Massachusetts Institute of Technology

Empirical Analysis of Sim-and-Real Cotraining

Adam Wei, Abhinav Agarwal, Boyuan Chen, Rohan Bosworth, Nicholas Pfaff, and Russ Tedrake

IROS
21 October, 2025

Adam Wei

Robot Data Diet

How can we obtain data for robot imitation learning?

Big data

Big transfer gap

Ego-Exo

Open-X

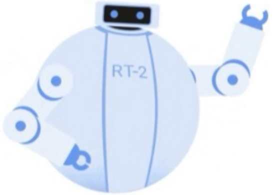

Small data

No transfer gap

Common Crawl

YouTube

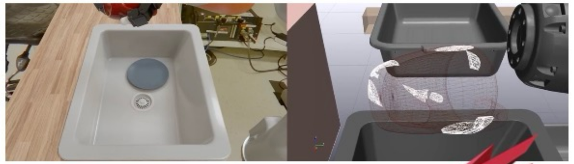

simulation

robot teleop

This talk!

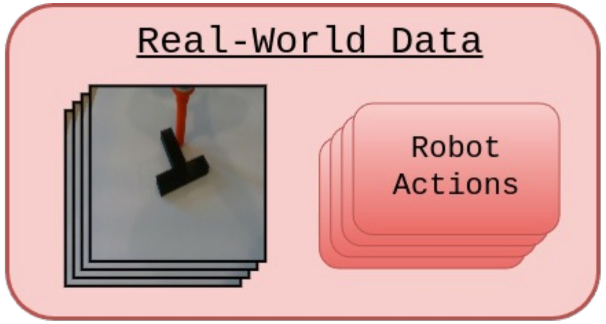

Sim-and-Real Cotraining

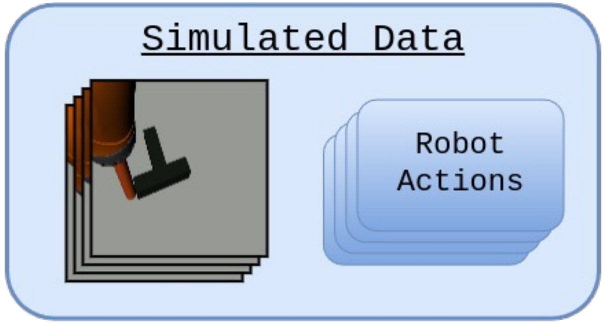

Cotraining: use ***both datasets*** to train a ***model*** that maximizes a real-world ***performance objective***

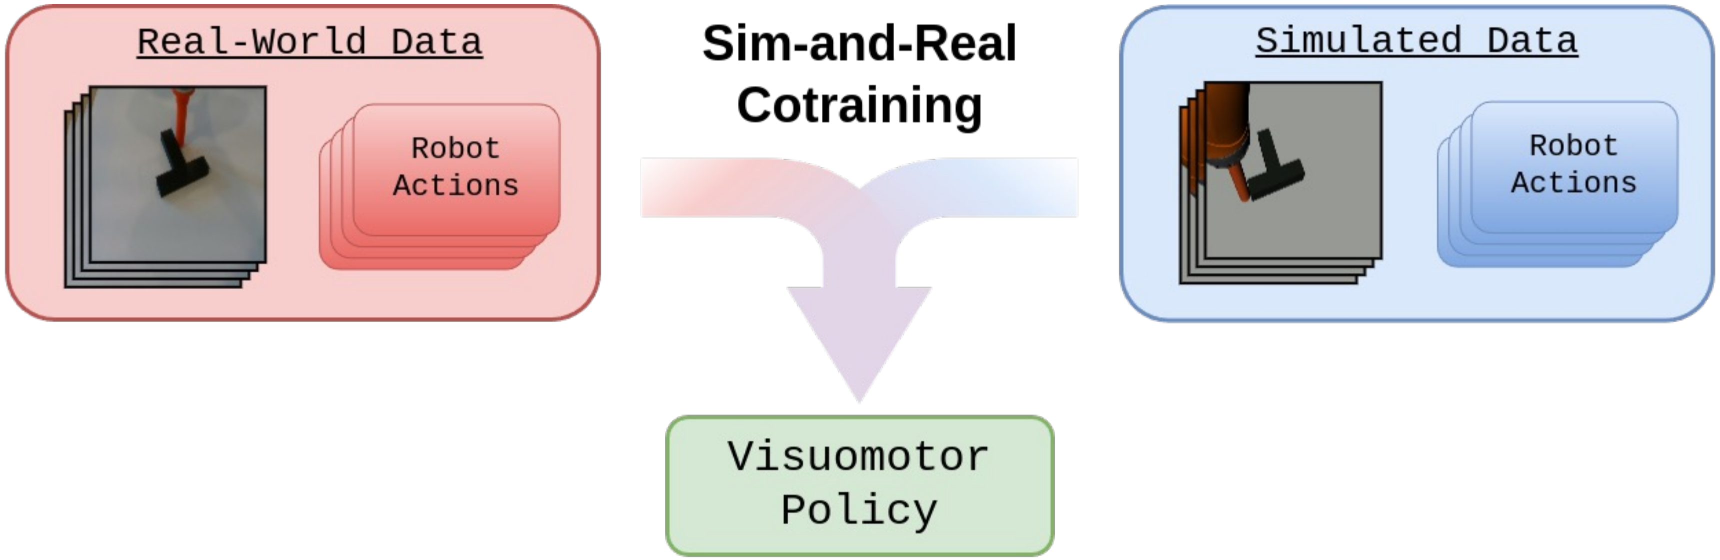

Cotraining: use ***both datasets*** to train a ***model*** that maximizes a real-world ***performance objective***

