Towards Humanoid Loco-Manipulation through Learning from Demonstration

Xiaopeng Wu, Austine Oloo, Zainab Altaweel, Xiaohan Zhang*, Shiqi Zhang*

School of Computing, Binghamton University

xwu23@binghamton.edu *Co-advising

Research Goal

How can visual teleoperation data be checked, prepared, and used to train task policies for closed-loop humanoid control?

Loco-manipulation requires coordinated movement, balance, and object interaction. We build a demonstration-learning pipeline as a foundation for acquiring these skills.

Current focus: connect visual teleoperation, demonstration data, and learned policies in a complete closed-loop workflow.

Demonstration-to-Policy Learning Pipeline

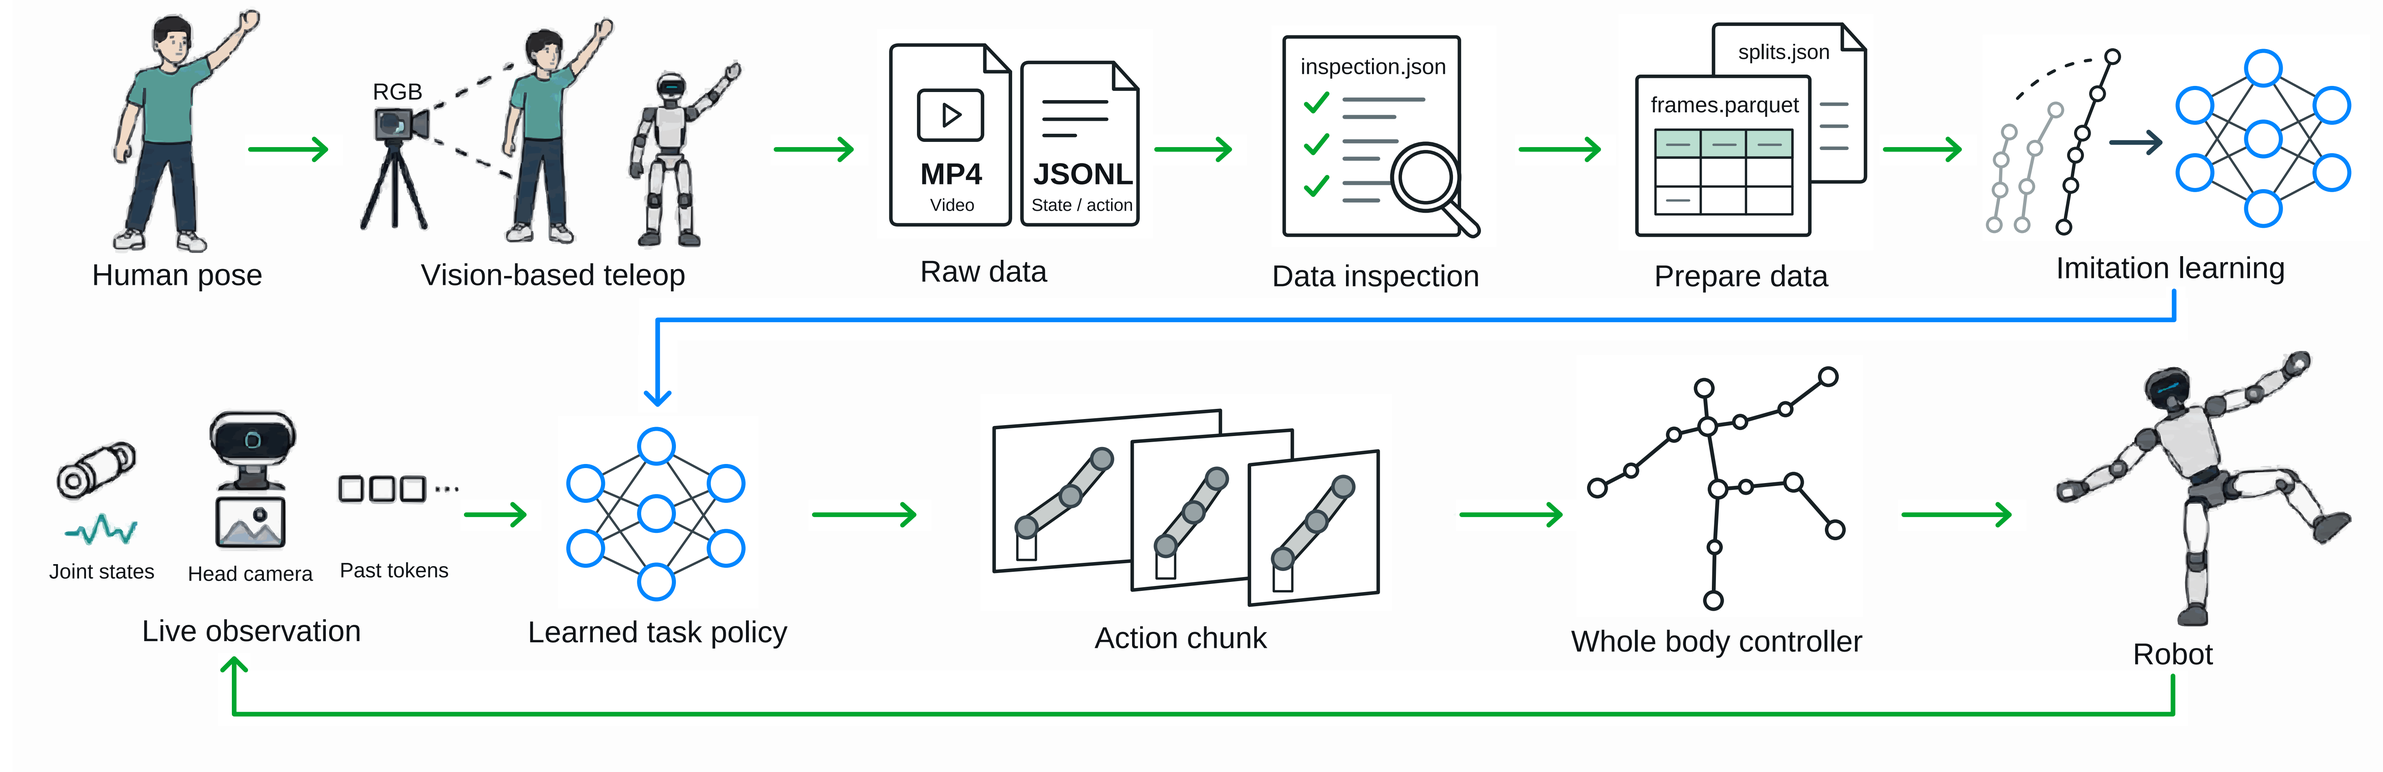

Human demonstration data are collected through vision-based teleoperation, inspected, and prepared for imitation learning. The learned policy uses live observations to predict action sequences. A whole-body controller executes these commands while robot feedback closes the loop.

Teleoperation: Bimanual Box Relocation

A human operator guides the robot to move the yellow box from the left of the black box to its right.

