

RESEARCH PORTFOLIO • 2026

DongGeon (Lloyd) Kwak

B.E. Computer Science and Engineering, SeoulTech

Robot learning and policy adaptation

kwaklloyd@gmail.com • lloydkwak.github.io • github.com/lloydkwak

RESEARCH INTERESTS

- Generalist robot learning
- Cross-embodiment learning
- Lifelong robot learning: adaptation, failure detection and recovery
- Long-horizon manipulation and task and motion planning

CONTENTS

FAVOR joint-failure correction 2–3

MTREA fine-tuning on a new arm 4

MakeYourBrick photo to LEGO model 5

RL-Arena, image detection 6

Logisim MLP, work experience 7

FAVOR: correcting a diffusion policy after a joint failure

QUESTION

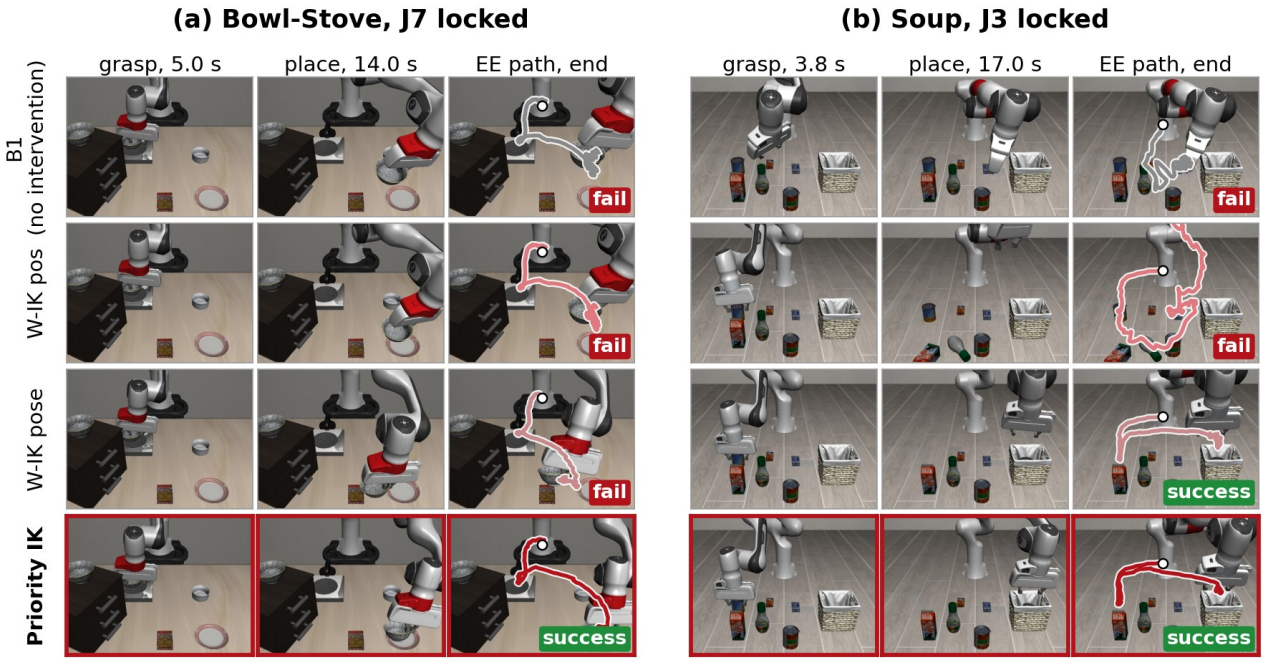
When one joint of a robot arm locks or loses range, can a pretrained diffusion policy still finish its task without retraining?

WHAT I DID

- Injected locked and range-limited faults into a Franka Panda running a diffusion policy (4 LIBERO tasks, 7 joints, 4 severities)
- Compared 7 corrections at 3 stages of inference on paired seeds
- Found that restoring end-effector position first (priority IK) works best from moderate faults on, and explained why

RESEARCH CONTEXT

Research feedback from Prof. Ji-Hyeong Han; weekly lab meetings. First-author manuscript in preparation.



The same episode under four methods for two faults. Only the bottom row (Priority IK) succeeds in both.