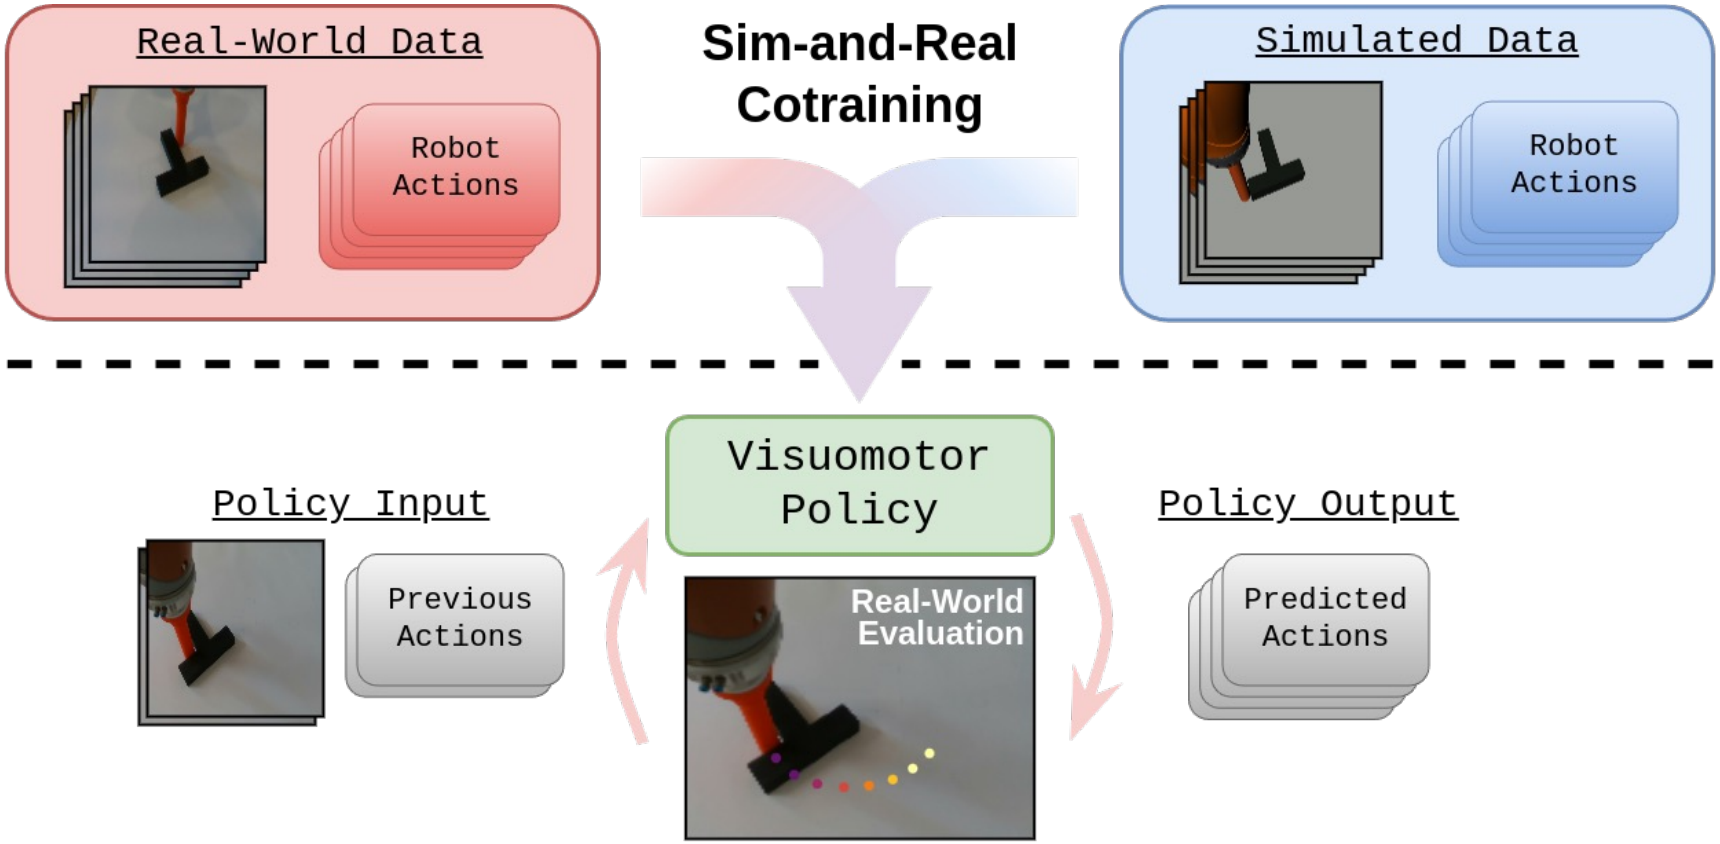

Cotraining: use ***both datasets*** to train a ***model*** that maximizes a real-world ***performance objective***

