Oregon State University

Friction-Aware RL for Pick and Place Robots

Roger Wienaah, Jamie Santos and Dax Lynch

Learning Based Control, Final Project Presentation

3rd December, 2025

Motivation

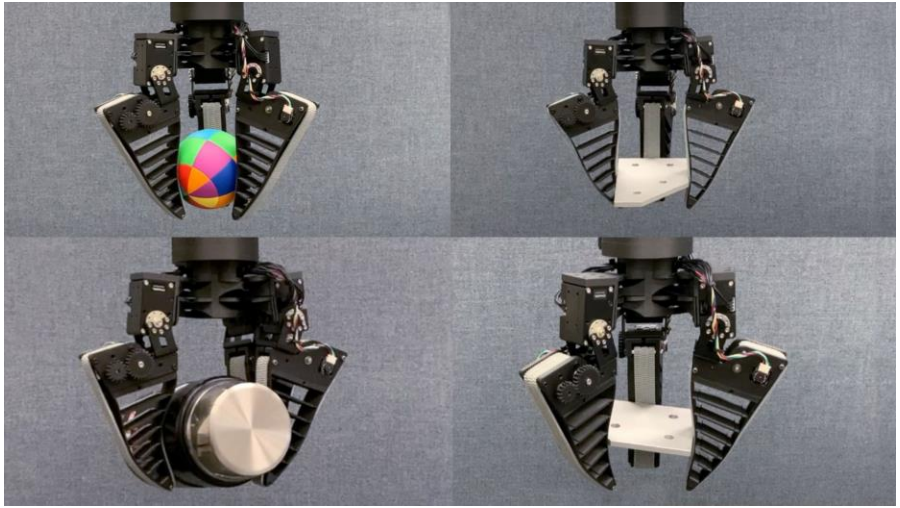


Fig 1. A robot grasping different objects [1]

Real World: objects vary widely in friction and grasp dynamics (e.g., ice vs. rubber).

E.g: Warehouse robots face stochastic textures when picking diverse items (e.g., Amazon fulfillment)

Question: How can an RL controller achieve robustness across materials with different friction dynamics?

RL for Robotic Manipulation

- Manipulator learns and adapts to different object shapes and sizes
- Most methods are vision-based and don’t account for contact dynamics.
- Tactile feedback is crucial for robust manipulation in real-world environments.

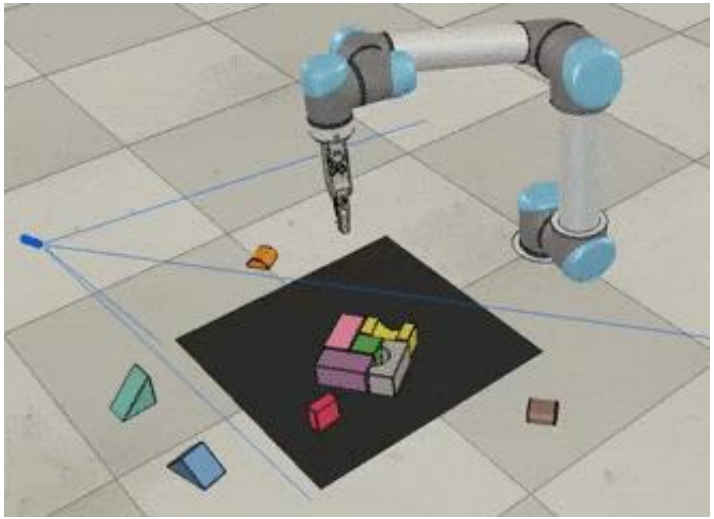


Fig 2. RL-based manipulation with vision [2]

Related Works (Current Solutions and Gaps)

Paper	Approach	Gap
Ding et al, 2021 [3]	DDPG + Tactile feedback	The approach does not **incorporate friction coefficients.**
Zhang et al, 2022 [4]	TD3 + Force Sensors	Focused on geometric alignment rather than material properties **(surface friction/texture)**.
Vulin et al, 2021 [5]	DDPG + Intrinsic reward and force inputs	Focused on exploring the surface of objects than **achieving stable grasp**

TQC + Friction-Awareness

Truncated Quantile Critics (TQC)

- Soft-Actor Critic with a distributional critic
  - Critic outputs multiple quantiles (not single Q-value)
- Uses multiple critic networks outputting quantiles to approximate the return distribution for robust learning.
- The algorithm truncates the most optimistic quantile estimates to reduce bias and improve stability during training.

