SURF Version

Teaching Robots to Write

Joel Castro, Austin Patel, Shuran Song

Background & Motivation

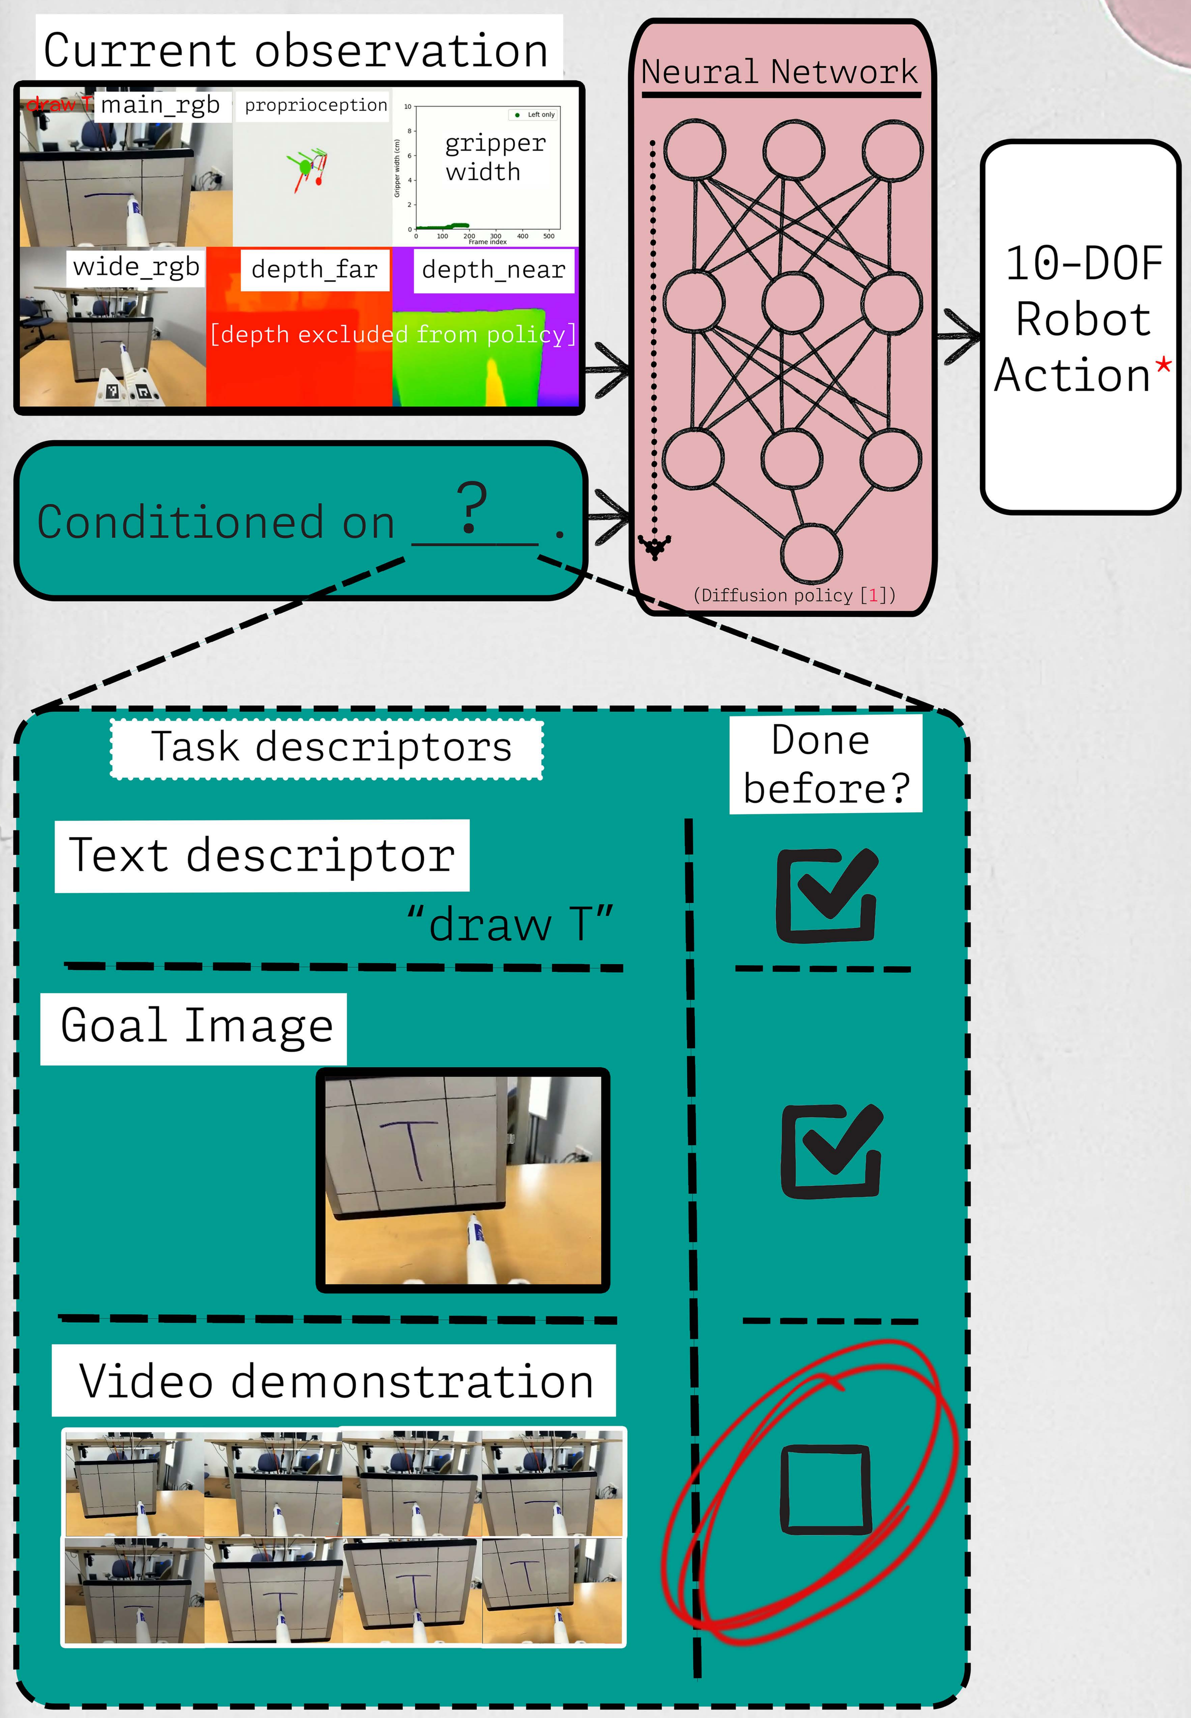

The drawing task will showcase benefits of this novel policy architecture.

