Dawit Chun

Taejae University, 22-8 Changdeokgung 5-gil, Jongno-gu, Seoul 03051, Republic of Korea
chundawit@taejae.ac.kr dchun4748@gmail.com +82 10-2353-0496
da-dawit.github.io github.com/da-dawit Google Scholar

Education

Taejae University Seoul, South Korea
B.S. in Data Science/AI & Natural Sciences (Physics) Expected August 2027

- GPA 4.46/4.5. Matriculation Award: admitted at age 15 to the first cohort as its youngest student, with the highest entrance score
- Founded in 2023, with all courses taught in English. Global rotation, one semester per city: Tokyo, San Francisco, New York, Shenzhen, Beijing (2025–2027)
- *Relevant coursework:* Machine Learning and Deep Learning, Artificial Intelligence I (Pace), Vision AI for Detection and Generation, Computational Decision Science, Experimental Design and Statistics, Mathematics for Social and Natural Science (calculus and linear algebra), Intro to Digital Electronics (University of San Francisco), Electronics and Optics, Physics in the Modern Society; *in progress (Fall 2026):* Computational Mathematics for Diverse Applications (linear algebra and statistics), Embodied Intelligence: AI in Physical Systems

Research Interests

I want to build robots that know how far to trust what they have learned: they apply past experience to new situations and know when not to.

- When to trust experience: per-memory reliability and surprise, failure prediction by AUROC and calibration, and the information-gathering action a robot should take after a failure, before it asks for help
- Learning intuitive physics, analogies by structure, and value from a robot's own contact with objects, with no human demonstrations
- Contact-rich manipulation on real robots: teacher-student learning, residual RL on frozen imitation policies, and sim-to-real transfer

Publications

Peer-Reviewed Papers

1. **D. Chun**, M. Jun. "PDPM: Perceptron-Driven Pipeline Morphing for Adaptive RISC-V Processors." *IEEE ICCD 2026*, pp. 396–405. Oral and poster presentation. **26% acceptance.**
2. D. Chun, **D. Chun**, Y. Cha, L. Jellibi, S.J. Park. "Watch-Out: Real-Time Tracking and Assistance for Wandering Alzheimer's Patients Using Wearable Technology." *ACM CHI 2025 EA*. doi:10.1145/3706599.3720308

Preprint

1. Y. Cha*, **D. Chun***, S.J. Park. "iPhoneme: Brain-to-Text Communication for ALS Using ConformerXL Decoding." *arXiv:2604.16441*, 2026. *Equal contribution. doi:10.48550/arXiv.2604.16441

Research Experience

MU: Learning Physics from Play, with Surprise-Gated Memory Reliability Project page
Independent project; paper in preparation for NeurIPS 2027 July 2026 – present

- MU (My Understanding) learns a world model and a hand model (where its fingertips actually go) from reward-free play in simulation (CausalWorld's TriFinger, Meta-World's Sawyer), with no human demonstrations, and plans by imagining outcomes. Its CausalWorld grip and its eight Meta-World motor primitives are hand-designed; it learns how to execute the primitives in play. A parent agent corrects child agents
- A surprise signal sets how much each memory is relied on. It raised first-try success after a world change by 59 to 75 percentage points, from at most 3% to 59–78%, on four of five tasks (10 seeds; no significant change on pushing)
- Empty and full jars look alike but need different grip forces. Child agents broke fewer objects than those without a parent agent while it watched and more after it left (5.7 against 3.55), because my rule took every memory as certain, read ordinary slips as a changed world, and treated every jar as full. Once each memory tracked its own reliability, the gap after the parent left closed (2.9 against 2.9) and most of its benefit disappeared
- On Meta-World ML45, the parent's pointing raised test success from 0.749 to 0.781, and schemas induced by comparing the successes it pointed out raised it to 0.817 (above 0.749 in 5 of 5 runs; re-trained MAML-TRPO, learning from raw actions: 0.427). Removing the parent still costs MU 6 percentage points (development split, 9 of 10 runs), and a child agent with a third of the play gained 8 percentage points in 9 of 10 runs and outscored its parent in 8 of 10. With predictions written down before each run, MU scored 0.882 on CausalWorld pushing and 0.644 on picking in unseen worlds with the hand-designed grip (learned grips, tested only on an earlier unseen set: 0.716 and 0.267), against nine 2020 model-free baselines that learn from raw actions (the per-protocol best of nine: 0.804 and 0.626). It recovered only 11% of failed first picks