Why Truncated Quantile Critics (TQC)

- Compared to SAC, TD3, and DDPG:
  - Mitigates overestimation bias in continuous control settings
  - Improved stability and robustness
  - Faster training than other algorithms

- Hyperparameters available

J. Gu, “Application of Robotic Arm Path Planning Based on TQC Algorithm,” vol. 53, pp. 216–220, Dec. 2023, doi: https://doi.org/10.1109/auteee60196.2023.10407890.

Experimental Setup

- Simulator: PandaGym (PyBulllet Physics Engine)
- Robot: Franka Emika Panda robotic arm



Fig 3. PandaGym Environment [6]

Experimental Setup

- Task: Pick and place objects with friction coefficients ranging from 0.05 to 2.0
- State Space and Inputs:
  - End-effector position/velocity
  - Gripper finger width/velocities
  - Object pos./rotation/velocities
  - Friction coefficient (μ)
- Baseline: TQC

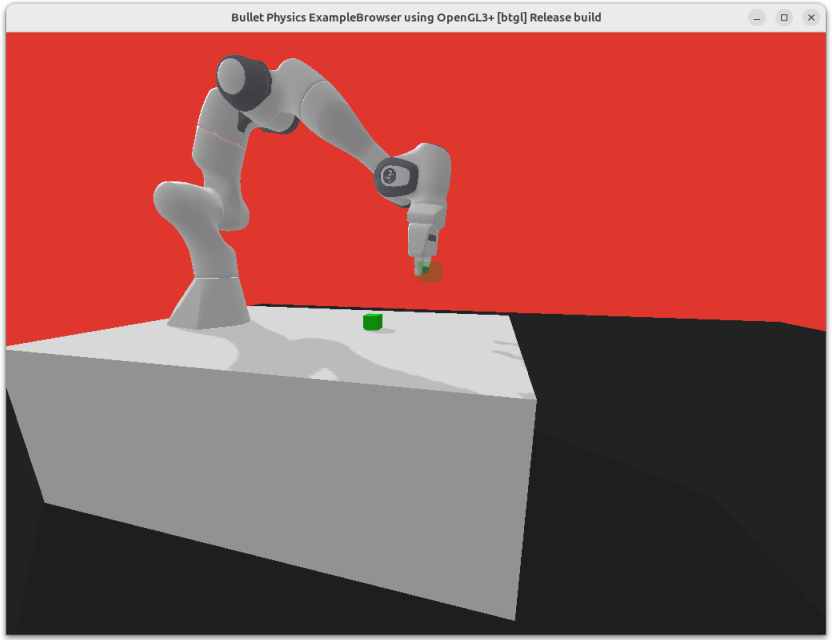


Fig 4. Project Simulation Setup [7]

Reward Structure

- **Dense reward**: receives reward for every timestep
  - Reward: (negative) distance to goal
- **Sparse reward:** receives reward only after reaching and successfully grasping the object
- Use sparse rewards
  - Successful completion prioritized



https://www.freepik.com/premium-vector/heap-gold-coins-clipart-vector-art-illustration_341610889.htm

Experiment 1: Friction Awareness Only

- What if we directly provide friction as input to TQC?
- **Idea:** use "sensor" to measure shear force between gripper and object
- **Hypothesis:** learn more specific gripper width and finger velocities
- Train controller on objects of varying frictions