Whiteboard Task

Hardware design (A)

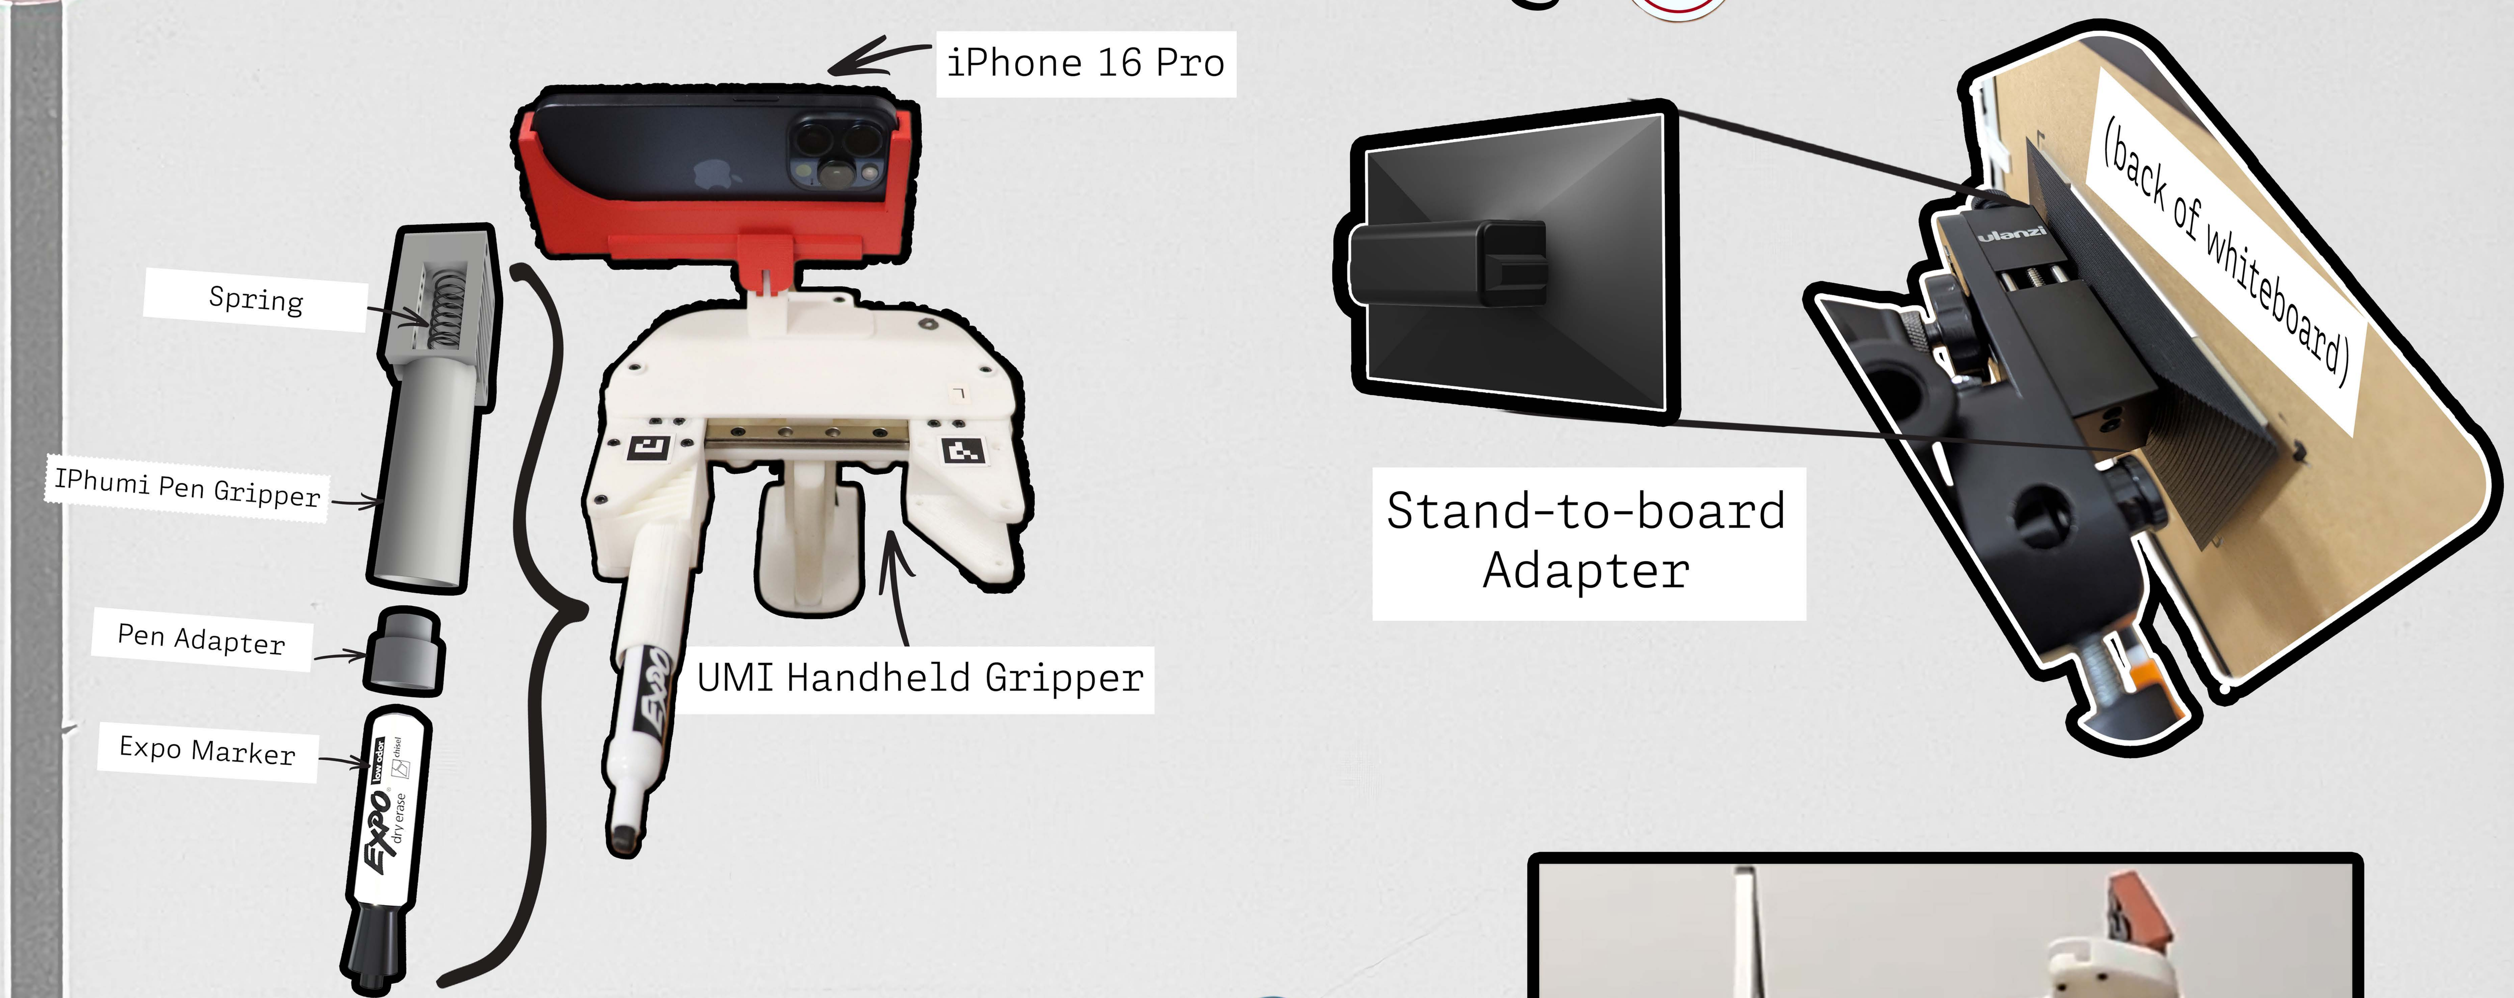

Collecting demonstrations (B)

