

Evaluating Hybrid Rollout Simulation for Efficient Robotic Reinforcement Learning Environments

Paper Submitted to National High School Journal of Science for peer review (Oct'25)

Author: Keegan Crasto

Summary:

Developed a hybrid rollout framework to accelerate reinforcement learning in robotic simulation. Using the PandaReach-v3 environment from Hugging Face's Panda Gym suite, the system integrates a trained policy network with a physics-based simulator. Between periodic PyBullet updates, the simulator performs several fast state prediction steps using the policy's actions, avoiding costly physics calculations while maintaining accuracy through periodic real-environment corrections. This approach significantly reduces computation time while preserving task performance, demonstrating a scalable method for faster robotic training.

Traditional robotic simulators like PyBullet provide accurate physical modeling but become computationally expensive when agents require millions of environment interactions. The hybrid rollout method addresses this limitation by combining learned policy predictions with periodic physics updates, capturing the efficiency of model-based simulation without losing the realism of physics-based environments. In experiments on PandaReach-v3, the framework achieved substantial speedups while maintaining high success rates, highlighting a practical trade-off between accuracy and efficiency. This approach contributes to the broader goal of making reinforcement learning in robotics more computationally scalable and accessible for large-scale experimentation and deployment.