

LEARNING TO DETUMBLE: ADAPTIVE POST-CAPTURE STABILIZATION OF UNCOOPERATIVE SPACE DEBRIS

Álvaro Belmonte-Baeza*, Celia Redondo-Verdú, José L. Ramón, Gabriel J. García, Miguel Cazorla, and Jorge Pomares

University of Alicante, Spain
*alvaro.belmonte@ua.es

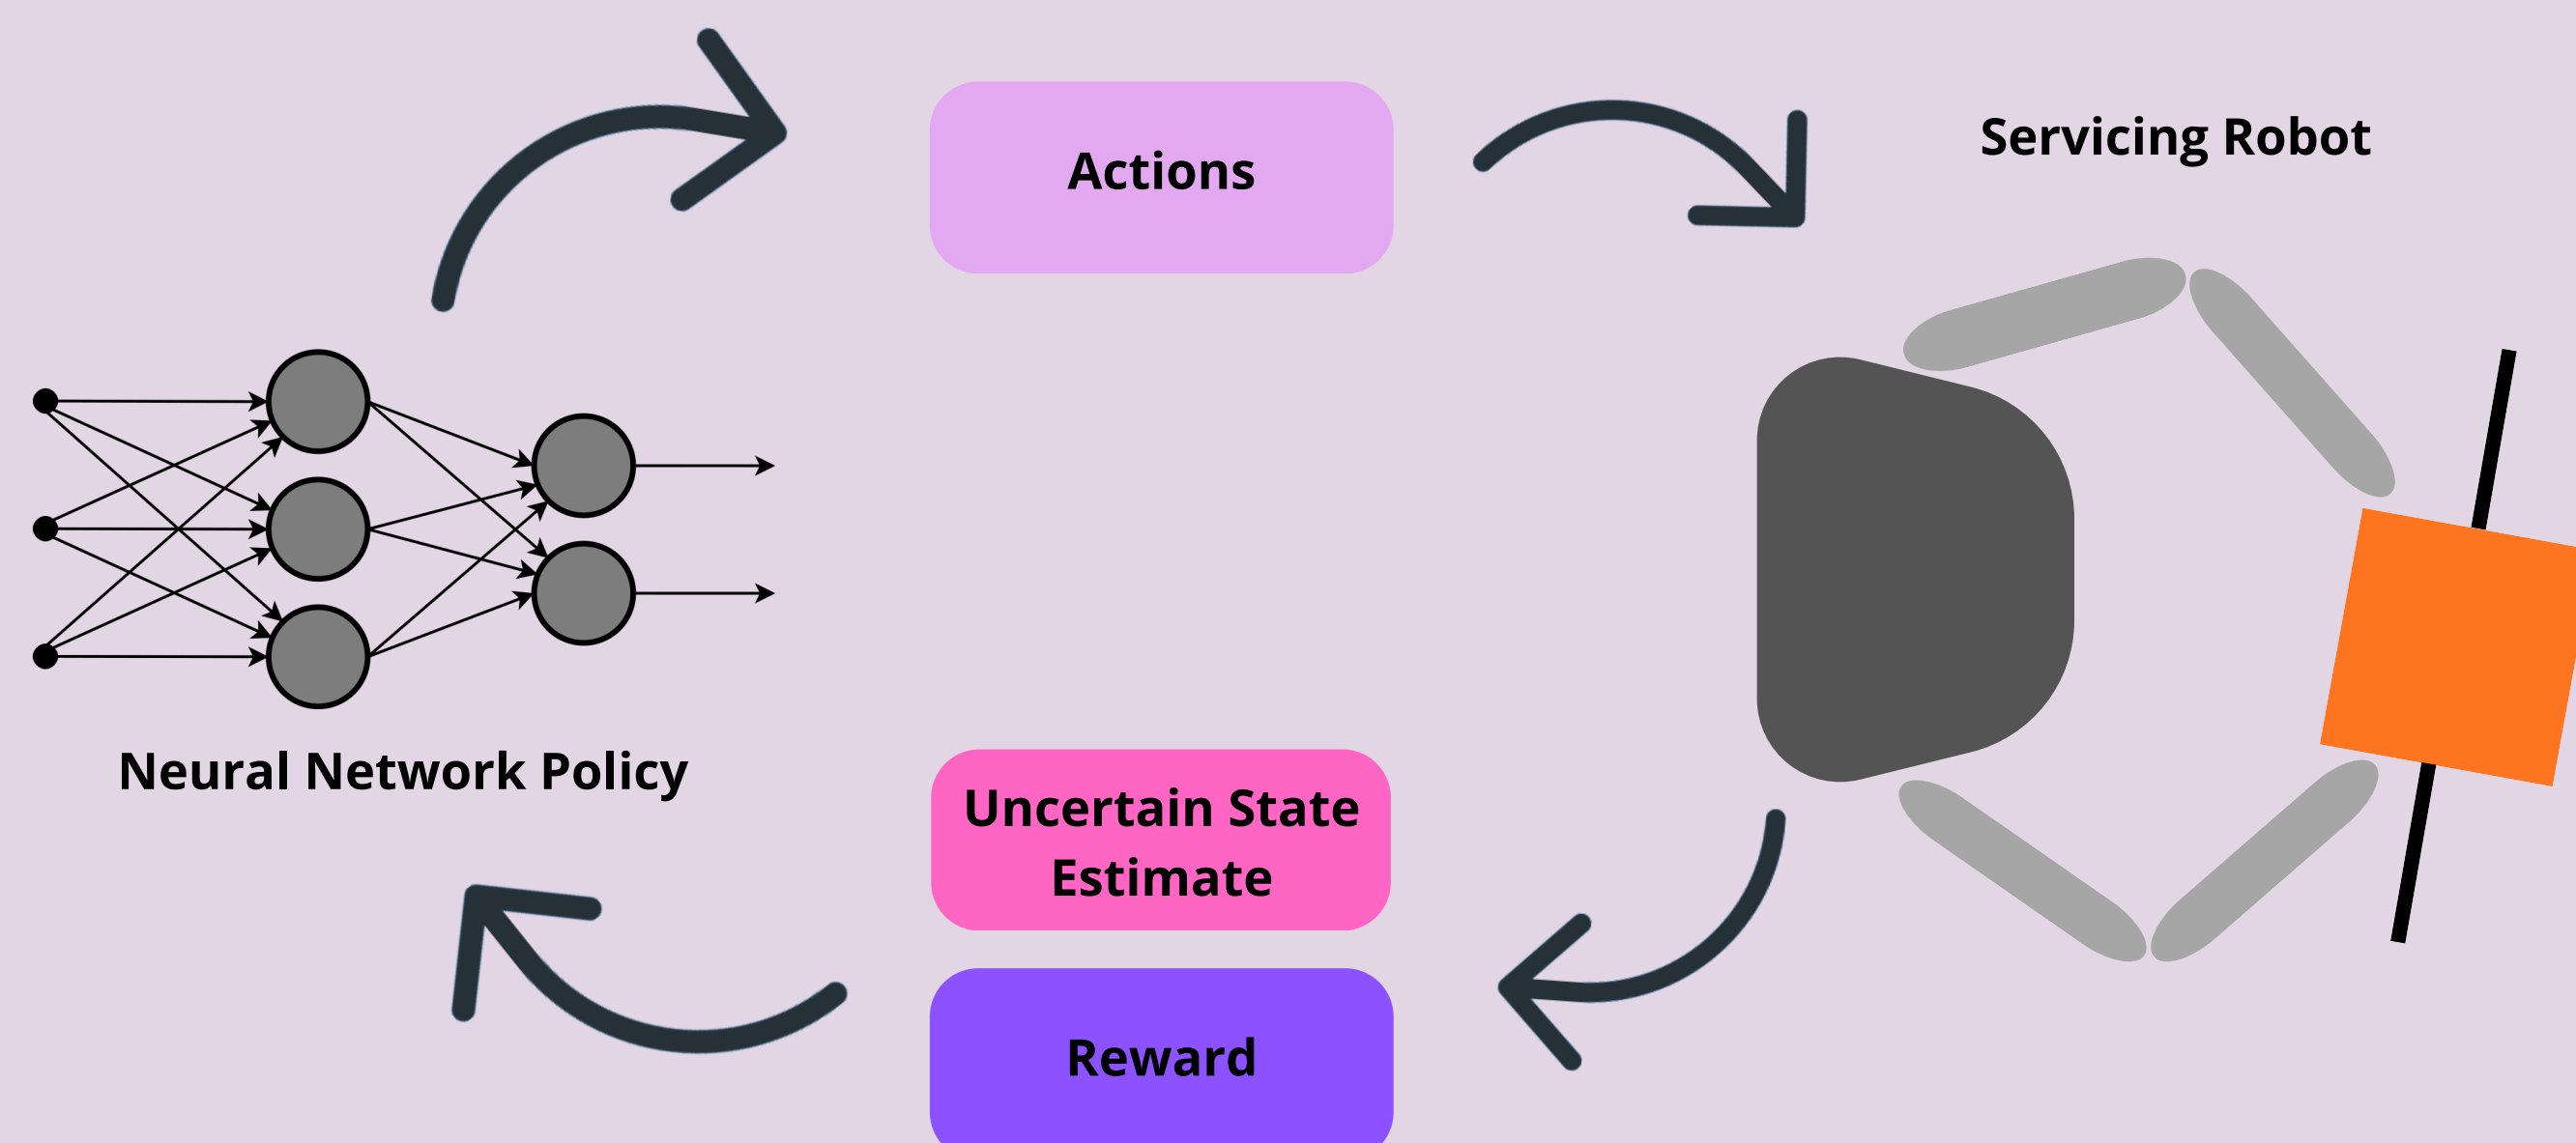
Motivation

Challenge

- Active Debris Removal (ADR) missions face a major technical hurdle in the autonomous capture and stabilization of non-cooperative, freely tumbling targets.
- Traditional model-based control methods struggle with dynamic uncertainties. These approaches require accurate state estimates, such as knowing the exact center of mass, inertia tensor, and friction constraints of the target.