Research Intern, Robot Learning and Physical AI **June 2026 – August 2026**
ROBOTIS, Humanoid Software Team; supervised by Sungho Woo Seoul
Built the teleoperation and correction pipeline, designed Prior-ACT, and trained and evaluated the policies below in simulation and on the real robots.
Video, technical story, pull requests, HG-DAgger code, Prior-ACT code

- **A residual actor trained in simulation raised Prior-ACT's real-robot success from 30 to 41 of 50 trials** (TD3+BC through Physical Intelligence's RL Token; one training seed; frozen policy; Isaac Sim bimanual two-bottle task calibrated to the AI Worker). Offline on LIBERO-10 with a TurboVLA base: 35; uniform-weight residual 59 (median of 3 seeds; no Q term); AWR with RECAP-style advantages 52. Adding TD3+BC's Q term there hurt, so the critic's share of the real-robot gain is open
- **After 13 correction episodes, end-to-end screwing went from 0 to 12 of 30 trials.** GR00T N1.7 on the AI Worker, fine-tuned on 31 demonstrations, inserted the screw in 12 of 30 trials, grasped the screwdriver in 3, and never finished. I took over by teleoperation at each failure (HG-DAgger, on a ROS 2 system I built, with auto-labeled interventions; HIL-SERL implemented, not yet validated); insertion then 30 of 30, screwdriver grasp 26 of 30 (not yet compared with 13 ordinary demonstrations)
- Designed **Prior-ACT**, an action-chunking transformer (ACT) policy whose latent comes from a frozen DINOv3 scene prior: 30 of 50 successes against vanilla ACT's 20 of 50 on the same AI Worker trials ($p \approx 0.07$, not significant; one training seed per policy) and 16% lower best held-out action error (0.1872 against 0.2238 rad); an ablation giving ACT the DINOv3 features without the prior has not been run
- Benchmarked four of 2026's newest vision-language-action models (VLAs) and world-action models (WAMs) on the 6-DOF OMY with a three-camera, 100-demonstration plate-organizing task: only GR00T N1.7 completed it, on the first try in 4 of 10 trials, 6 of 10 with age-weighted temporal ensembling (MolmoAct2, VLA-JEPA, FastWAM: none). Published as ROBOTIS's AI MANIPULATOR #13 video and docs
- Open source: two merged code PRs moved Cyclo Intelligence to LeRobot 0.6.1 and added MolmoAct2, VLA-JEPA, and FastWAM; three merged documentation PRs include the technical story

First Author, PDPM **2025 – 2026**
Computer architecture, with Prof. Minhee Jun (Catholic University of America) Paper, code

- 89.40% configuration accuracy on 301 benchmarks with deployable features (94.28% with trace-derived features; nested 5-fold cross-validation) for a RISC-V core switching among 3-, 5-, and 7-stage pipelines by a 10-weight perceptron, within 0.26% of oracle runtime, using 1.30% of FPGA LUTs with its mode-control FSM

Equal-Contribution First Author, iPhoneme **2025 – 2026**
Brain-computer interfaces, with Prof. Sung Jun Park (Taejae University) arXiv, code

- Decoded attempted speech offline from public intracortical recordings (T15: 256 electrodes, 45 sessions, 8,071 trials) with a ConformerXL decoder and a 6-gram phoneme language model: 7.86% phoneme error rate on four validation sessions, an optimistic estimate since they were also used to tune the decoder (the test set has no public labels). I designed the decoder (192.9M parameters: a dilated-convolution and GRU prenet for neural jitter, CTC subsampling, 12 Pre-RMSNorm blocks); my co-first author built the gaze-assisted interface

Research Intern (part-time), Hardware R&D **Dec 2024 – Aug 2025**
Ubion Corporation, DeepCo AI education kits Seoul
Car code, board demo video, app code (phone 1, phone 2), DeepCo Studio

- Rebuilt the DeepCo self-driving car around the Raspberry Pi 5 (PCB, servo driver, battery pack, laser distance sensor), driven by a LeNet classifier, a 2D CNN regressor, and a YOLOv8-lite detector
- Proposed the DeepCo Board, an ESP32-S3 AI education board, and built a two-floor smart-home demo on two boards for the COEX Education Fair 2025

Teaching and Service

Guest Instructor **July 2026**
Geumcheon Science Cube, Seoul
Taught 20 elementary students to train a self-driving car's AI (LeNet-5), from data collection to driving

Teaching Assistant, Taejae University **Fall 2025, Spring 2026**

- **NS200: Physics in the Modern Society** (San Francisco): ran a weekly TA session, graded all problem sets, wrote extra thermodynamics worksheets, and helped students design mechanics experiments
- **XD221: Principles of Quantitative Thinking in the Sciences** (New York): ran weekly calculus and linear algebra sessions, taught supplementary lectures, and advised students' deep learning projects