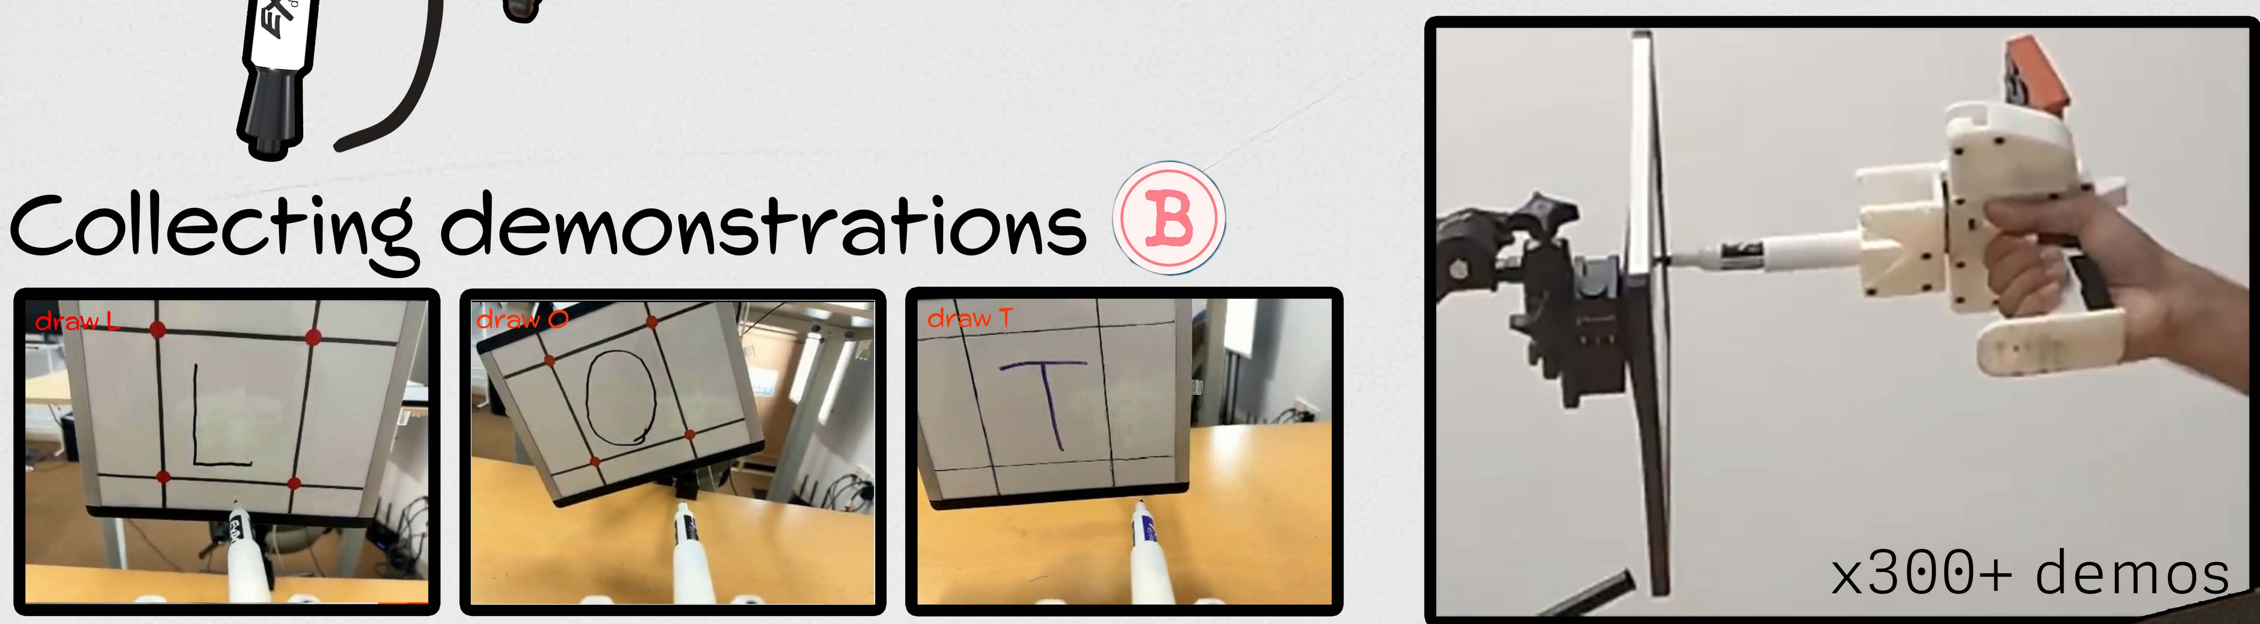

NN Training (C)

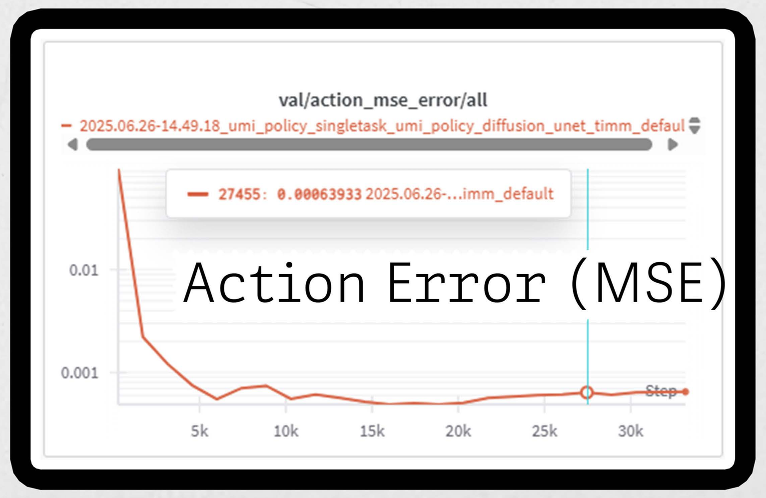

Rollout: hand policy control of robot (D)