Cotraining: use ***both datasets*** to train a ***model*** that maximizes a real-world ***performance objective***

Performance Objective:
Success rate on planar pushing from pixels

Focusing on a single canonical task enables controlled and thorough analysis

Cotraining: use ***both datasets*** to train a ***model*** that maximizes a real-world ***performance objective***

Performance Objective:

Success rate on planar pushing from pixels

Model:

Diffusion Policy

$$\mathcal{L}_{\mathcal{D}^{\alpha}} = \alpha \mathcal{L}_{\mathcal{D}_R} + (1-\alpha)\mathcal{L}_{\mathcal{D}_S}$$

Real-World Dataset:

Simulated Dataset:

Model-based Motion Planner

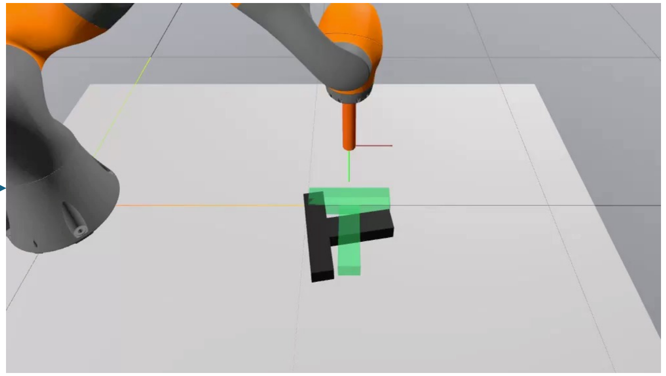

sim2real gap!

Does Cotraining Improve Performance?

Policy trained with
50 real demos, 0 sim demos

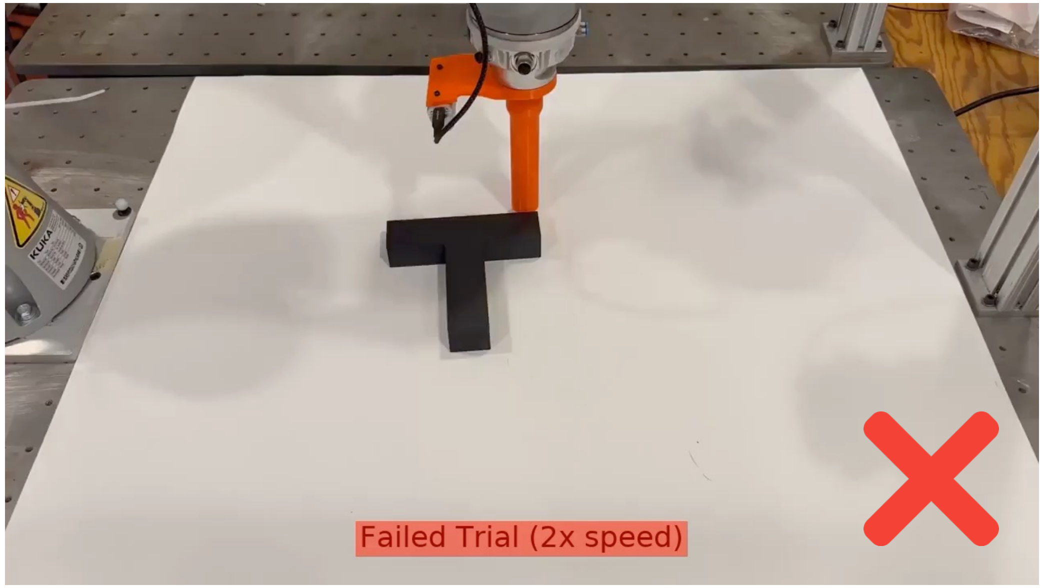

Success rate: **10/20**

Policy cotrained with
50 real demos, 2000 sim demos

Success rate: **18/20**

1.8x improvement!

Does Cotraining Improve Performance?

Policy trained with
10 real demos, 0 sim demos

Policy cotrained with
10 real demos, 2000 sim demos

*Success rate: **2/20*** → *Success rate: **14/20***

7x improvement!

