

Improving Trajectory Tracking in Reinforcement Learning by Augmenting States with Future Targets

Tuesday, 31 March 2026 12:00 (2 hours)

Standard Reinforcement Learning (RL) for trajectory tracking typically relies on myopic state representations, providing agents only with the current target. This forces a reactive control paradigm, resulting in lag and overshoot during dynamic transitions. To address this, we propose augmenting the standard RL state space, which traditionally contains only the current reference, with future target information, e.g., a finite-horizon sequence of future targets or target velocities.

We evaluate this predictive state representation on a real-world industrial testbed (Quanser Aero 2) using continuous S-curve trajectory profiles. Preliminary experiments demonstrate a significant performance improvement: augmenting the state with five future targets at 0.1 s intervals reduced the average tracking error from 2.60° (baselines) to 0.34°. These results suggest that simple state-augmentation enables model-free agents to learn sophisticated anticipatory behaviors, i.e. initiating control actions before target changes occur, without explicit model-based planning.

Student

Yes

Primary author: SCHÄFER, Georg

Co-authors: Dr REHRL, Jakob (Salzburg University of Applied Sciences); HIRLAENDER, Simon (PLUS University Salzburg); Dr HUBER, Stefan (Salzburg University of Applied Sciences)

Presenter: SCHÄFER, Georg

Session Classification: Poster session