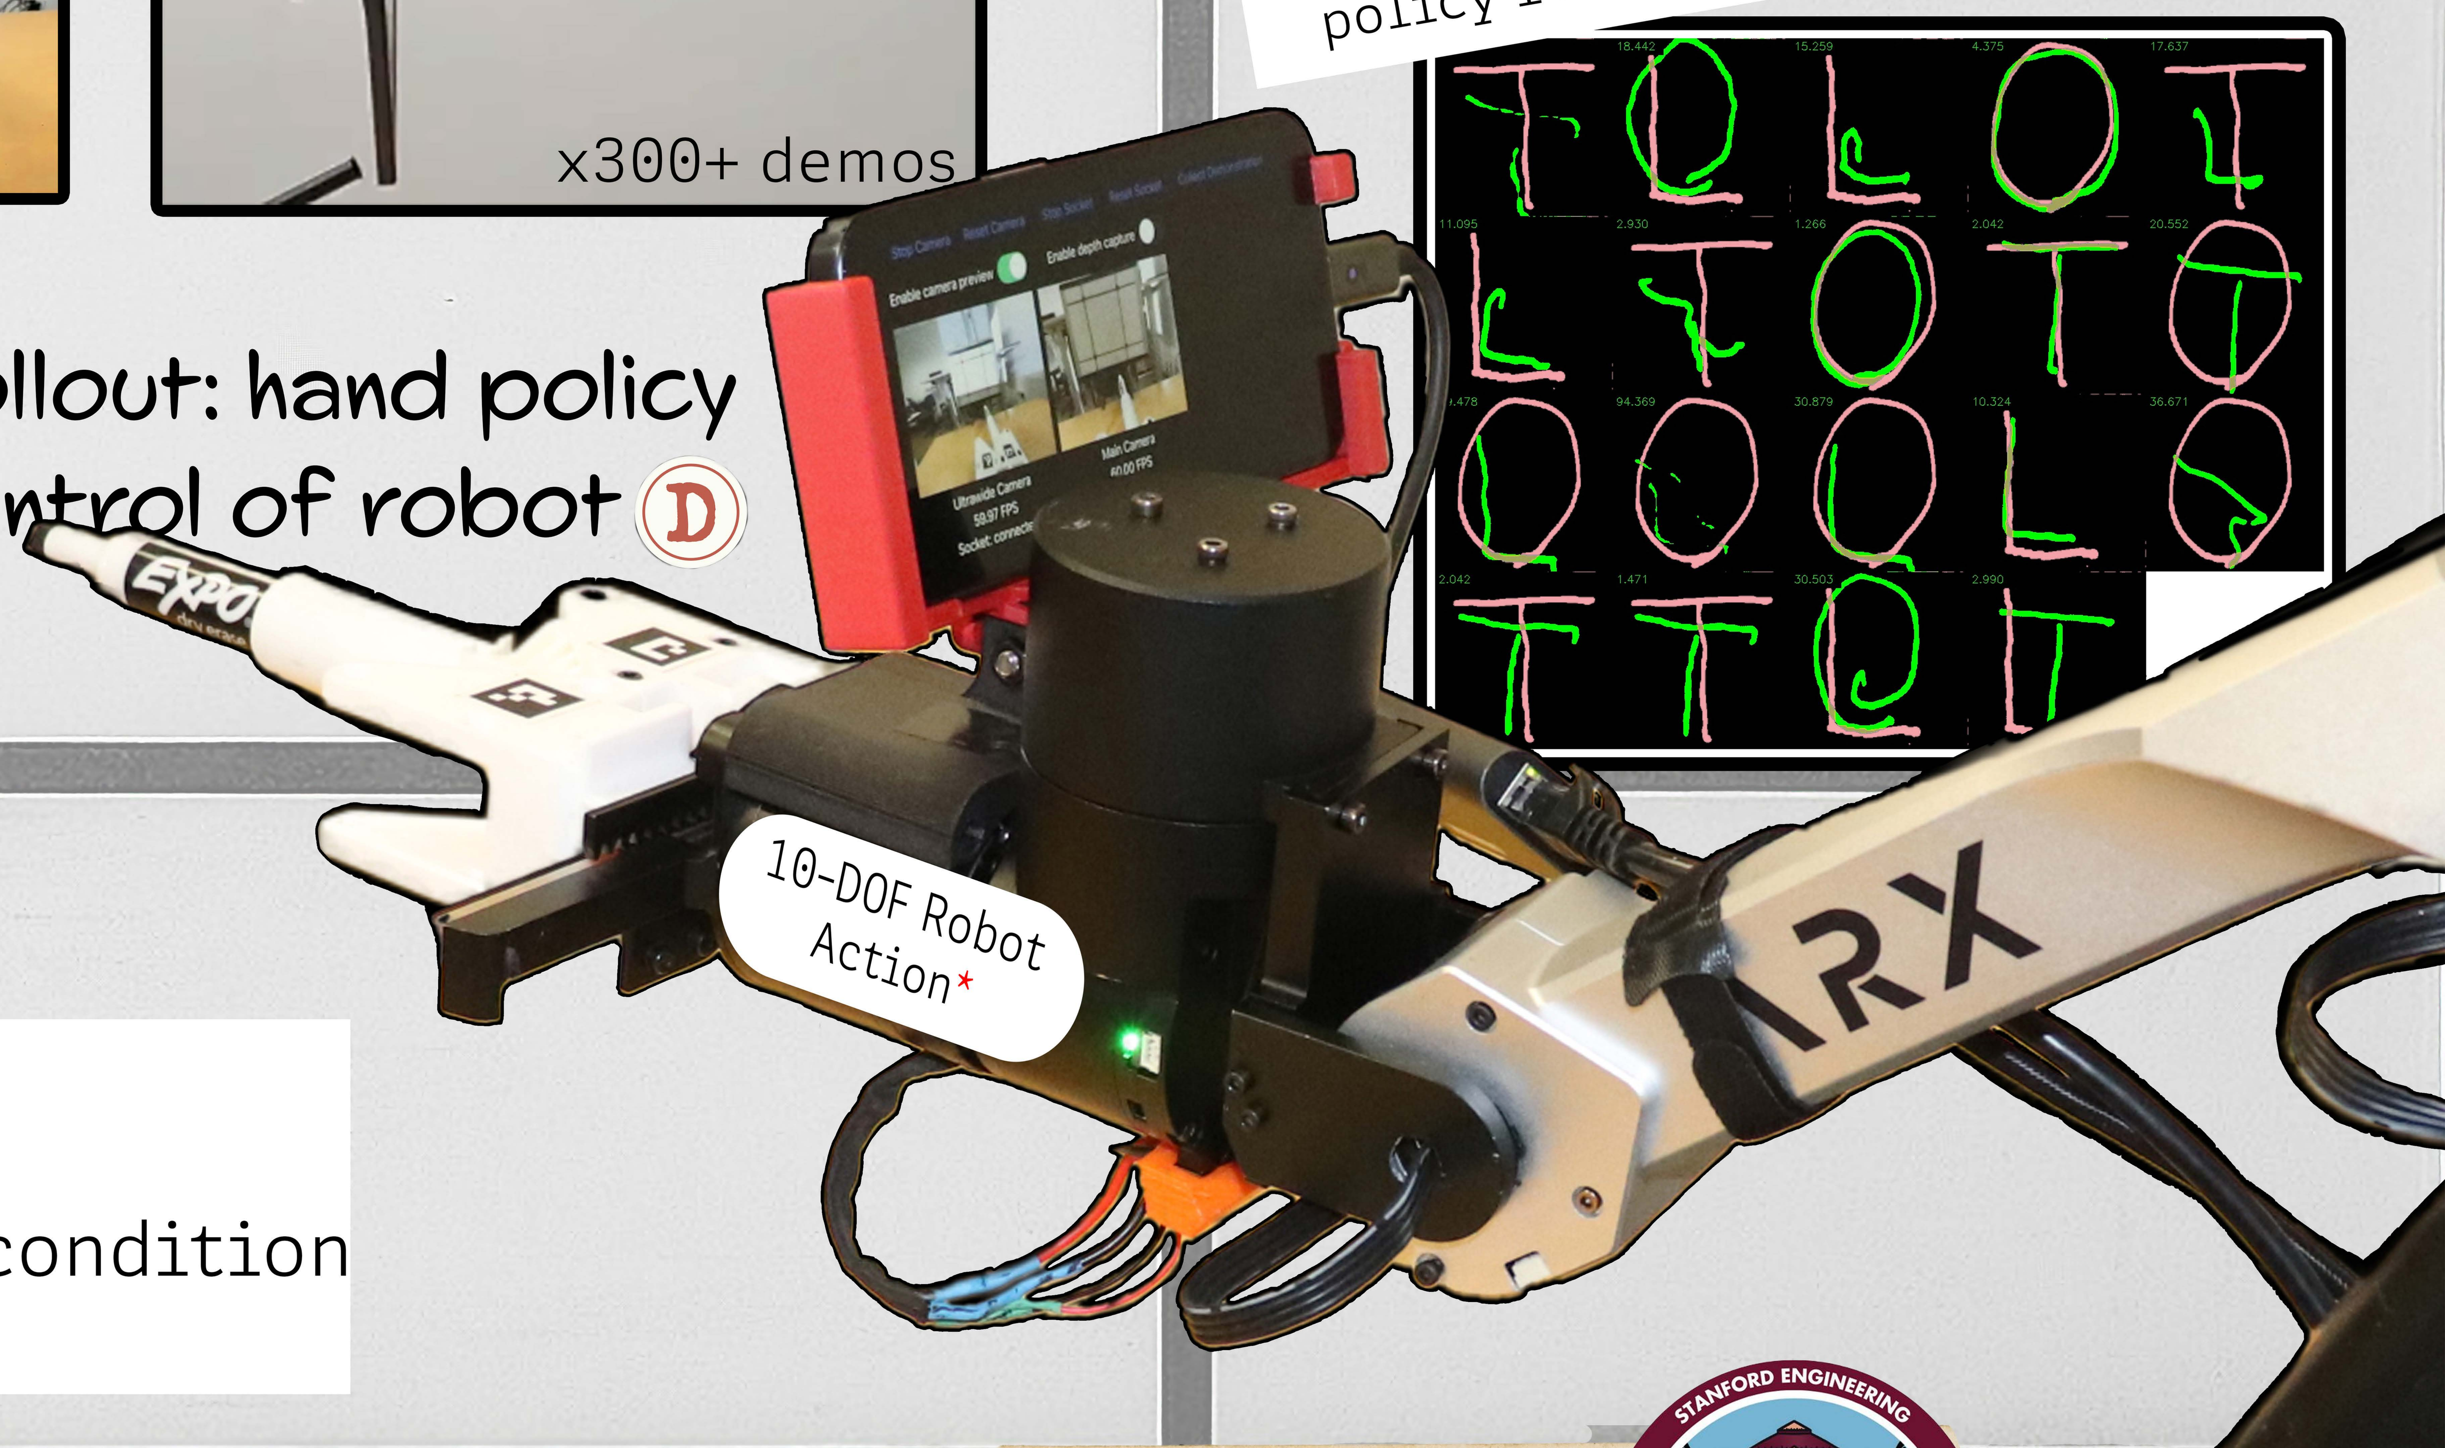

Evaluation Metric

Scoring drawings (E)

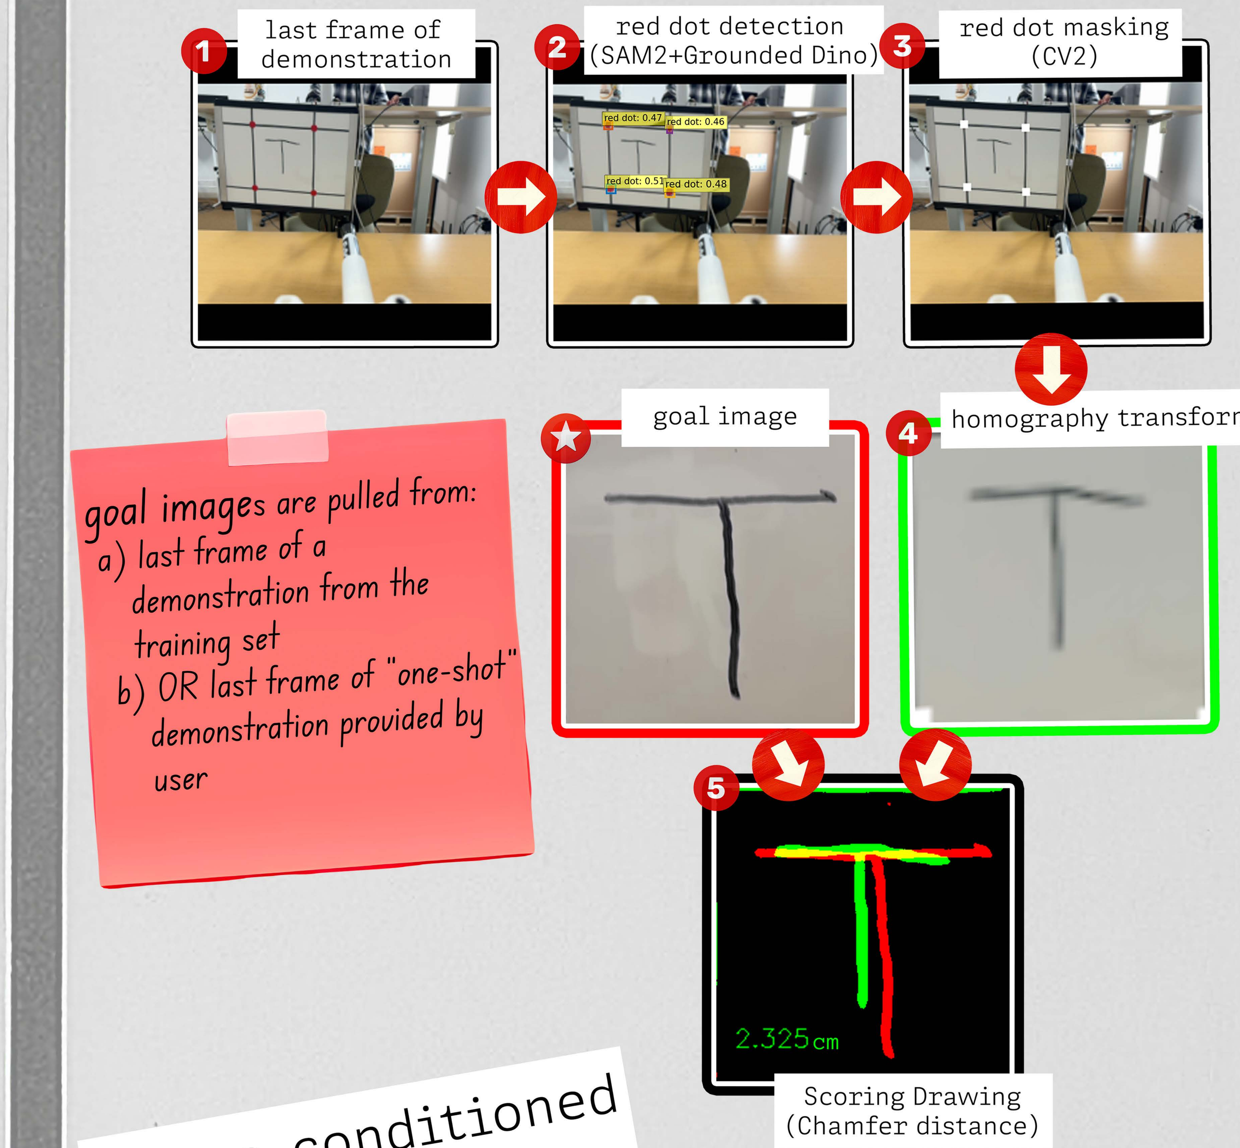

Image-conditioned policy rollout

Future Work

- Train and evaluate video-conditioned policy

Results

- Infrastructure in place for training and evaluation
- Evaluation revealed image-condition is often ignored

Thank You!

