

Dongyoon Hwang

PH.D. CANDIDATE IN ARTIFICIAL INTELLIGENCE

Seongnam, Republic of Korea

✉ godnpeter@kaist.ac.kr | 🏠 Homepage | 📧 godnpeter | 🎓 Google Scholar | 📅 Updated: August 24, 2026

Summary

I am interested in building general-purpose robots that can perceive, reason, and act in the unstructured physical world. My research centers on two directions: (i) robot foundation models that perceive and act in metric 3D space, and (ii) utilizing human motion as guidance for robot embodiment learning. I believe that policies anchored in explicit spatial structure, combined with motion priors from humans, offer a scalable path to broadly capable robots. I am currently a Ph.D. candidate at the DAVIAN Lab, KAIST AI, advised by Prof. Jaegul Choo. Previously, I was a research intern at Krafton AI and Naver Labs Europe.

Field of Interest

MY RESEARCH INTERESTS INCLUDE, BUT ARE NOT LIMITED TO:

- Building robot foundation models that perceive and act in metric 3D space [O1, C1, P1]
- Reinforcement learning from human priors for dexterous manipulation [O2, C3, P2]
- Visual representation learning for embodied agents [C2, C5, C7]

Education

Korea Advanced Institute of Science and Technology (KAIST)

Seongnam, Republic of Korea

INTEGRATED M.S./PH.D. IN ARTIFICIAL INTELLIGENCE (ADVISOR: PROF. JAEGUL CHOO, DAVIAN LAB)

Sep. 2021 - Present

Korea University

Seoul, Republic of Korea

B.S. IN COMPUTER SCIENCE

Mar. 2015 - Aug. 2021

Work Experience

Research Intern

Seoul, Republic of Korea

KRAFTON AI

Nov. 2024 - Nov. 2025

- Built an AI co-playable character (CPC) for PUBG Ally on top of LLM/VLM-based agent systems. [demo]
- Developed an LLM-based chess-reasoning agent post-trained with reinforcement learning [W1].

Research Intern

Grenoble, France

NAVER LABS EUROPE

Nov. 2022 - Feb. 2023

- Studied representation learning for reinforcement-learning-based recommender systems [C8].

Ongoing Projects

[O1] PointmapVLA: Aligning 3D Perception with Action for Vision-Language-Action Models

Dongyoon Hwang*, Dongjin Kim*, Byungkun Lee, Jaegul Choo

Ongoing research project, 2026

[O2] FlashDexRetarget: Fast and Stable Motion Tracking via RL for Dexterous Manipulation

Kyungmin Lee*, Sibeon Kim*, Dongyoon Hwang*, Yoonsang Oh*, Jaegul Choo, Hojoon Lee

Ongoing research project, 2026

Selected Publications

[C1] 3D HAMSTER: Hierarchical VLAs through 3D Trajectory Guidance

Dongyoon Hwang*, Byungkun Lee*, Dongjin Kim, Hyojin Jang, Hoiyeong Jin, Jueun Mun, Minho Park, Hojoon Lee, Hyunseung Kim, and Jaegul Choo

International Conference on Intelligent Robots and Systems (IROS), 2026 [paper]

[C3] SimBa: Simplicity Bias for Scaling Up Parameters in Deep Reinforcement Learning

Hojoon Lee*, Dongyoon Hwang*, Donghu Kim, Hyunseung Kim, Jun Jet Tai, Kaushik Subramanian, Peter R. Wurman, Jaegul Choo, Peter Stone, and Takuma Seno

International Conference on Learning Representations (ICLR), **spotlight**, 2025 [paper]

- [C2] **Adapting Pretrained ViTs with Convolution Injector for Visuo-Motor Control**
Dongyoon Hwang*, Byungkun Lee*, Hojoon Lee, Hyunseung Kim, and Jaegul Choo
International Conference on Machine Learning (ICML), 2024 [paper]

Publications

(* equal contribution)