Correcting position first worked best from moderate faults on

0.37 → 0.60
mean success, 112 conditions

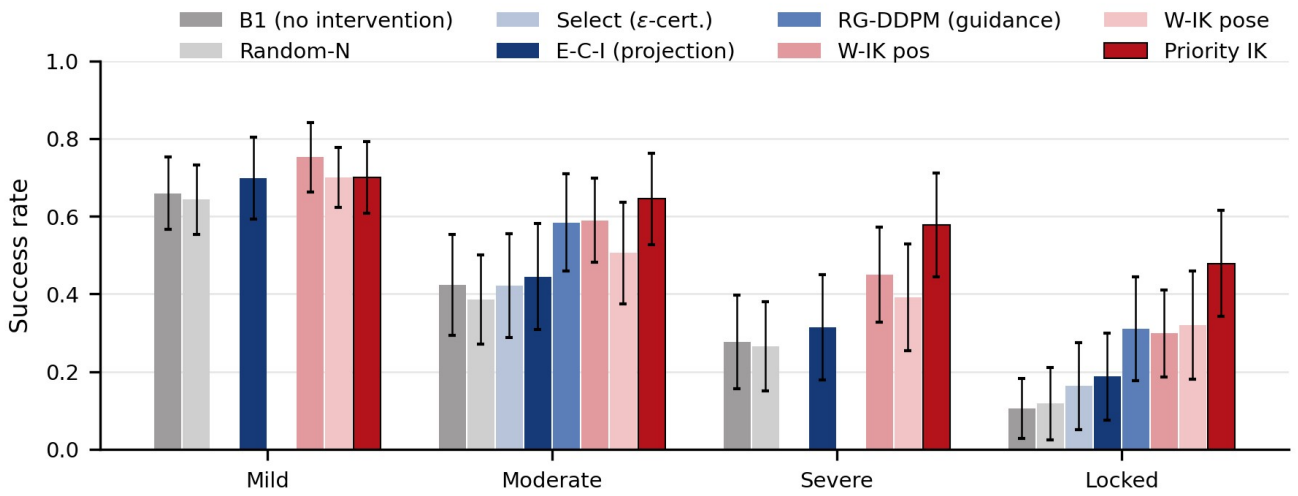
0.11 → 0.48
with a locked joint

0.94
AUC: residual error predicts recovery

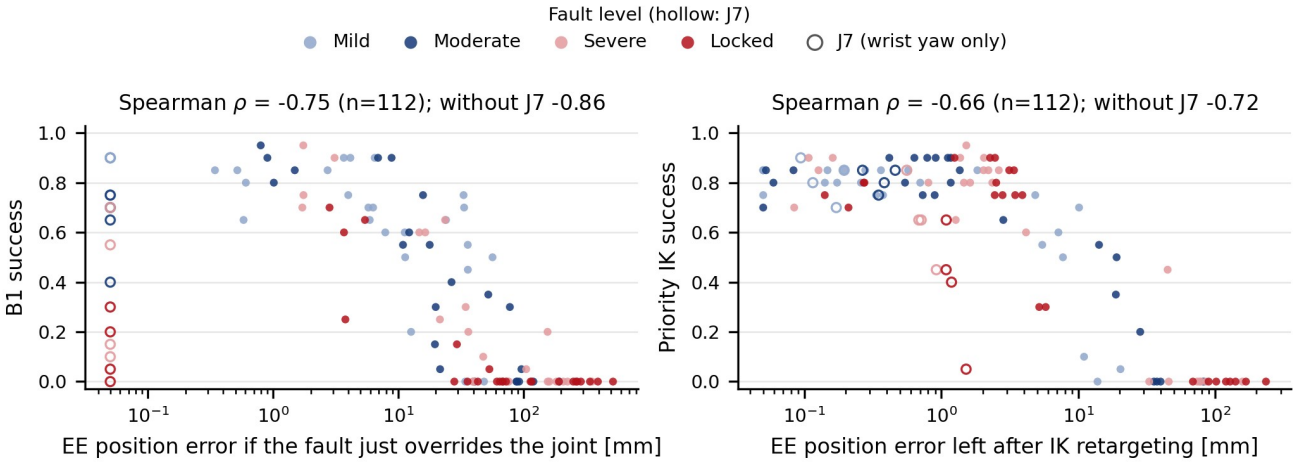
0.69
joint locks mid-task

LIMITATIONS

- Needs an accurate lock angle (0.01 rad error: 0.67 → 0.39)
- Locked joints 2 and 4 cannot be compensated
- Simulation only so far; real xArm7 tests in progress