Shuran Song et al. "Diffusion Policy.", Robotics: Science and Systems, 2023 [1]

SEED Version

Teaching Robots to Write, Toward Improved Diffusion Policy

Joel Castro,[1] Austin Patel,[2] Shuran Song,[2]

[1]UC Berkeley; [2]Stanford University

Motivation

Task descriptors are used to prompt machine learning models to perform a desired task. For example, one uses a *text* task descriptor to ask an LLM to write things. In robotics, task descriptors include text (e.g. "put the cup on the plate"), but can come in many forms. For example, one can use a *goal image* to describe to an ML how you want it to manipulate the world around it (e.g. a picture of a cup on a plate). However, what if I wanted a robot to perform a task that required knowledge of intermediate steps (not just a descriptor of the end goal)? A state of the art Diffusion Policy[1], fails to perform such tasks given text or image task descriptors.

Approach

Our group, REAL Labs, is working to solutions to these limitations. Specifically, a Diffusion Policy architecture that is capable of taking an entire *video demonstration* as a task descriptor. To test this novel architecture, we needed a robot task that indeed requires knowledge of intermediate steps. We came up with a robot writing task. Components were CADed and 3D printed. An IPhone mounted to a hand-held gripper was used for task demonstrations data collection. Next, 3 independent baseline single-task Diffusion Policies were trained on demonstrations of drawing "T", "L", & "O" to verify task/hardware robustness and feasibility. Then, 1 goal-image-conditioned policy was trained on demonstrations of all 3 letters. An evaluation metric was developed to score trained policies.

