Donghu Kim

✉ quagmire@kaist.ac.kr | 🏠 i-am-proto.github.io | 🔗 I-AM-PROTO

Research Interest

Efficient Reinforcement Learning
- Making RL work with as little samples and/or compute as possible.
- Related: Simba, SimbaV2, AtariPB

Plasticity
- Maintaining plasticity (the ability to train) when the data distribution is constantly shifting/expanding.
- Related: Dynamic MoE, Catastrophic Interference, Hare&Tortoise

Education

KAIST   *Seongnam, Korea*
M.S. Candidate in AI (GPA: 3.82/4.3, Advisor: Jaegul Choo)   *Mar. 2024 - Present*

Korea University   *Seoul, Korea*
B.S. in Computer Science (Major GPA: 4.5/4.5, Cumulative GPA: 4.37/4.5)   *Mar. 2018 - Feb. 2024*

Work Experience

Krafton AI   *Seoul, Korea*
Research Intern @ Physical Intelligence Team   *Jun. 2025 - Sep. 2025*
- Developed locomotion RL policy for Unitree G1 robots to navigate company environment.

Publications & Preprints

SimbaV2: Hyperspherical Normalization for Scalable Deep Reinforcement Learning   *ICML'25*
Hojoon Lee*, Youngdo Lee*, Takuma Seno, **Donghu Kim**, Peter Stone, Jaegul Choo   ***Spotlight***
- arXiv / project page / code

SimBa: Simplicity Bias for Scaling Up Parameters in Deep Reinforcement Learning   *ICLR'25*
Hojoon Lee*, Dongyoon Hwang*, **Donghu Kim**, ... , Jaegul Choo, Peter Stone, Takuma Seno   ***Spotlight***
- arXiv / project page / code

ATARI-PB: Investigating Pre-Training Objectives for Generalization in Pixel-Based RL *ICML'24*
Donghu Kim*, Hojoon Lee*, Kyungmin Lee*, Dongyoon Hwang, Jaegul Choo   *Poster*
- arXiv / project page / code

Do's and Don'ts: Learning Desirable Skills with Instruction Videos   *NeurIPS'24*
Hyunseung Kim, Byungkun Lee, Hojoon Lee, Dongyoon Hwang, **Donghu Kim**, Jaegul Choo   *Poster*
- arXiv / project page

Slow and Steady Wins the Race: Maintaining Plasticity with Hare and Tortoise Networks *ICML'24*
Hojoon Lee, Hyeonseo Cho, Hyunseung Kim, **Donghu Kim**, Dugki Min, Jaegul Choo, Clare Lyle *Poster*
- arXiv / code

Projects

Dynamic Mixture-of-Experts   *2025*
- Explored dynamically adding new experts to Mixture-of-Experts layers to maintain plasticity in environments with severe distribution shifts (e.g., Craftax).
- report / slides

KAN RL   *2024*

- Implemented Kolmogorov-Arnold Network in sequential Atari environments and investigated its relevance to catastrophic forgetting and plasticity.
- report (Colab)

RL Basic Tutorial   *2024*

- Developed and delivered a series of three lectures on reinforcement learning for a government-funded bootcamp program in Korea.
- page / material1 (Korean) / material2 (Korean)

Character-level BERT   *2022*

- Proposed a character-level tokenizer for BERT to enhance robustness against character-level attacks common in spam emails.
- report / code

Honors & Awards

Dean's List (4.5/4.5 semester GPA)	Korea University	*2019, 2022*
Starcraft AI Competition Silver Prize ($2000)	NCsoft AI Fellowship	*2019*
Presidential Science Scholarship ($40000 in total)	Korea Student Aid Foundation	*2018*

Academic Services

Reviewer	International Conference on Learning Representations **(ICLR)**	*2025*
Reviewer	Conference on Neural Information Processing Systems **(NeurIPS)**	*2025*