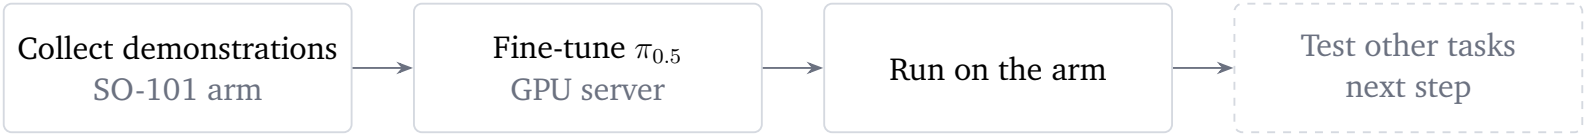
Success rate by fault level for all methods. Priority IK is dark red.



Success drops sharply once the end-effector error left after correction exceeds about 10 mm.

2 MTREA • SEP. 2026 – PRESENT • CAPSTONE DESIGN • TWO-PERSON TEAM

MTREA: does fine-tuning on a new arm reduce other skills?



QUESTION

When $\pi_{0.5}$ is fine-tuned on one task with a new low-cost arm, does it lose performance on tasks it could do before?

PROTOCOL (WITH MY TEAMMATE)

Fine-tune on pick-and-place only, then test on other tasks. Naive full fine-tuning is the baseline, since zero-shot stays near zero on a new body; a model fine-tuned on each held-out task is the ceiling.

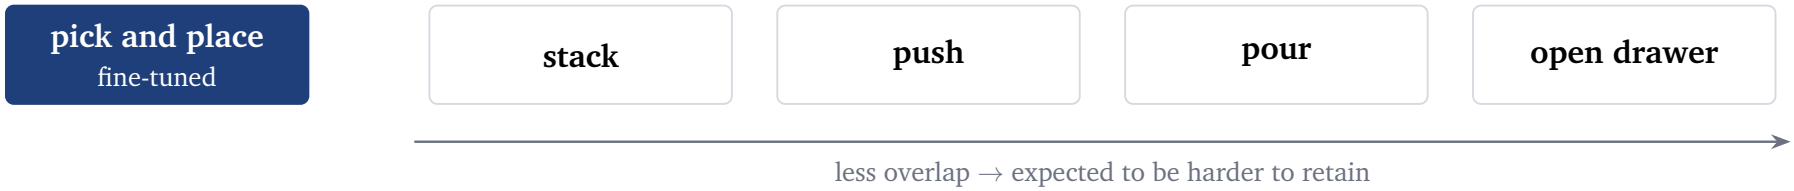
MY PART

Calibration and teleoperation with a safety lock, training on a remote GPU server, and asynchronous serving of the policy on the arm. Chose full fine-tuning at batch 16 from memory measurements on a 40 GB GPU slice.

STATUS

The pipeline works end to end. Evaluation on the other tasks is next.

HELD-OUT TASKS, ORDERED BY EXPECTED MOTION OVERLAP WITH THE FINE-TUNING TASK



MakeYourBrick: from a photo to a LEGO model

A web app that turns a single photo into a LEGO model.

HOW IT WORKS

- The user selects an object
- SAM2 segments it
- SAM 3D Objects reconstructs it in 3D
- The shape is converted into colored bricks (LDraw format)

PROBLEM I FIXED

Shadows in the reconstructed texture turned into black bricks; softening shadows before color quantization removed them.

RESULT

The model shown uses 10,263 bricks.



Input photo and the generated model.

4 RL-ARENA • 2025 • OPEN SOURCE SOFTWARE COURSE

RL-Arena: a platform for ranking RL agents

Users train reinforcement-learning agents, submit them, and see them ranked against each other.