Student Body Vice President, Taejae University **2025**

Volunteer Instructor, Arduino, Circuits, and Coding **July 2022 & 2023**
STEM Teacher Training, KOICA / Korea Productivity Center, Phnom Penh, Cambodia
Presented the Arduino, circuit, and coding sessions to 30 high-school science teachers each summer and helped them build and program Arduino robots for their classrooms; the course materials were adopted in four additional provinces

Fellowships and Training

Oxford–Stanford–Taejae Fellowship — **Mar & Aug 2025**
Kellogg College, Oxford & Taejae University, Seoul — Full funding ($4,000)
One of 8 Taejae representatives in Kellogg College's Sustainability Changemakers UK 2025 program, on AI governance and climate policy; won 1st place in the AI-Powered Sustainability Policy Debate in Seoul, donated the $2,000 prize to WWF Korea, and presented our recommendations to Korea's Deputy Prime Minister

DYROS Robotics Bootcamp — **August 2025**
Seoul National University, Profs. Hyeongki Lee and Jaeheung Park
Selected for an intensive two-session program: control algorithms, motion planning, ROS, inverse kinematics, and actuator control

KAIST AI and Data Science Masterclass — **Mar 2023 – Jan 2024**
KAIST Center for Gifted Education — Grade: A+
Selected for the DICIA-KAIST AI Master Class on computer vision and algorithmic thinking

Honors & Awards

Talent Award of Korea (Minister of Education Award), University and Young Adult category — 2025
One of 60 recipients; donated the entire 2 million KRW prize to the Geumcheon Future Scholarship Foundation.
1st Place (1 of 84), ACM CHI 2025 Student Design Competition, Yokohama — 2025
1st Place, Oxford–Stanford–Taejae AI Sustainability Policy Debate — 2025
Silver Medal, International Robot Olympiad, AI Autonomous Driving, Busan — 2025
Gold Medal, International Robot Olympiad, AI Autonomous Driving, Athens — 2024
Excellence Paper Award, KSCI Winter Conference — 2024
Mayor's Commendation, Gangnam-gu, for volunteer work — 2023

Before university (2021–2023): Gold Prize, Aramco Coding World Cup; Gold and Bronze Prizes, KSME Student Design Competition; 1st Place, Geumcheon Hackathon; KAIST President's Award. Five IRO medals in total, 2023–2025.

Other Publications

Peer-Reviewed Papers

1. **D. Chun**, D. Chun, W. Hwang. "AI-Driven Smart IoT Security Window Optimization System for Disaster Response." *J. Korea Society of Computer and Information*, 30(6), 2025. doi:10.9708/jksci.2025.30.06.147
2. D. Chun, **D. Chun**, W. Hwang. "Smart Emergency Window System for Rapid Evacuation in Disasters." *J. Korea Society of Computer and Information*, 30(4), 2025. doi:10.9708/jksci.2025.30.04.093

Patent Application

1. **D. Chun**, D. Chun, W. Hwang. "AI-Based Integrated Smart Security Window and Automatic Control System." Korean Patent Application 10-2026-0029503, filed Feb 13, 2026.

Books (Co-author)

1. **D. Chun** et al. *Youth Conversations on Sustainable Development*. Office for Government Policy Coordination, 2025. Pub. No. 11-1092000-100007-01.
2. W. Hwang, **D. Chun**, D. Chun. *ChatGPT: Practical AI for the Classroom*. ISBN 9788970930084, 2024.

Skills

Programming: Python, C++, C, VHDL, Verilog, SystemVerilog, Arduino, ESP32, Raspberry Pi
Tools: PyTorch, TensorFlow, OpenCV, ROS 2, MuJoCo, Isaac Sim, Isaac Lab, LeRobot, W&B, SolidWorks, 3D printing, Git, LaTeX
Domains: VLAs and WAMs, imitation and reinforcement learning, human-in-the-loop learning, sim-to-real, CPU architecture, IoT

Certifications

GPA 4.46/4.5 — TOEFL iBT 118/120 — HSK Level 4
OPIc Advanced Low — Generative AI Expert Level 2

Volunteer Work

U9 Change (2021 to now): connecting Korean sponsors with the children they support in Kenya and South Sudan; Excellence Award, 2023. Hosanna Church meal delivery for elderly residents living alone (2021–2023); Gangnam-gu amphibian conservation (2023); Geumcheon science classrooms.