Updates for Learning Efficient TAMP

Minseo Kwon, Tom Silver
Princeton University



Motivation: Limb Manipulation

- Caregiving and rehabilitation using a robot require moving a person's limb by direct, forceful contact.
- Long-horizon: where to stand, where to grasp, how to move the limb $\Rightarrow$ Task and Motion Planning (TAMP).
- The human's range of motion (RoM) differs across people



Problem Formulation

TAMP domain (PDDLStream): The plan-limb-motion stream is an MPC that searches over robot torques under the human and robot limits.

Domain	Streams
move-base	plan-base-motion
grasp	sample-grasp plan-grasp-motion
move-limb	plan-limb-motion

Task: $\tau = (q_h^0, q_r^0, q_h^{\text{goal}}) \sim P(\tau)$ move the limb to a goal config


Research Question

Learn the human's parameters θ
What is the human's range of motion?


Learn the TAMP solver's parameters ψ
How can it plan fast for this person?

How can we learn to plan faster for the new user?

Human Model




Approach



Experimental Setup

- 5 simulated people $\times$ 15 episodes $\cdot$ left arm (3 joints) $\cdot$ 4 dreams / episode $\cdot$ 600 s per plan
- Tested on a left-arm manipulation environment in PyBullet sim and TidyBot

Proof of Concept: Dreaming



- Evaluated only the dreaming and execution steps, on a fixed set of 6 goals
- Dreams make planning more efficient:** with dreams, the robot reuses plans prepared while dreaming, while planning without memory or with experience-only memory stays slow.
- Future work:** since dreamed plans already speed up planning, learn a model from them instead of caching, so it generalizes to unseen goals and beliefs.

Contact: mk1406@princeton.edu