WHO BUILT WHAT

Environment package
PyPI, me

Match runner
isolated containers, me

Backend
ELO matchmaking, mostly me

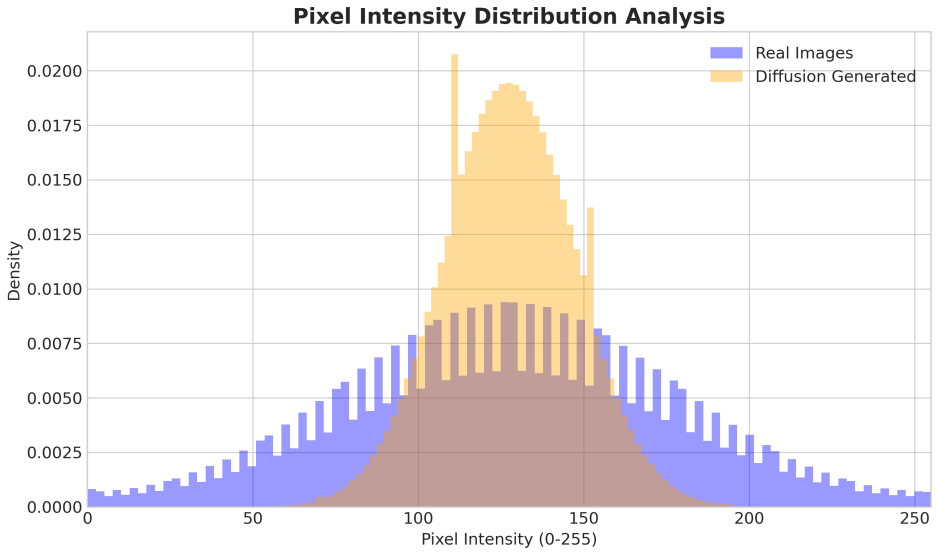
Web interface
shared

NOTE

Two-person team. Selected as an outstanding project in the course.

5 AI-GENERATED IMAGE DETECTION • 2025 • TEAM OF TWO

Detecting images made by diffusion models

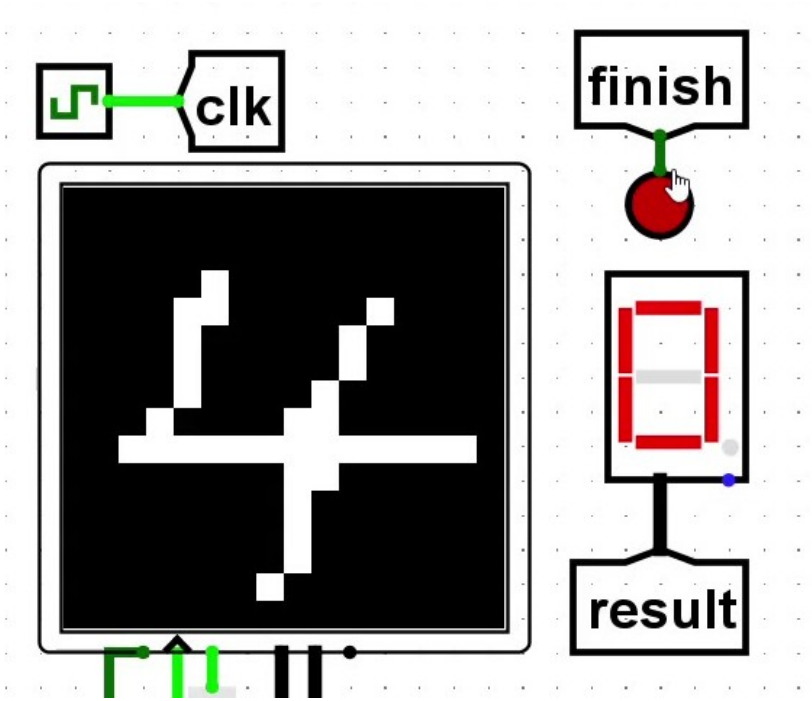


Generated images have a narrower spread of pixel intensities than real photos.

Preprocessing based on this difference raised test accuracy from 72.6% to 96.3% (ResNet-50) and from 60.7% to 88.9% (EfficientNet-B4) on 115k images.

6 LOGISIM MLP • 2025 • DIGITAL LOGIC AND SYSTEMS COURSE

A digit classifier built from logic gates



The circuit after classifying a hand-drawn “4”.

A 256-128-10 MLP with about 34,000 trained weights in ROM, built to understand how a neural network computes at the hardware level. Selected as an outstanding project in the course.

7 WORK EXPERIENCE • 2022 – 2023

Software developer

Laver (Jun. – Oct. 2022) and CHASER (Nov. 2022 – Jan. 2023).

- Object-detection system, from model development to integration
- Backend services and APIs
- IFC-to-XML converter for building information models

PATENT

Co-inventor (third of three) on Korean Patent 10-2647135, *Real-time crack detection system for construction sites*, registered March 2024. Mobile robots record video at construction sites, and cracks are detected and tracked across frames.