Results

Designed Novel Robot Writing Task & Hardware

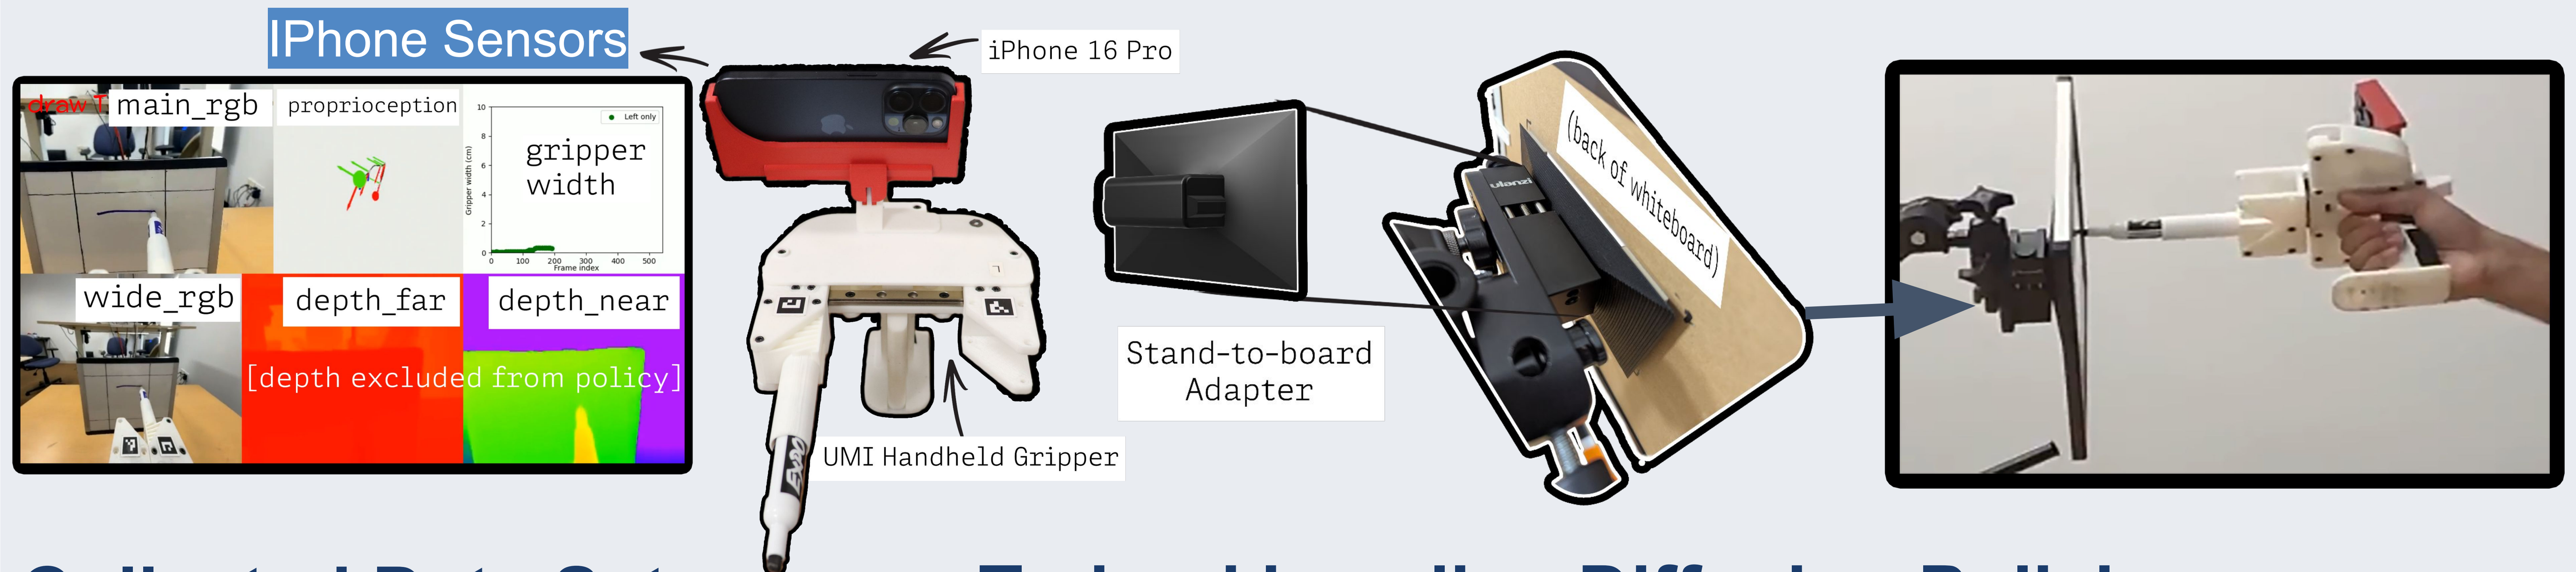

Collected Data Set

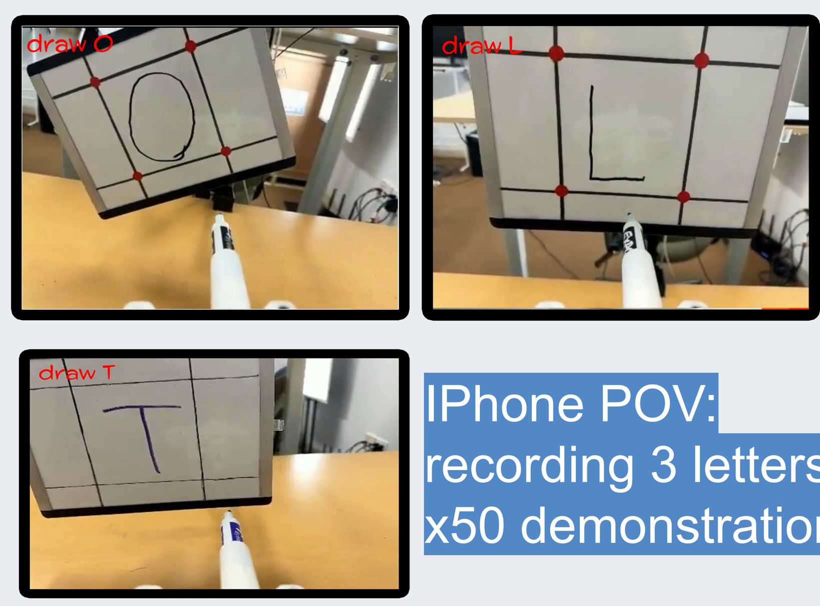

IPhone POV: recording 3 letters x50 demonstrations

Trained baseline Diffusion Policies

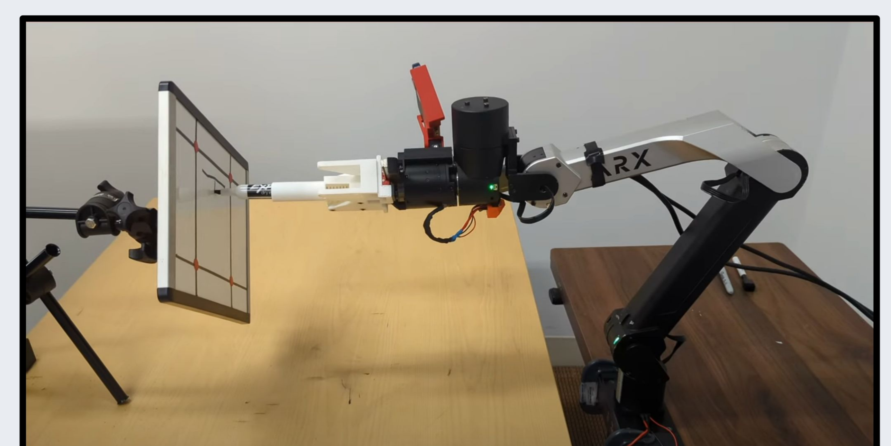

E.g. rollout of baseline policy trained on the "draw T" task alone

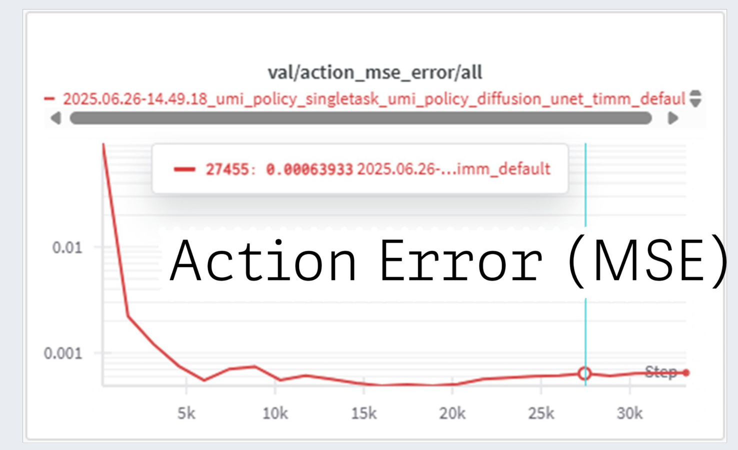

Validation action MSE Error over training steps.