Our approach

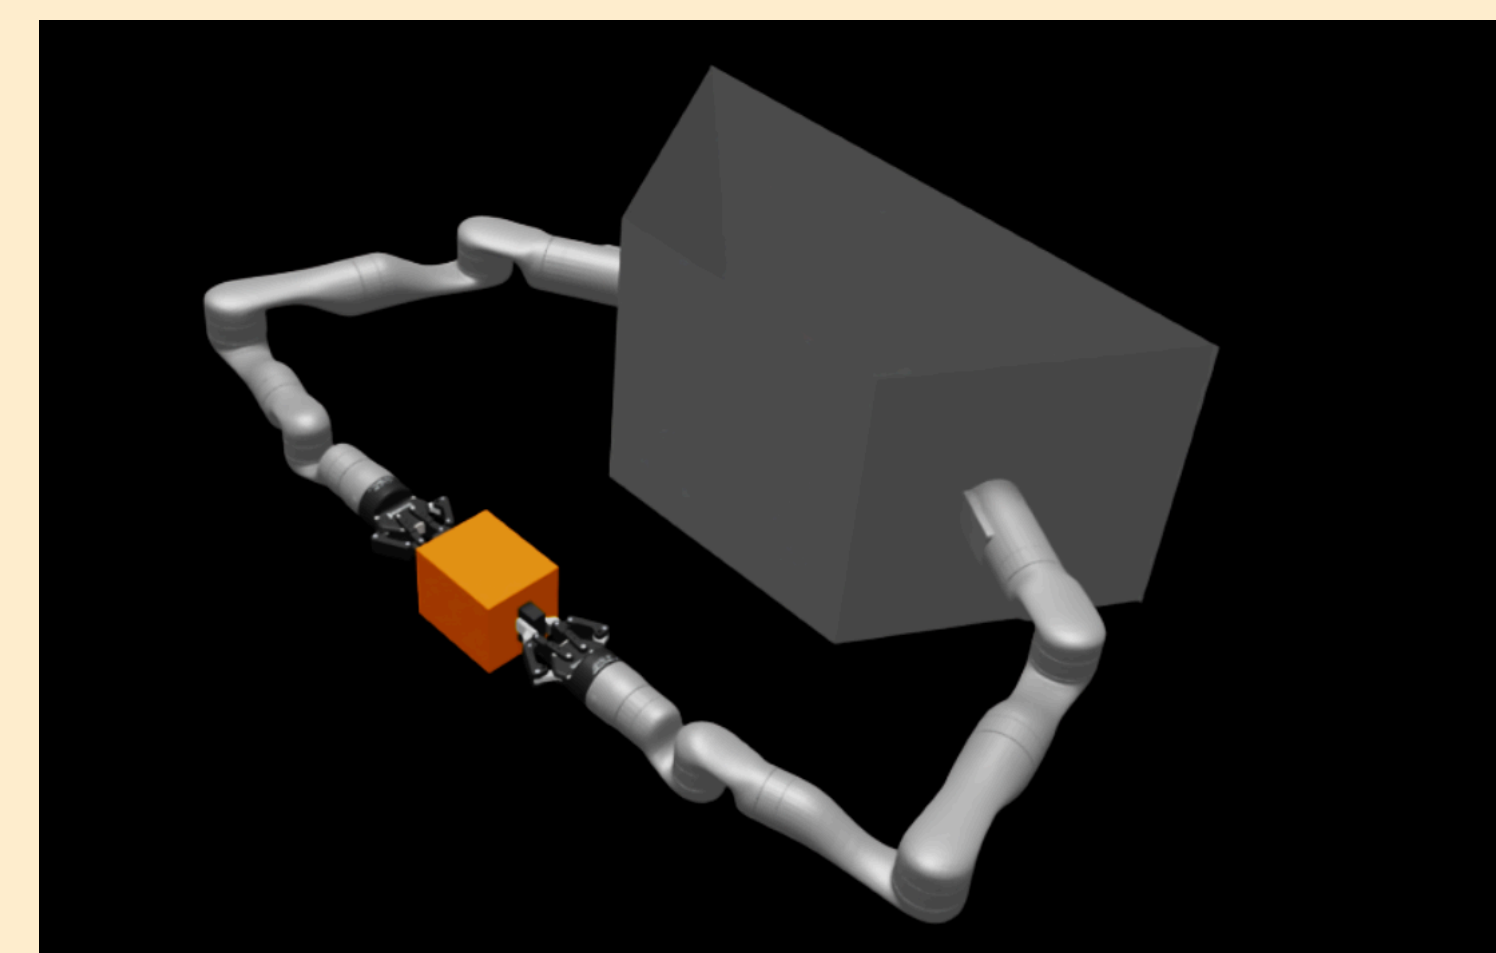
- We introduce a robust Reinforcement Learning (RL) framework specifically designed for post-capture detumbling.