Key Takeaways: Performance Gains

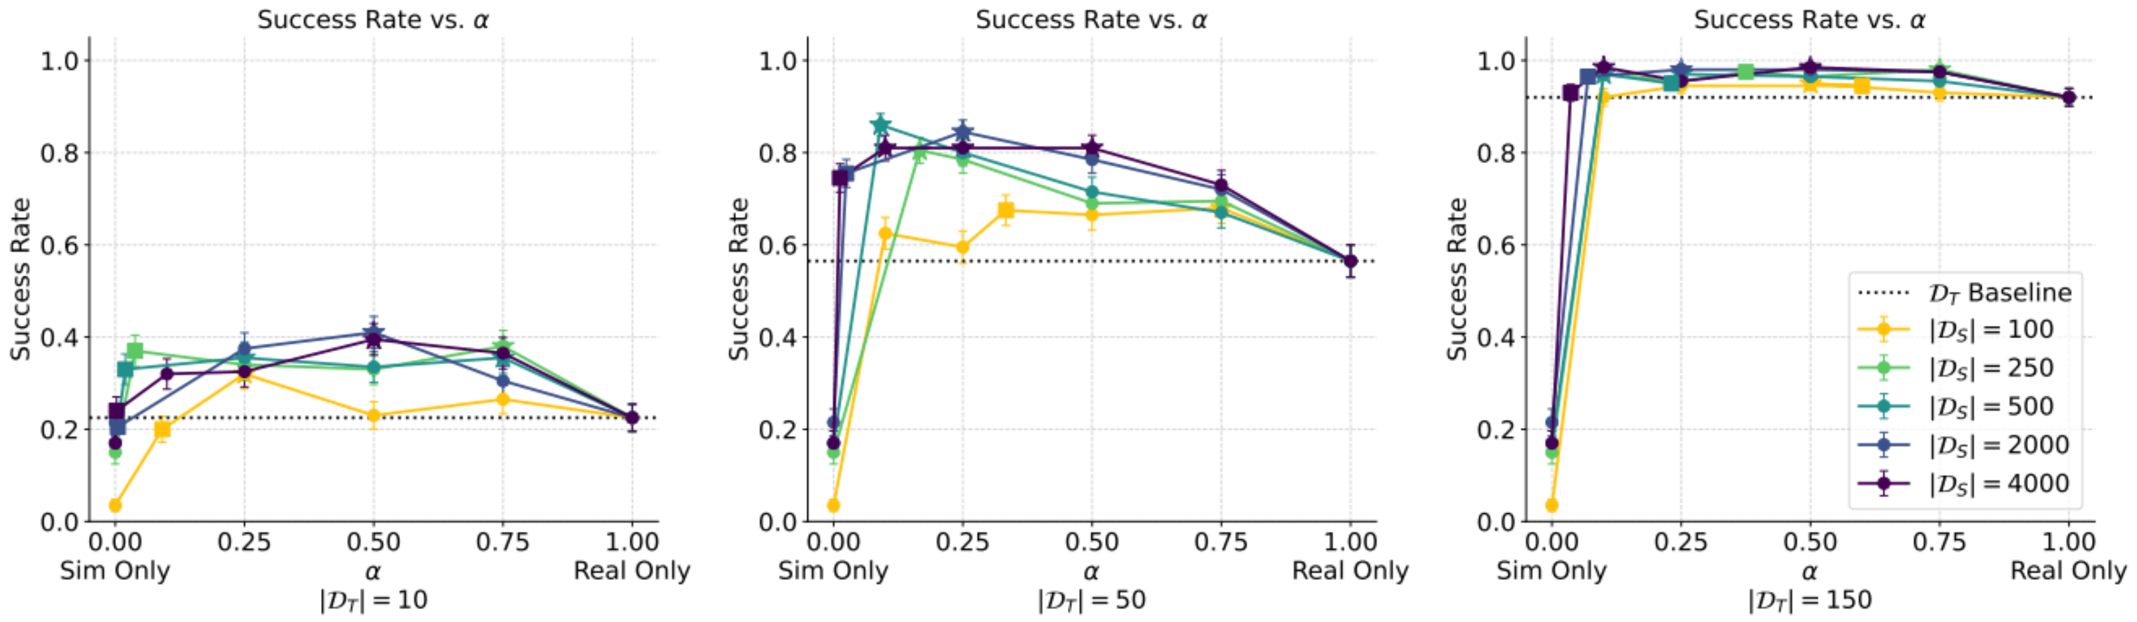

- Cotraining improved performance *up to 7x!*
- Cotraining is most effective in the *low to medium* data regime.
- Scaling simulated data alone is insufficient!

The Effect of Sim2Real Gaps (i.e. distribution shifts)

Which *sim2real gaps* affect the *value* of simulated data?