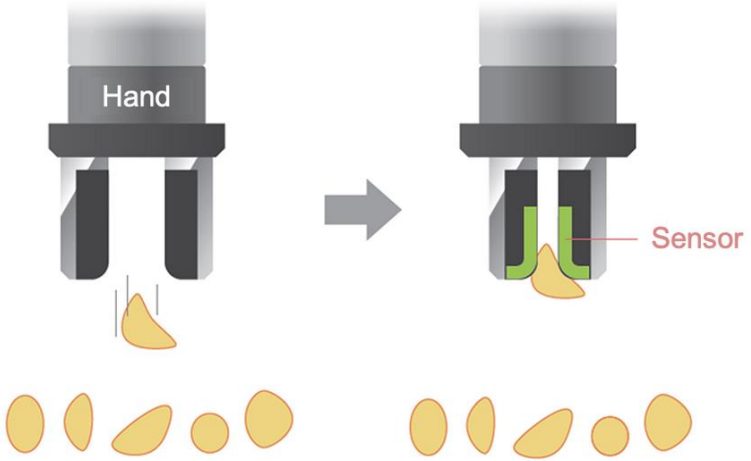
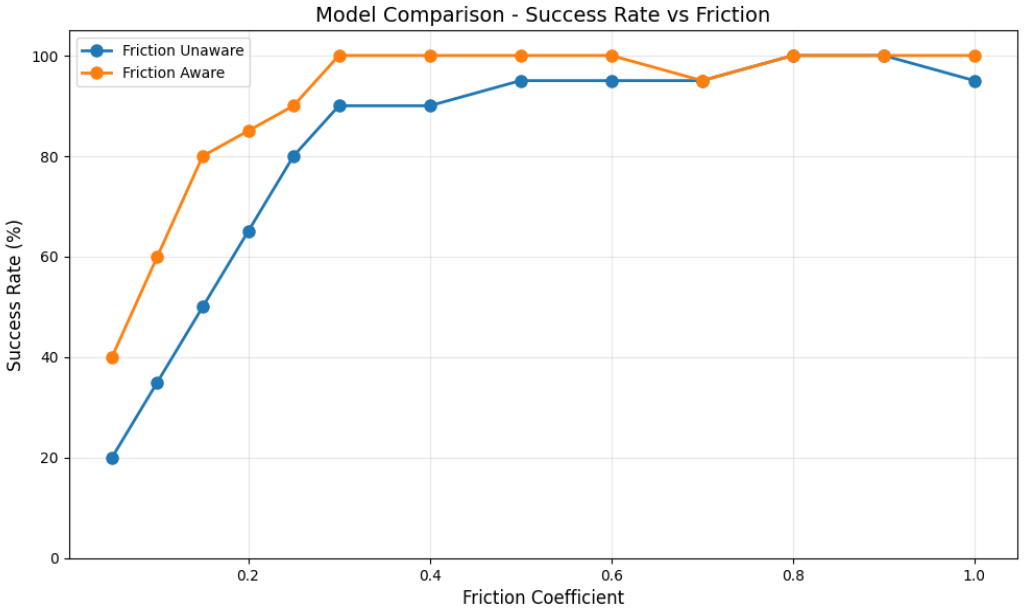


Fig 5. Friction-aware grasping [8]

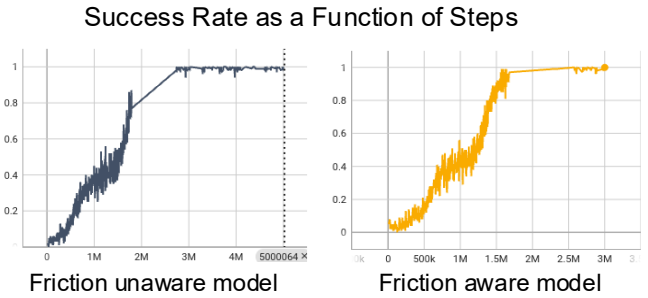
"Film-type tactile sensor," Nisha Co, Ltd. https://connect.nissha.com/filmdevice/en/shear-force-sensor/robot