Developed Task Evaluation Metric

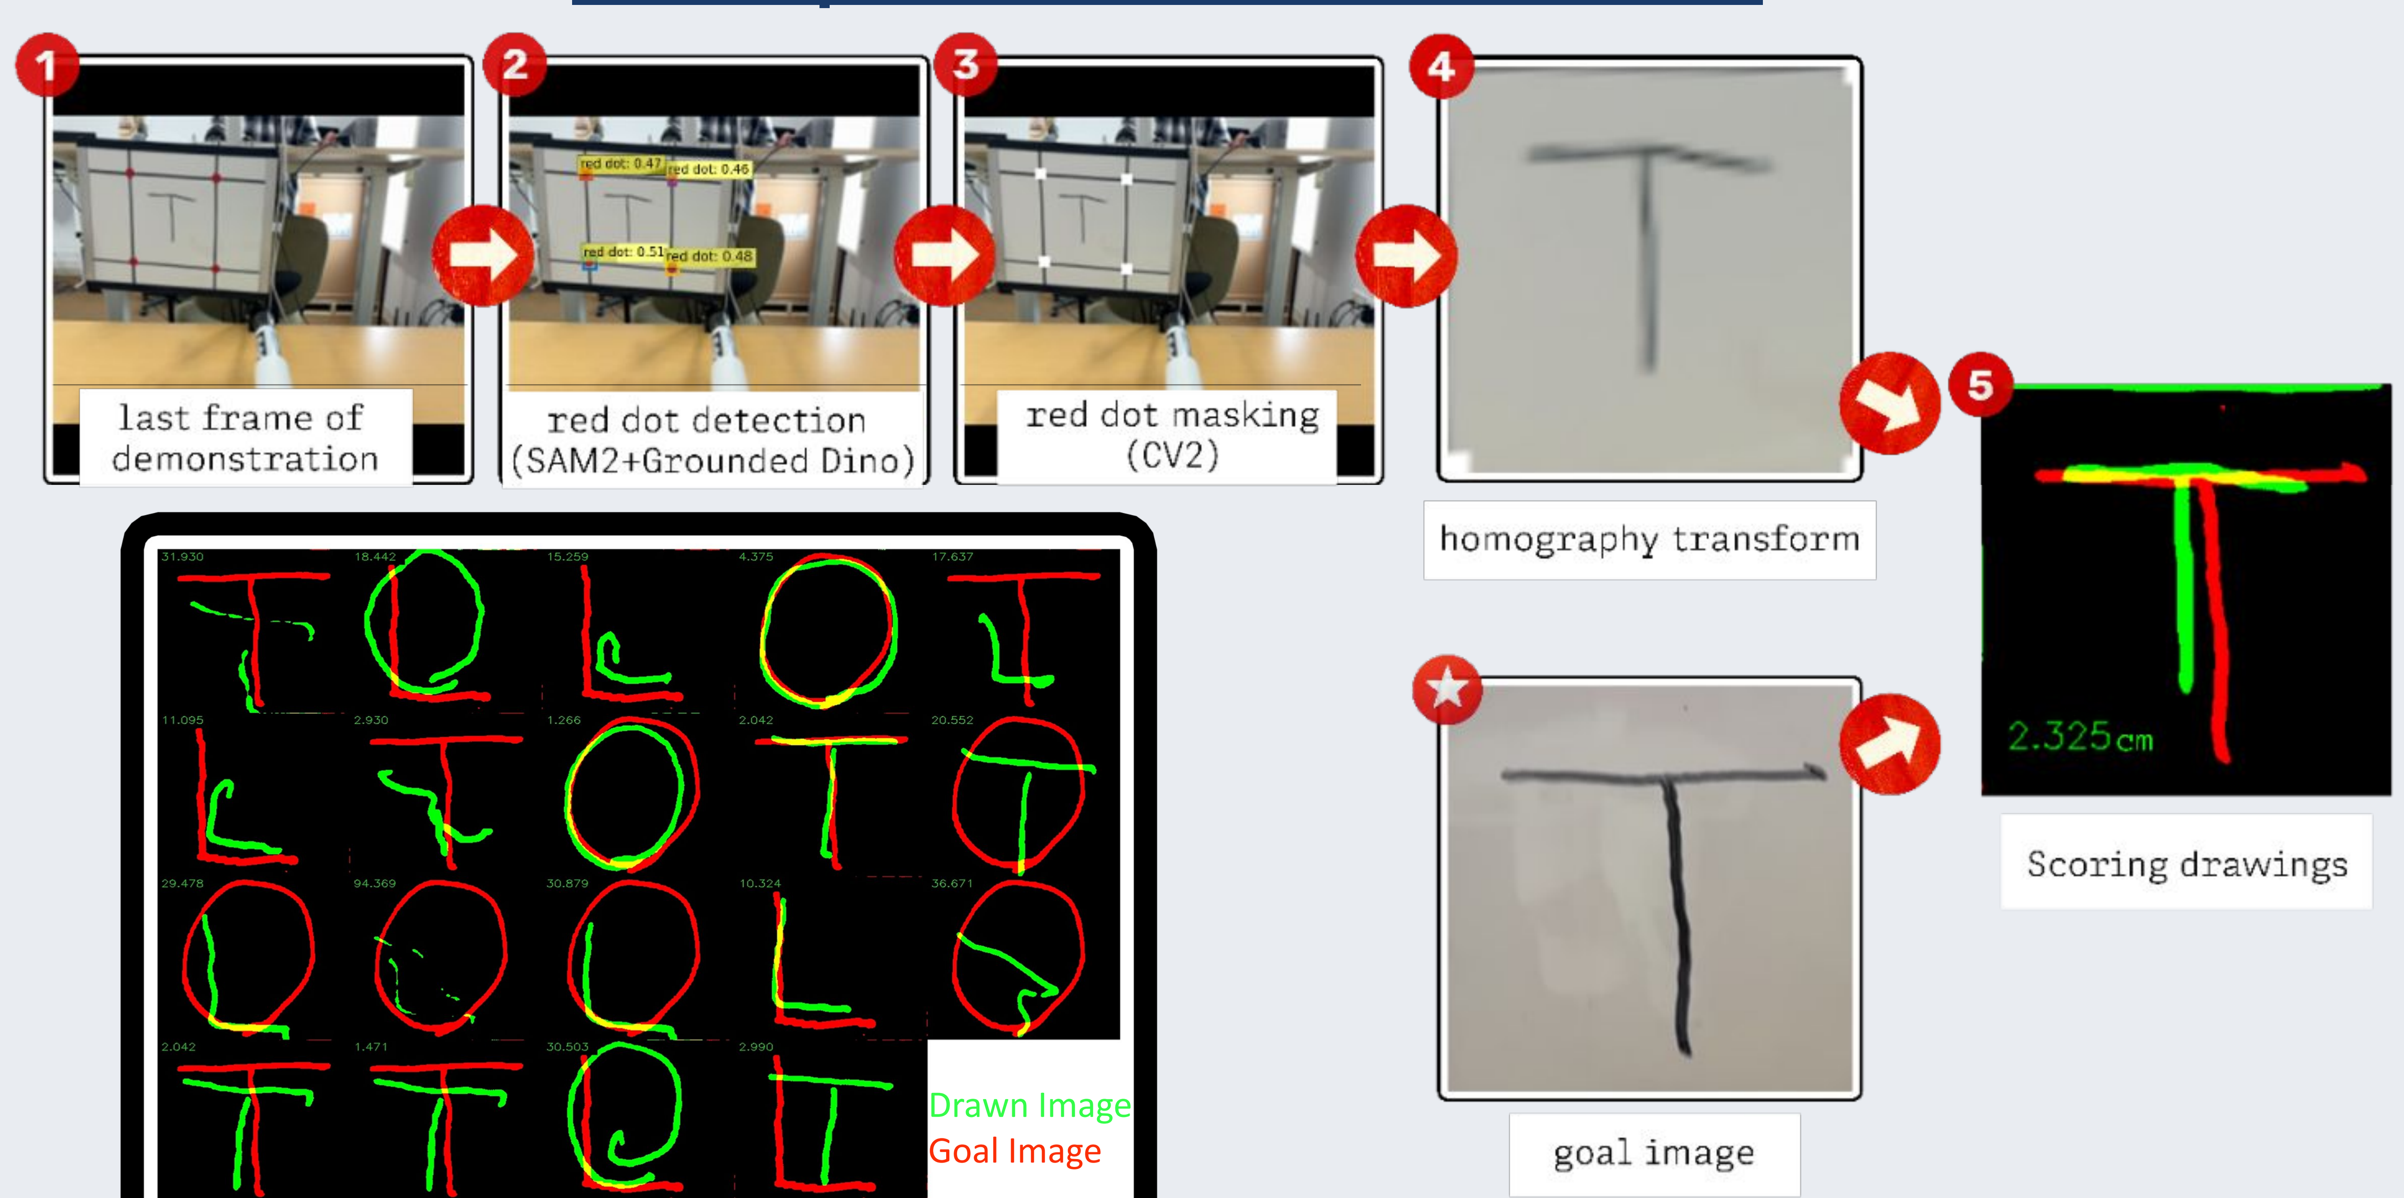

Evaluation Metric applied to rollout of image-conditioned Diffusion Policy

Conclusion

Running the goal-image-conditioned Diffusion Policy trained on demonstrations "T", "L", and "O" revealed that the policy largely ignored the goal-image and drew letters from its dataset at seemingly random. This was the expected result for goal-image-conditioning on a task like drawing letters because of the need for knowledge of intermediate steps.

Next Steps

With the resulting task design, manufactured hardware, compiled dataset, and a complete task evaluation metric, the infrastructure is in place to compare our group's novel video-conditioning architecture with existing approaches (such as the goal-image-conditioned Diffusion Policies we have evaluated here). Once our video-conditioning policy is complete, we will run it through a parallel analysis and compare, anticipating significant performance improvements. This will allow us to clearly showcase the benefits of our novel policy architecture, establishing its dominance for complex robot tasks, and opening the door to more generalizable vision-to-action models that can perform even unseen tasks from single demonstrations.

Acknowledgements

Research funded by the SEED Scholars Honors and SURF REU program.

Shuran Song et al. "Diffusion Policy.", Robotics: Science and Systems, 2023 [1].