Example: Analyzing Color Shift

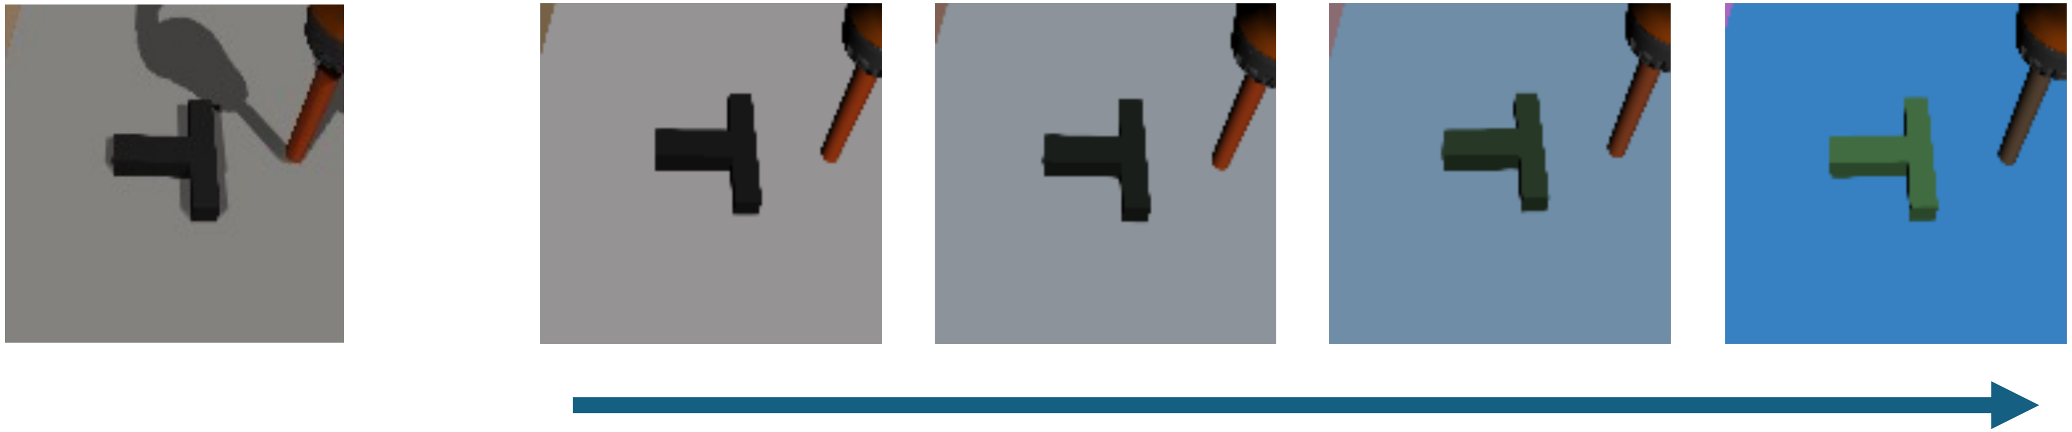

“True” Color

Increasing “Color Shift”

Ex. Analyze policies trained on increasing intensities of color shift

The Effect of Sim2Real Gaps (i.e. distribution shifts)

Visual Gaps
Color Shift
Color Randomization
Camera Pose Shift

Physical Gaps
Center of Mass Shift

Task Gaps
Target Shift
Object Mismatch

Key Takeaways: Sim2Real Gaps

- *Cotraining still improves performance…*
  but all gaps reduce the value of sim data

- Physics & task gaps were most impactful

- *Better rendering* improves performance, but *perfect rendering* hurts performance!

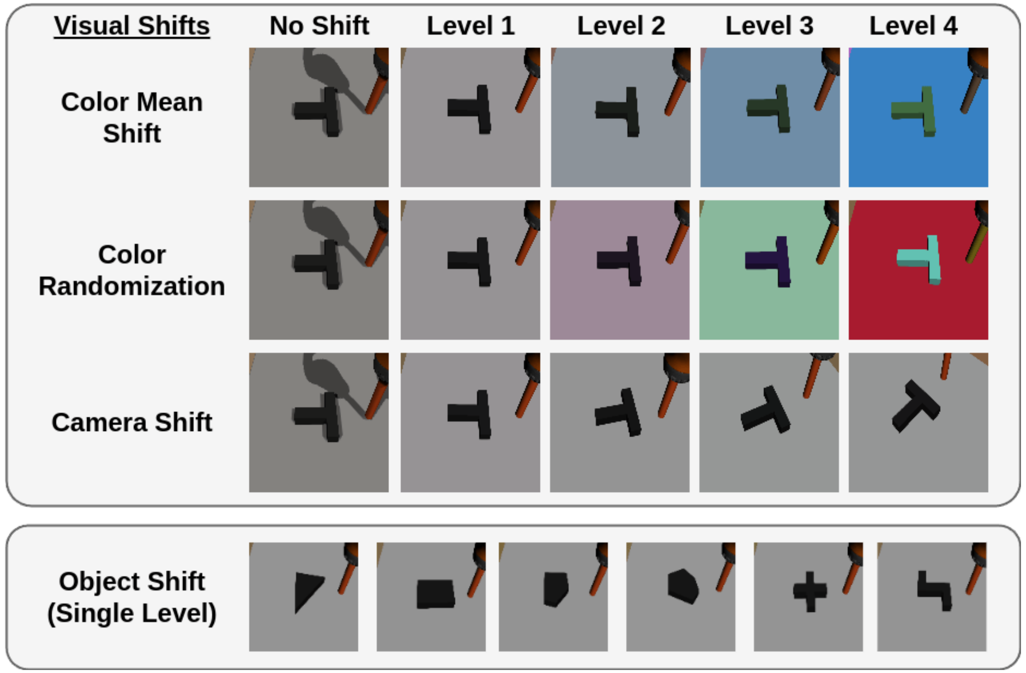

Sim vs Real "Expert"