Oregon State University

Results: Experiment 1, Friction Awareness



- **Friction aware:** μ as input
- **Result:** μ-aware outperforms or matches baseline



Oregon State University

Experiment 2: Train for Joint-Level Control

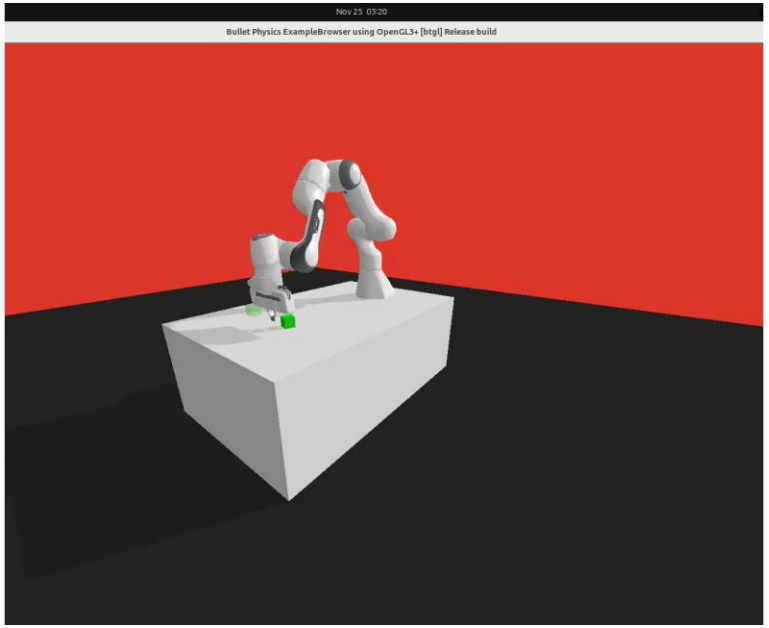
- Two modes of control:
  - End-effector position (inverse kinematics)
  - Joint-level torque commands
- What if we "open up wrist" via torque commands?
- Hypothesis: 6 DoF wrist -> more gripping strategies to explore