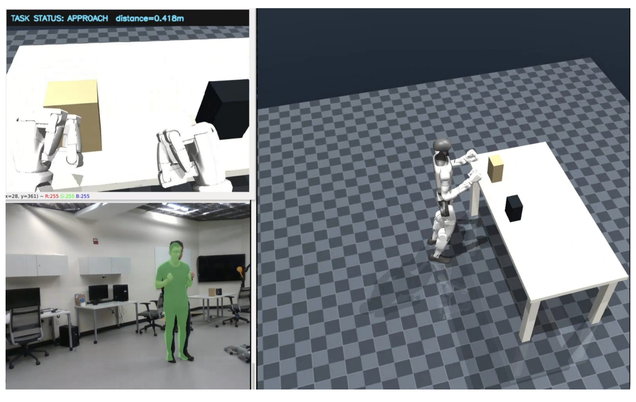

01 Approach

Approach the yellow box

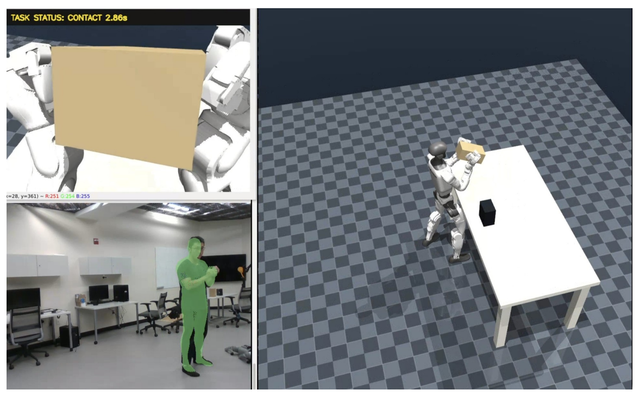

02 Grasp

Hold the box with both hands

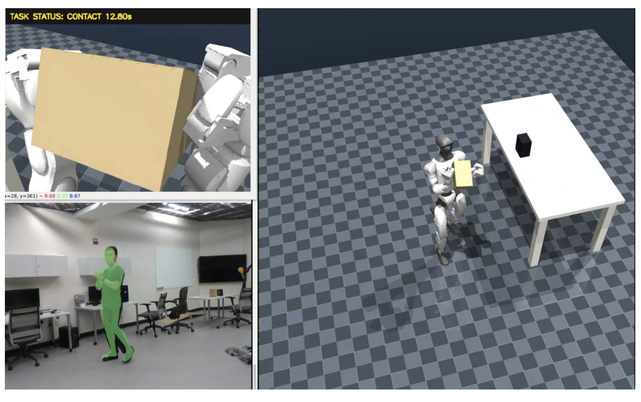

03 Carry

Move while holding the box

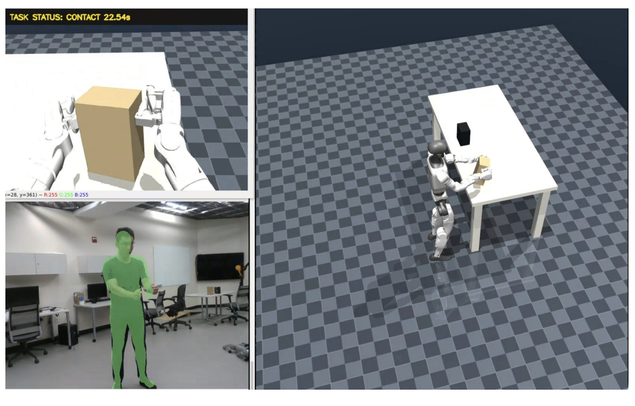

04 Place

Lower the box on the other side

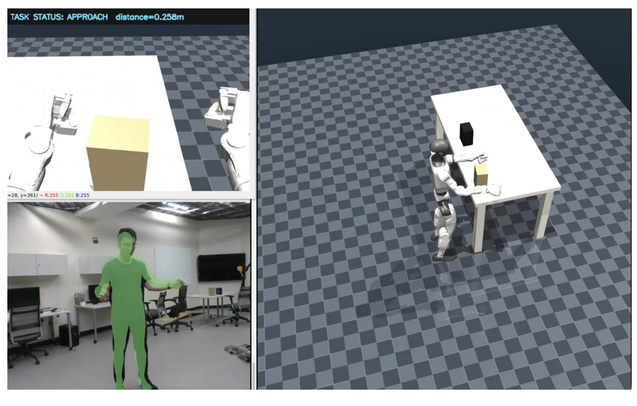

05 Release

Leave the box beside the black box

Expert collection

An RGB camera captures the operator's motion. Vision-based teleoperation translates it into commands for a whole-body controller, enabling coordinated robot movement and object handling.

Recording demonstrations

We record the demonstrated commands alongside camera images and robot and task states for inspection and learning.

Policy Inference: Autonomous Target Interaction