Real-World Demos

- Fixes orientation first, then translation

Sim Demos

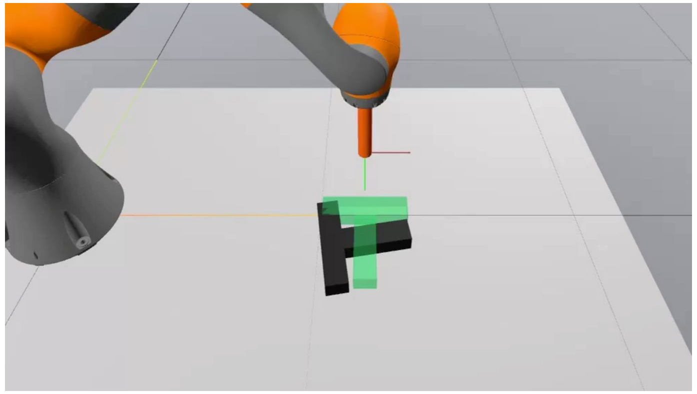

- Fixes orientation and translation *simultaneously*

Sim & real behaviors are characteristically different

Sim vs Real "Expert"

Distinctly more similar to real-world expert!

Real-World Demos

50 demos

Sim Demos

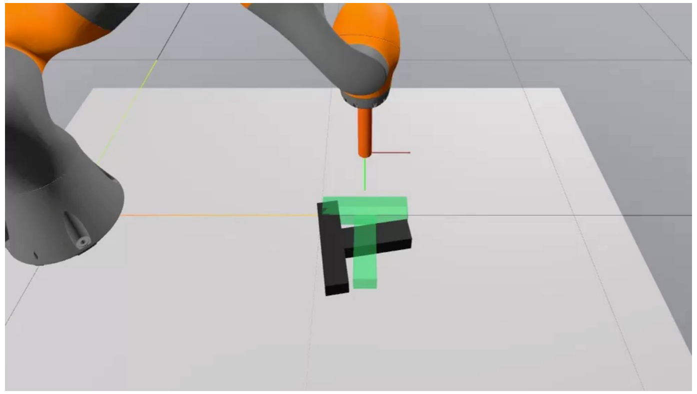

2000 demos

Cotrained Policy

Binary probes show that policies are learning to **distinguish sim from real**!

Sim & Real Discernability

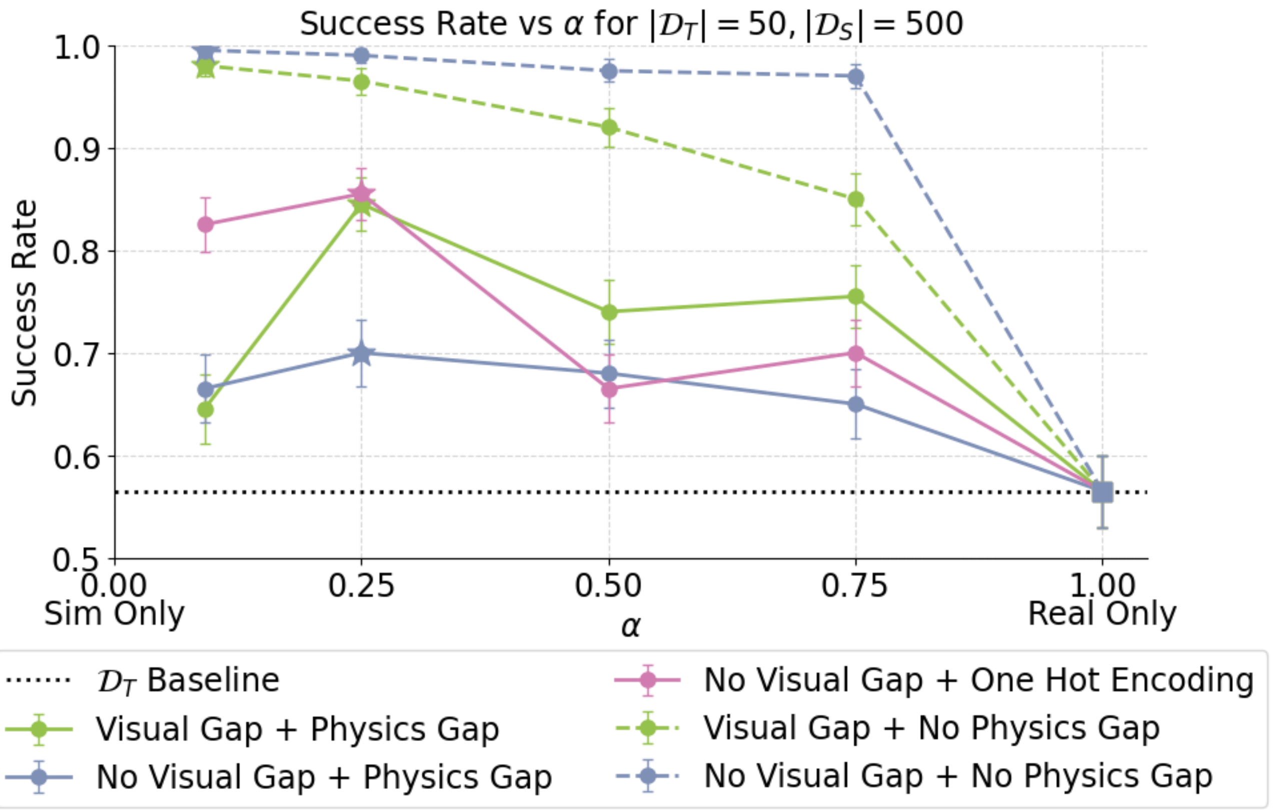

High-performing policies must learn to *identify sim vs real*

since the *physics* of each environment *requires different actions*

Positive Transfer: Scaling Law

Scaling sim data improves real-world test loss according to a power law!