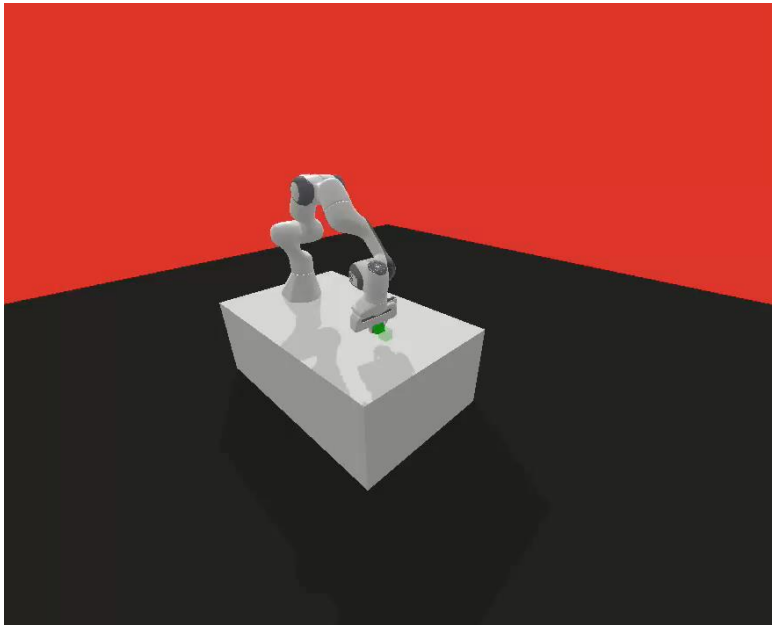
https://en.wikipedia.org/wiki/Gimbal

Results: End Effector vs. Joint-Level Control

Delta End Effector Mode



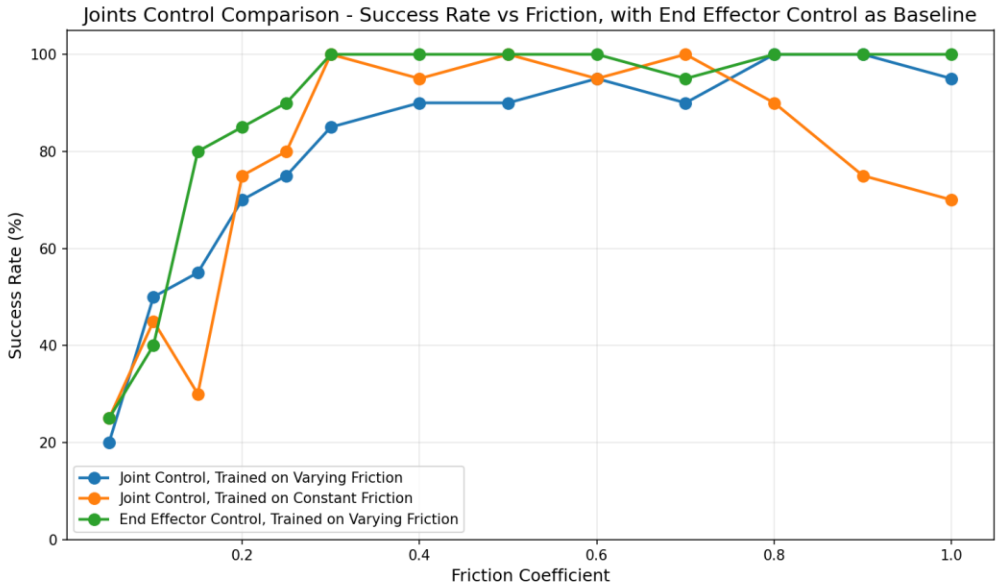
Joint Control Mode



*Note the constant orientation of the end effector

Oregon State University

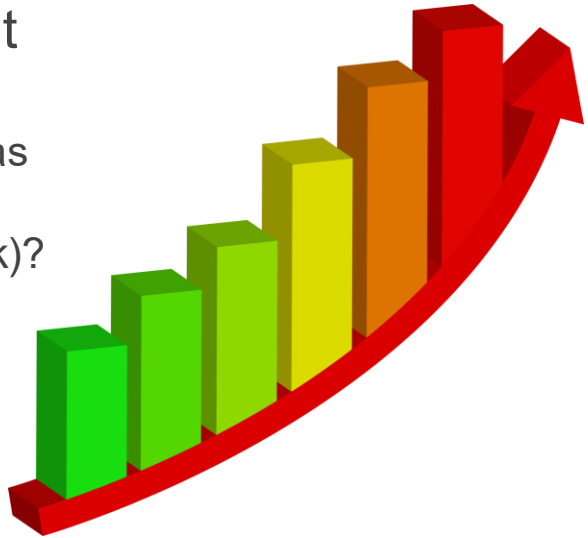
Results: Experiment 2, Joint-Level Control



- **Baseline:** μ-aware, end-effector control
- **Test:** μ-aware, joint-torque outputs
- **Results:**
  - Joint-level able to grasp from different approach angles (wrist not "locked")
  - Currently similar or slightly worse result, slower convergence
  - Still running tests, we hope to show how the grasping strategies change