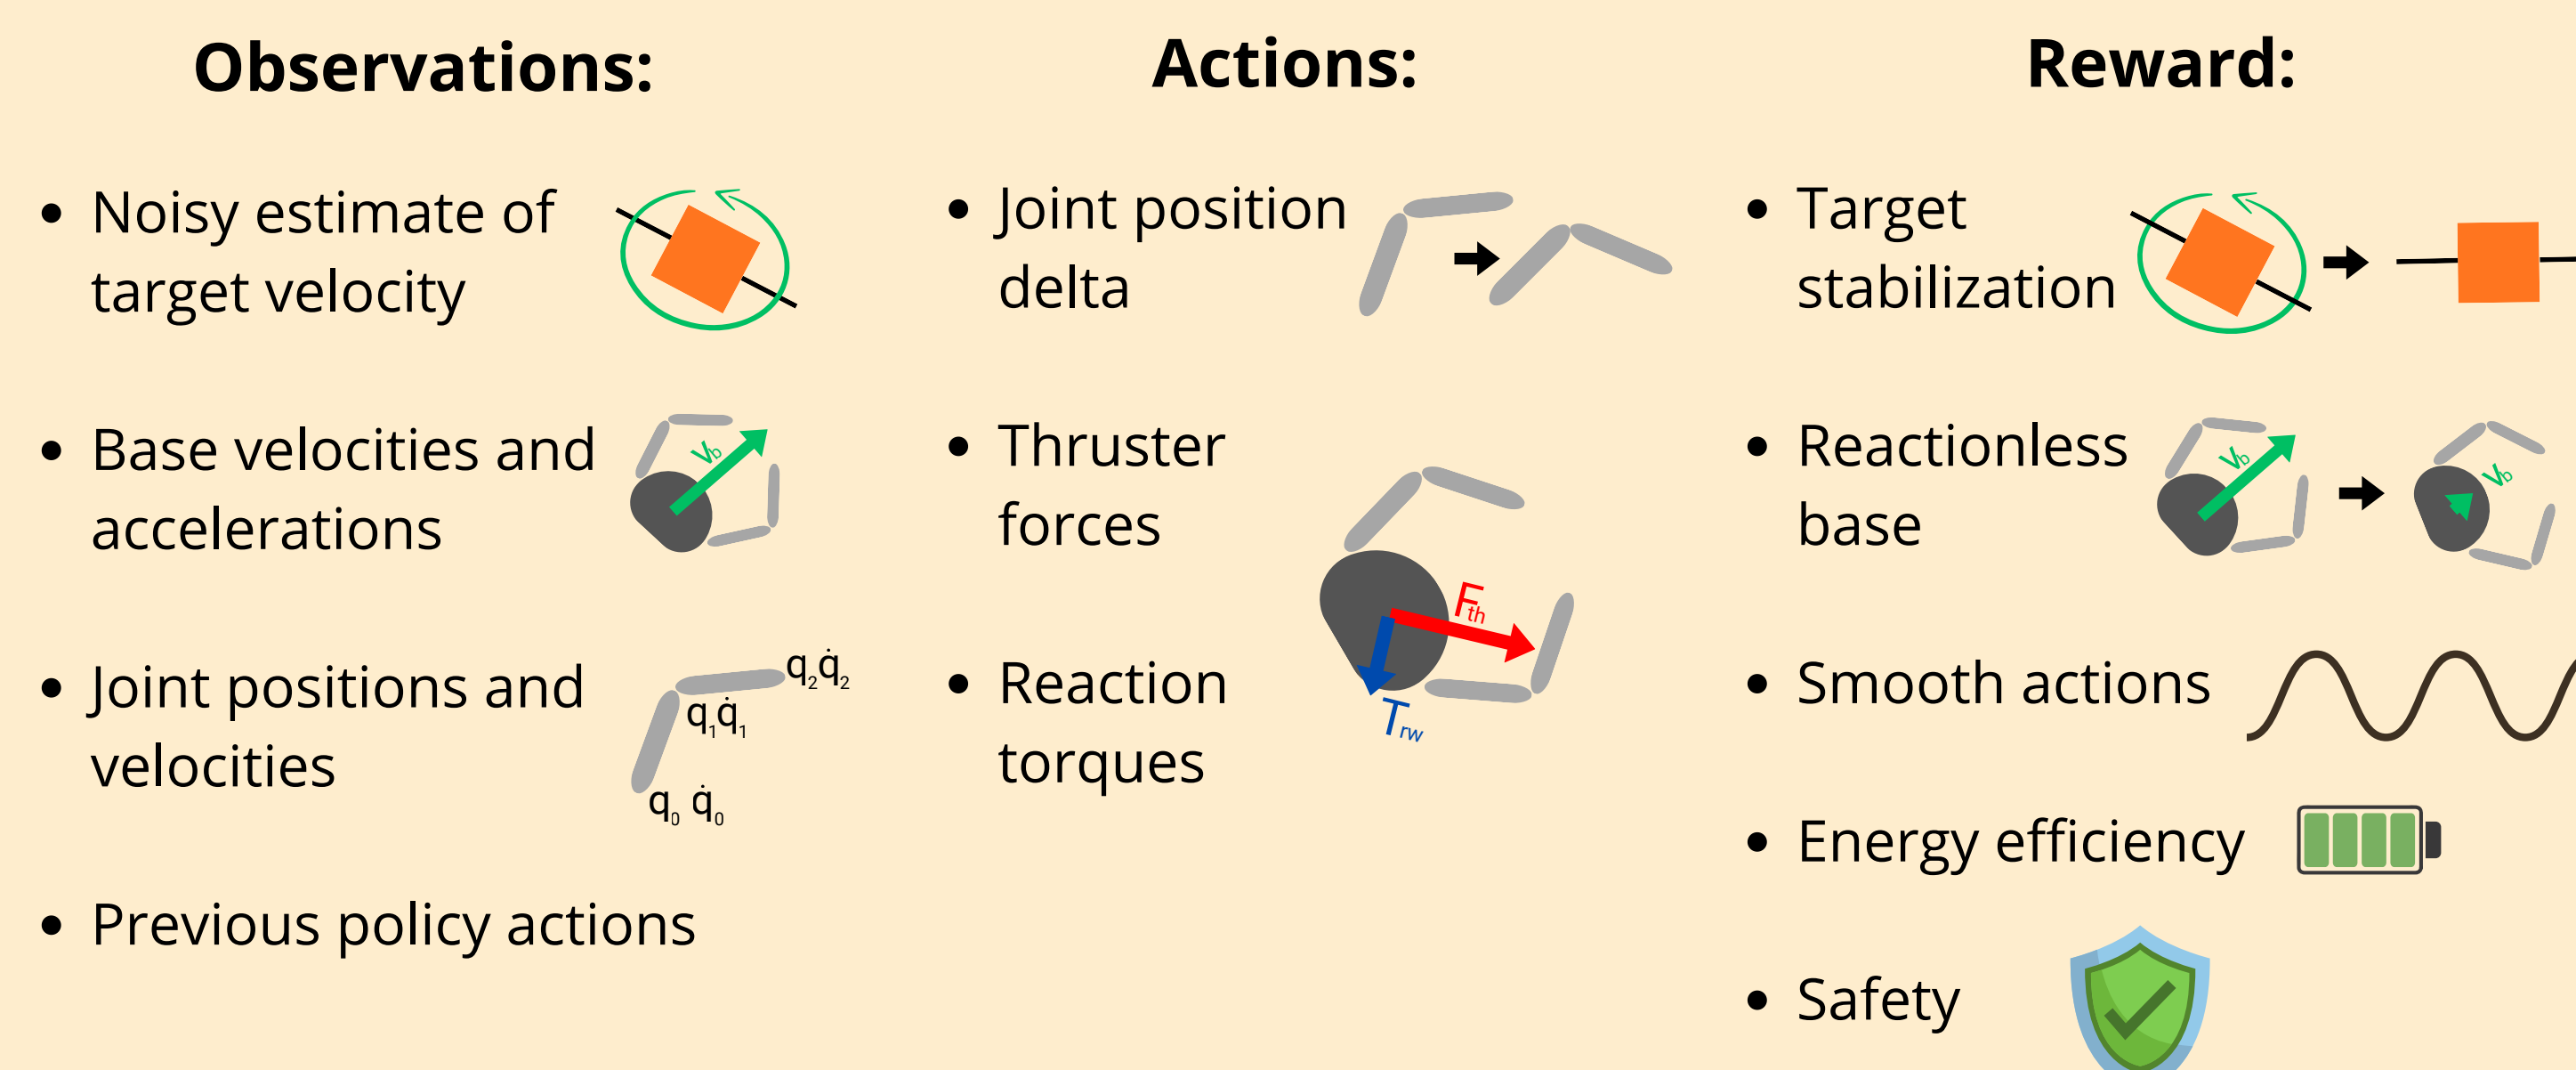
RL for Space Debris Stabilization

System and Task Design

- The robot is composed by a base satellite equipped with two 7-DoF arms.
- The base satellite uses symmetric proportional thrusters and reaction wheels to compensate dynamic coupling induced by the target
- Tumbling debris with uncertain inertial state.




MDP Formulation



Preliminary Results

Experimental Setup:

- Simulate the target stabilization task with Isaac Sim and Isaac Lab.
- Initial target velocities randomized within ± 0.01 m/s and ± 0.05 rad/s.
- Train in simulation with 4096 parallel environments to collect interaction data.



Quantitative Evaluation

- Run trained policy on 1000 environments, with randomized initial velocities for the target debris.
- We report mean velocity residuals for both the target debris and the servicing satellite across all samples.
- Debris velocity residuals:**
 - Linear: $1.7 \pm 0.7 \times 10^{-3}$ m/s
 - Angular: $1.4 \pm 0.7 \times 10^{-3}$ rad/s
- Base velocity residuals:**
 - Linear: $1.6 \pm 0.6 \times 10^{-3}$ m/s
 - Angular: $1.1 \pm 0.4 \times 10^{-3}$ rad/s

Qualitative Evaluation