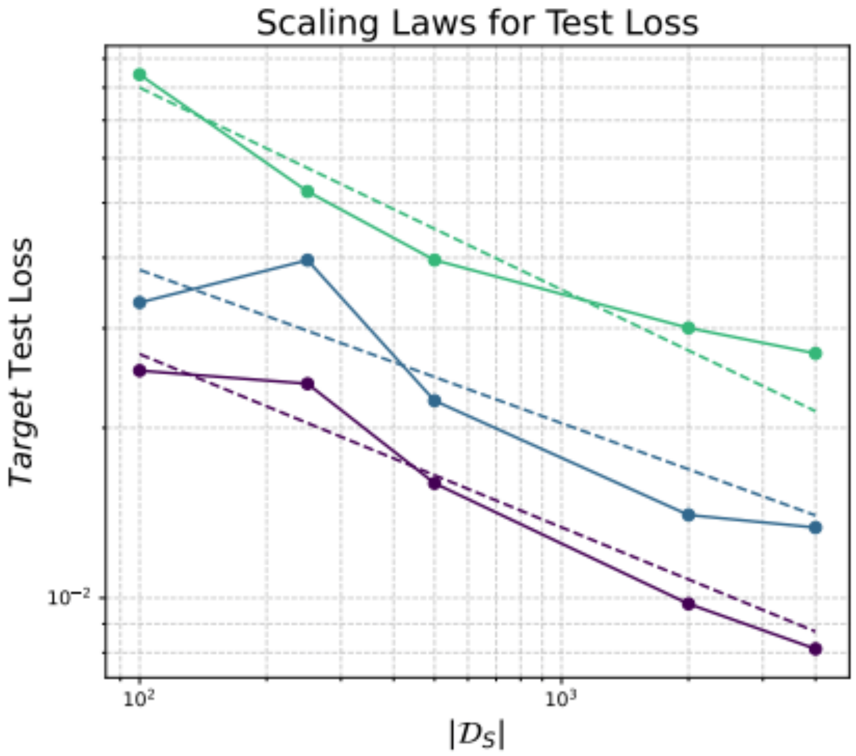

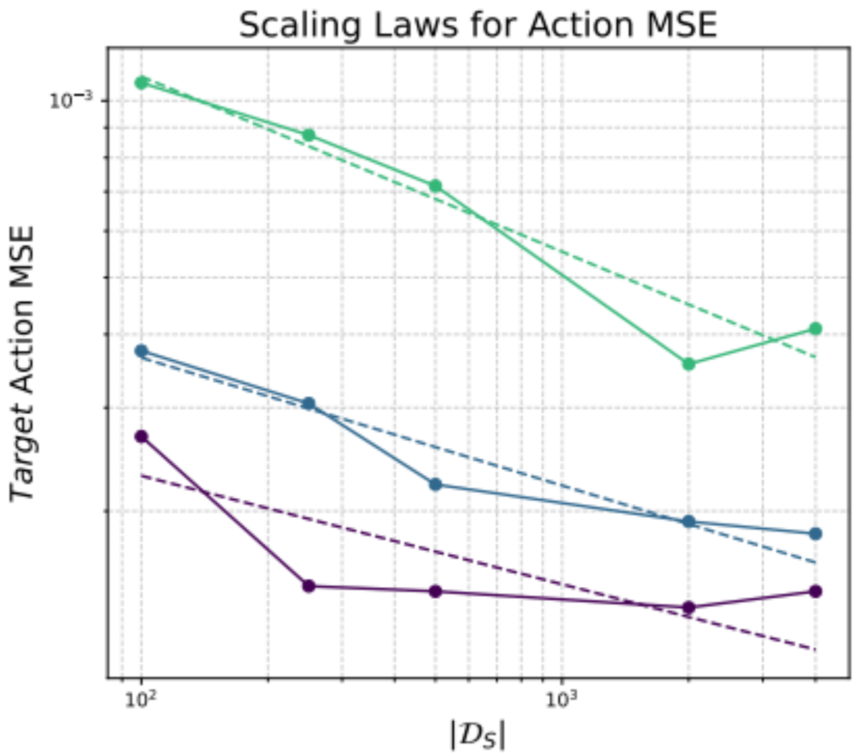

$$\mathcal{L}_{\mathcal{D}_T^{\text{test}}} \propto |\mathcal{D}_S|^{-0.332} \cdot |\mathcal{D}_T|^{-0.397}, \quad R^2 = 0.945$$
$$\text{MSE}_{\mathcal{D}_T^{\text{test}}} \propto |\mathcal{D}_S|^{-0.285} \cdot |\mathcal{D}_T|^{-0.587}, \quad R^2 = 0.975$$

Positive Transfer: Scaling Law

Scaling sim data improves real-world test loss according to a power law!