ROBOTICS

- [C1] **3D HAMSTER: Hierarchical VLAs through 3D Trajectory Guidance**
Dongyoon Hwang*, Byungkun Lee*, Dongjin Kim, Hyojin Jang, Hoiyeong Jin, Jueun Mun, Minho Park, Hojoon Lee, Hyunseung Kim, and Jaegul Choo
International Conference on Intelligent Robots and Systems (IROS), 2026 [paper]
- [P1] **See like a Robot: Robot-Centric Pointmaps for Vision-Language-Action Models**
Byungkun Lee*, **Dongyoon Hwang***, Dongjin Kim, Hojoon Lee, Minho Park, and Jaegul Choo
arXiv preprint, 2026 [paper]
- [P2] **PHUMA: Physically-Grounded Humanoid Locomotion Dataset**
Kyungmin Lee*, Sibeon Kim*, Minho Park, Hyunseung Kim, **Dongyoon Hwang**, Hojoon Lee, and Jaegul Choo
arXiv preprint, 2025 [paper]
- [C2] **Adapting Pretrained ViTs with Convolution Injector for Visuo-Motor Control**
Dongyoon Hwang*, Byungkun Lee*, Hojoon Lee, Hyunseung Kim, and Jaegul Choo
International Conference on Machine Learning (ICML), 2024 [paper]

REINFORCEMENT LEARNING

- [C3] **SimBa: Simplicity Bias for Scaling Up Parameters in Deep Reinforcement Learning**
Hojoon Lee*, **Dongyoon Hwang***, Donghu Kim, Hyunseung Kim, Jun Jet Tai, Kaushik Subramanian, Peter R. Wurman, Jaegul Choo, Peter Stone, and Takuma Seno
International Conference on Learning Representations (ICLR), spotlight, 2025 [paper]
- [C4] **Do's and Don'ts: Learning Desirable Skills with Instruction Videos**
Hyunseung Kim, Byungkun Lee, Hojoon Lee, **Dongyoon Hwang**, Donghu Kim, and Jaegul Choo
Neural Information Processing Systems (NeurIPS), 2024 [paper]
- [C5] **Investigating Pre-Training Objectives for Generalization in Vision-Based Reinforcement Learning**
Donghu Kim*, Hojoon Lee*, Kyungmin Lee*, **Dongyoon Hwang**, and Jaegul Choo
International Conference on Machine Learning (ICML), 2024 [paper]
- [C6] **DISCO-DANCE: Learning to Discover Skills through Guidance**
Hyunseung Kim*, Byungkun Lee*, Hojoon Lee, **Dongyoon Hwang**, Sejik Park, Kyushik Min, and Jaegul Choo
Neural Information Processing Systems (NeurIPS), 2023 [paper]
- [C7] **On the Importance of Feature Decorrelation for Unsupervised Representation Learning in Reinforcement Learning**
Hojoon Lee, Koanho Lee, **Dongyoon Hwang**, Hyunho Lee, Byungkun Lee, and Jaegul Choo
International Conference on Machine Learning (ICML), 2023 [paper]

OTHERS

- [W1] **Can Large Language Models Develop Strategic Reasoning? Post-training Insights from Learning Chess**
Dongyoon Hwang*, Hojoon Lee*, Jaegul Choo, Dongmin Park, and Jongho Park
Conference on Language Modeling (CoLM), SCALR workshop, 2025 [paper]
- [C8] **SARDINE: Simulator for Automated Recommendation in Dynamic and Interactive Environments**
Romain Deffayet, Thibaut Thonet, **Dongyoon Hwang**, Vassilissa Lehoux, Jean-Michel Renders, and Maarten de Rijke
ACM Transactions on Recommender Systems (TORS), 2024 [paper]

- [C9] **Towards Validating Long-Term User Feedbacks in Interactive Recommender Systems**
 Hojoon Lee, **Dongyoon Hwang**, Kyushik Min, and Jaegul Choo
ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR), short paper, **best short paper honorable mention**, 2022 [\[paper\]](#)

- [C10] **DraftRec: Personalized Draft Recommendation for Winning in MOBA Games**
 Hojoon Lee*, **Dongyoon Hwang***, Hyunseung Kim, Byungkun Lee, and Jaegul Choo
ACM Web Conference (WWW), 2022 [\[paper\]](#)

- [P3] **MOI-Mixer: Improving MLP-Mixer with Multi Order Interactions in Sequential Recommendation**
 Hojoon Lee, **Dongyoon Hwang**, Sunghwan Hong, Changyeon Kim, Seungryong Kim, and Jaegul Choo
arXiv preprint, 2021 [\[paper\]](#)

Honors & Awards

- 2025 **Spotlight Presentation (top 5%), *SimBa* [C3]**, ICLR
- 2022 **Best Short Paper Honorable Mention Award [C10]**, ACM SIGIR
- 2021 **Graduation Project Award, Silver Prize (2,000 USD)**, Korea University