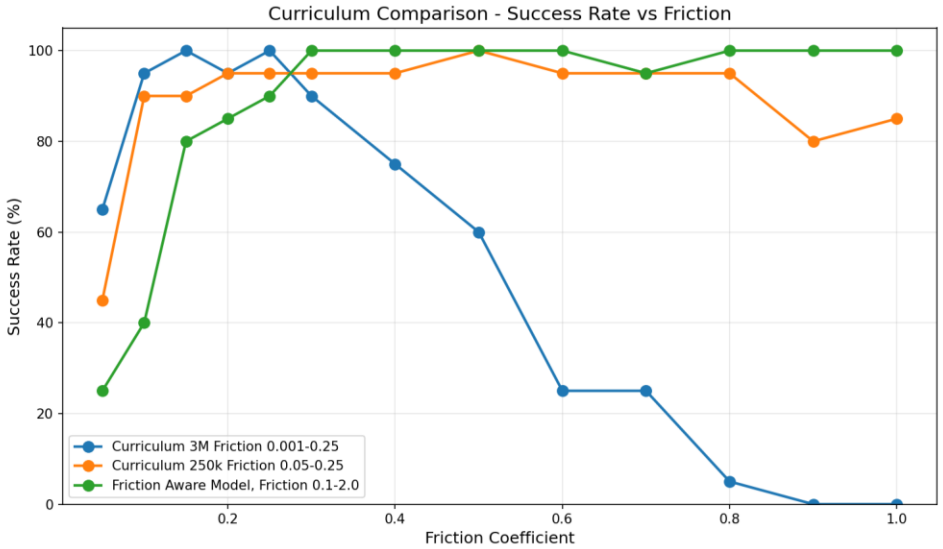
Experiment 3: Implement Training Curriculum

- Start w/ controller trained on varying friction
- Level up difficulty: keep training w/ lowest range of friction coefficients
  - What if we train for the same amount of time as before (~3 million time steps)?
  - What if we train for a fraction of the time (250k)?



https://www.vecteezy.com/free-png/3d-increase-arrow

Results: Experiment 3, Training Curriculum



- Baseline: μ-aware
- Test 1: train 3M time steps
  - μ = [0.001, 0.25]
- Test 2: train 250k time steps
  - μ = [0.05, 0.25]
- Result:
  - 3M: much improved at low levels, fails after -> overfitted
  - 250k: still much improved over μ-aware, **generalizes much better**

Conclusion

Experiment	Conclusion
μ-aware TQC	• Only adding μ as input helps to increases robustness and convergence speed
Joint-Level Control	• ~50% longer training time • No corresponding performance benefit, YET
Training Curriculum	• Simplest, most effective change • Can likely be fine-tuned to produce more robust results across entire range

Future Work

- Implement custom reward using friction bonus
  - Reward proportional to inverse of friction (lower friction, higher reward)
- Improve training curriculum
  - Tune parameters to flatten success curve near 100%

Questions?

References

1. SRI International. “SRI Showcases Its High Dexterous Robot Grasper and Hand - SRI.” SRI, 25 Apr. 2023, www.sri.com/press/story/sri-showcases-its-high-dexterous-robot-grasper-and-hand/. Accessed 3 Dec. 2025.
2. A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez and T. Funkhouser, "Learning Synergies Between Pushing and Grasping with Self-Supervised Deep Reinforcement Learning," 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
3. Z. Ding, Y. -Y. Tsai, W. W. Lee and B. Huang, "Sim-to-Real Transfer for Robotic Manipulation with Tactile Sensory," 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
4. Dong, Siyuan & Jha, Devesh & Romeres, Diego & Kim, Sangwoon & Nikovski, Daniel & Rodriguez, Alberto. (2021). Tactile-RL for Insertion: Generalization to Objects of Unknown Geometry. 6437-6443. 10.1109/ICRA48506.2021.9561646.
5. Vulin, Nikola & Christen, Sammy & Stevsic, Stefan & Hilliges, Otmar. (2021). Improved Learning of Robot Manipulation Tasks Via Tactile Intrinsic Motivation. IEEE Robotics and Automation Letters. PP. 1-1. 10.1109/LRA.2021.3061308.
6. “Film-type tactile sensor,” Nisha Co, Ltd. https://connect.nissha.com/filmdevice/en/shear-force-sensor/robot