

Humanoid robots performing loco-manipulation tasks must handle dynamic and unpredictable interaction wrenches while maintaining balance. Although Model Predictive Control (MPC) provides a physically consistent stability prior, its effectiveness is limited by an execution gap arising from unmodeled dynamics such as structural compliance, contact mismatch, and tracking delays.

This work proposes a pre-inverse kinematics (pre-IK) task-space residual reinforcement learning framework to bridge this gap. The method augments an MPC-based locomotion planner with a high-bandwidth residual policy that injects corrections into task-space variables before inverse kinematics. This design preserves geometric feasibility while enabling direct adaptation of contact-related behaviors, such as foot placement, ankle orientation, and center of mass (CoM) regulation.

The proposed framework is evaluated in both simulation and hardware experiments under significant interaction disturbances. Results demonstrate that the controller maintains stability under disturbances exceeding the training range (up to ± 160 N) and successfully generalizes to real-world scenarios, including walking with a 10 kg freely swinging payload. These findings highlight the effectiveness of combining model-based control with learning-based adaptation for robust humanoid loco-manipulation.