

SEONVIN CHO

Integrated M.S./Ph.D. Student in Electronic Engineering, Hanyang University
seonbin0319@hanyang.ac.kr | Website | GitHub | Seoul, Republic of Korea

Research Interests

Reinforcement Learning · Sequential Decision Making · Generative Models

Education

Hanyang University, Seoul, Republic of Korea Sep. 2024 – Present
Integrated M.S./Ph.D. Student in Electronic Engineering
Advisor: Prof. Songnam Hong
GPA: 4.45 / 4.50

Hanyang University, Seoul, Republic of Korea Mar. 2020 – Aug. 2024
B.S. in Electronic Engineering
Summa Cum Laude, Early Graduation
GPA: 4.26 / 4.50

Research

Recent Work

PathBridger: Subgoal Bridges for Offline Goal-Conditioned Reinforcement Learning 2026
Soohyun Choi*, **Seonvin Cho***, Songnam Hong[†] · *Preprint* · *Equal contribution · [†]Corresponding author
[Code] [arXiv]

- Developed a hierarchical offline goal-conditioned RL framework that connects subgoal selection to short-horizon execution through explicit state-space bridges.
- Combined generative endpoint proposals, transitive value-based selection, and inverse-dynamics decoding to produce executable short-horizon action chunks.

Multi-step Proximal Policy Improvement in Offline Reinforcement Learning 2026
Soohyun Choi*, **Seonvin Cho***, Songnam Hong[†] · *Preprint* · *Equal contribution · [†]Corresponding author
[Code] [arXiv]

- Developed a geometric view of offline actor updates, interpreting a broad class of behavior-anchored objectives as a single proximal policy improvement step on a statistical manifold.
- Proposed MPI, a plug-in refinement that composes sequential re-centered proximal steps and improves TD3+BC, ReBRAC, and IQL on many D4RL tasks.

Federated Learning under Noisy Labels 2023
Undergraduate Capstone Project
[Code] [Technical Report]

- Analyzed loss-value distributions in federated learning to identify unreliable samples under label noise.
- Received the Best Award (Department) and Excellence Award (College of Engineering).

Ongoing Research

MART: Multi-Step Actor Refinement Trajectories Present
Developing a fixed-budget actor-refinement path that divides a nominal improvement horizon across predecessor-centered actors and separates critic-facing exposure from deployment reach.

AMO: Adaptive Multiscale Policy Optimization Present
Studying adaptive proximal policy improvement that learns a total horizon T with a cross-fitted outer objective, using a conservative T/N -policy for Bellman targets and a full- T policy for execution.

Extensions of PathBridger Present
Exploring offline-to-online learning, longer-horizon control, and action-free settings where the data do not provide actions.

Honors & Awards

BK Research Assistantship (BK-RA), Hanyang University Spring & Fall 2026
IEEE Antennas and Propagation Society Undergraduate Summer Research Scholarship 2024
Excellence Paper Award, KIEES [Code] Winter & Fall 2024
Best & Excellence Awards, Federated Learning Capstone Project 2023
Excellence Award, Embedded Creative Robot Challenge 2022
Merit-Based & Academic Scholarships, Hanyang University 2021 – 2024