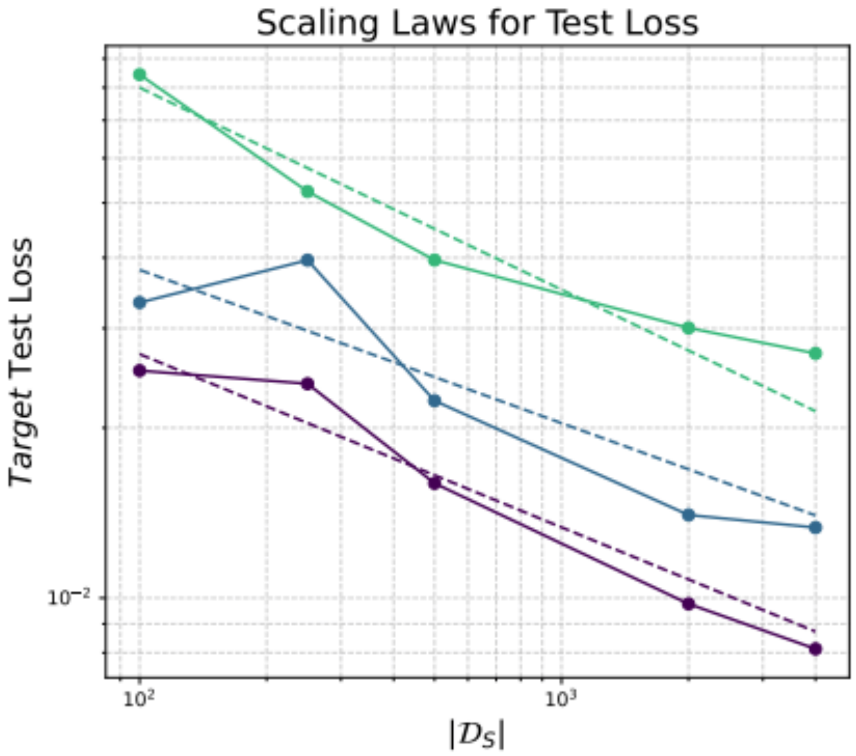

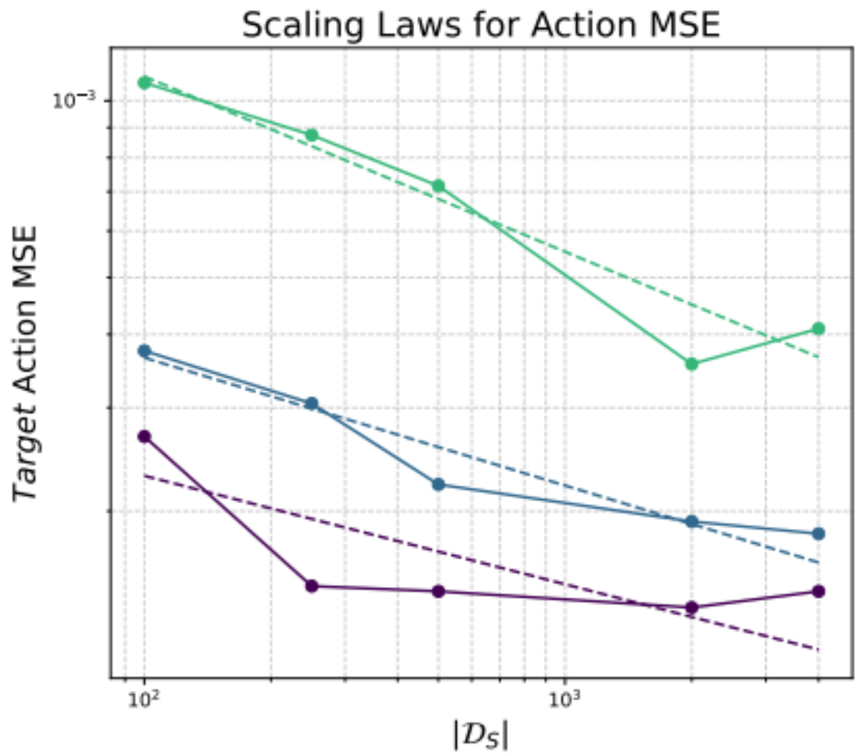

$$\mathcal{L}_{\mathcal{D}_T^{\text{test}}} \propto |\mathcal{D}_S|^{-0.332} \cdot |\mathcal{D}_T|^{-0.397}, \quad R^2 = 0.945$$

$$\text{MSE}_{\mathcal{D}_T^{\text{test}}} \propto |\mathcal{D}_S|^{-0.285} \cdot |\mathcal{D}_T|^{-0.587}, \quad R^2 = 0.975$$

A sim demo is worth ~0.5-0.8 real demos

Empirical Analysis of Sim-and-Real Cotraining

- Simulation is a promising tool for *scaling data generation in robotics*
- We study the *principles* and *mechanisms* of cotraining from both sim and real data

Our Paper
Scan to learn more!

Personal Website (Adam Wei)
Feel free to reach out!