Under the control of the learned policy, the robot approaches and touches the red box, then returns to standing.

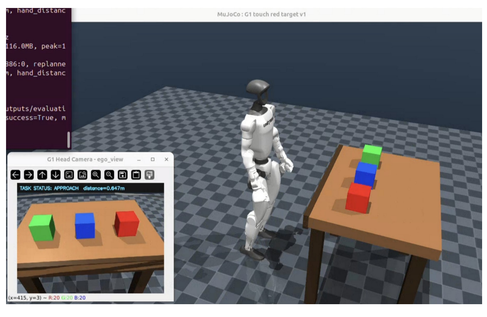

01 Stand

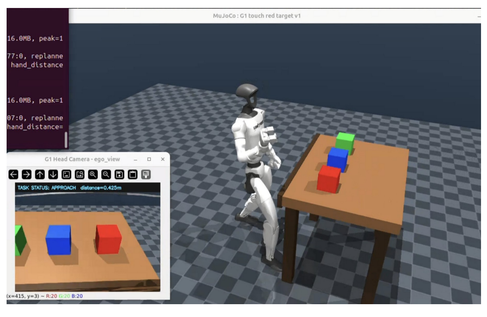

02 Locate & approach

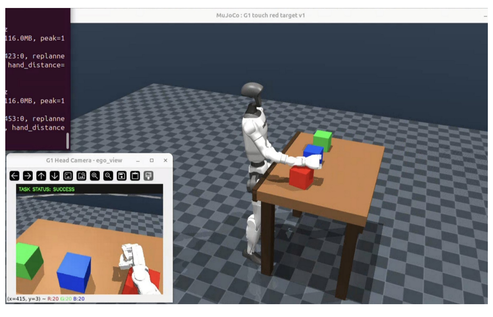

03 Touch the red box

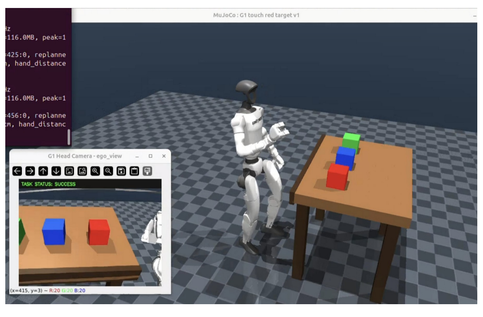

04 Return to standing

Online rollout

The learned policy predicts action sequences from camera images, robot state, task-phase information, and action history. A whole-body controller executes the commands. New observations provide feedback for subsequent predictions.

Conclusion

We connect visual teleoperation data with imitation learning and closed-loop humanoid control.

Initial simulation results demonstrate autonomous target interaction, providing a starting point for learning more complex loco-manipulation tasks.