

Doohyuk Jang

Ph.D. Student, Graduate School of AI, KAIST

Daejeon, South Korea | +82-10-3757-5227 | jadohu@kaist.ac.kr | Homepage | GitHub | LinkedIn | Google Scholar

SUMMARY

Ph.D. student at the Machine Learning and Intelligence Lab (MLILAB), KAIST, advised by Prof. Eunho Yang. My research centers on advancing the reasoning capabilities of large language models, mainly through reinforcement learning-based post-training. More recently my interest has been moving toward physical AI, carrying that reasoning beyond text into embodied agents that perceive, act, and learn from continuous interaction with the physical world. I also work on inference efficiency for foundation models, through techniques such as speculative decoding and knowledge distillation. Three first or co-first author papers at ICML, ICLR, and EMNLP Findings.

SKILLS

Programming: Python, C/C++, Bash

ML Frameworks: PyTorch, Hugging Face Transformers, DeepSpeed, vLLM

LLM Post-Training: Reinforcement Learning with Verifiable Rewards (RLVR), PPO, GRPO, Supervised Fine-Tuning, Knowledge Distillation

Efficient Inference: Speculative Decoding, KV Cache and Memory Optimization

Tools and Infrastructure: Linux, Git, Docker, Weights & Biases

Languages: Korean (Native), English (Fluent)

PUBLICATIONS

* denotes equal contribution. First and co-first author work is listed first.

- Verifying Meta-Awareness via Predictive Rewards in Reasoning Models**
Yoonjeon Kim*, **Doohyuk Jang***, Eunho Yang
International Conference on Machine Learning (ICML), 2026. paper
- Reasoning Model is Stubborn: Diagnosing and Mitigating Reasoning Rigidity in Large Language Models**
Doohyuk Jang*, Yoonjeon Kim*, Chanjae Park, Hyun Ryu, Eunho Yang
Findings of the Association for Computational Linguistics: EMNLP, 2026. paper | project page
- LANTERN: Accelerating Visual Autoregressive Models with Relaxed Speculative Decoding**
Doohyuk Jang*, Sihwan Park*, June Yong Yang, Yeonsung Jung, Jihun Yun, Souvik Kundu, Sungyub Kim, Eunho Yang
International Conference on Learning Representations (ICLR), 2025. paper
- Discounted Beta-Bernoulli Reward Estimation for Sample-Efficient Reinforcement Learning with Verifiable Rewards**
Haechan Kim, Soohyun Ryu, Gyouk Chu, **Doohyuk Jang**, Eunho Yang
International Conference on Machine Learning (ICML), 2026. paper
- ReviewScore: Misinformed Peer Review Detection with Large Language Models**
Hyun Ryu, **Doohyuk Jang**, Hyemin S. Lee, Joonhyun Jeong, Gyeongman Kim, Donghyeon Cho, Gyouk Chu, Minyeong Hwang, Hyeongwon Jang, Changhun Kim, Haechan Kim, Jina Kim, Joowon Kim, Yoonjeon Kim, Kwanhyung Lee, Chanjae Park, Heecheol Yun, Gregor Betz, Eunho Yang
Findings of the Association for Computational Linguistics: EMNLP, 2026. paper
- PromptKD: Distilling Student-Friendly Knowledge for Generative Language Models via Prompt Tuning**
Gyeongman Kim, **Doohyuk Jang**, Eunho Yang
Findings of the Association for Computational Linguistics: EMNLP, 2024. paper
- SeamsTalk: Seamless Talking Face Generation via Flow-Guided Inpainting**
Yeongho Jeong, Gyeongman Kim, **Doohyuk Jang**, Jaeryong Hwang, Eunho Yang
IEEE Access, 2024. paper

WORKSHOP PAPERS

1. **LANTERN++: Enhancing Relaxed Speculative Decoding with Static Tree Drafting for Visual Auto-regressive Models**
Sihwan Park*, **Doohyuk Jang***, Sungyub Kim, Souvik Kundu, Eunho Yang
ICLR 2025 Workshop on Scalable Optimization for Efficient and Adaptive Foundation Models (SCOPE), 2025.
Oral presentation. [paper](#) | [code](#)

PREPRINTS

1. **Med-PerSAM: One-Shot Visual Prompt Tuning for Personalized Segment Anything Model in Medical Domain**
Hangyul Yoon, **Doohyuk Jang**, Jungeun Kim, Eunho Yang
Preprint, 2024. [paper](#)

WORK EXPERIENCE

NAVER Cloud, Bundang-gu, Gyeonggi, South Korea

WBL Residency, Alignment Learning

Oct 2025 – Feb 2026

- Researched post-training methods to improve the performance and safety of HyperCLOVA X, as part of the Korean Proprietary AI Foundation Model project.
- Focused on improving coding ability during post-training.

Synopsys Korea, Bundang-gu, Gyeonggi, South Korea

Internship

Mar 2021 – Aug 2021

- Improved code coverage for chipfinish.
- Tested IR-drop aware DECAP cell insertion.

PROJECTS

AITRICS

Conversational language model for virtual doctors

Apr 2024 – May 2024

- Developed a conversational language model for virtual doctor scenarios.

EDUCATION

Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea

Ph.D. Student in Artificial Intelligence

Sep 2024 – Present

- Lab: Machine Learning and Intelligence Laboratory
- Advisor: Eunho Yang

M.S. in Artificial Intelligence

Mar 2023 – Aug 2024

- Lab: Machine Learning and Intelligence Laboratory
- Advisor: Eunho Yang

B.S. in Electrical Engineering and Computer Science (Double Major)

Mar 2018 – Feb 2023

TEACHING

Teaching Assistant, Machine Learning for AI (AI501), KAIST

Fall 2024

Teaching Assistant, Advanced Machine Learning for AI (AI601), KAIST

Spring 2025, Fall 2025

ACADEMIC SERVICE

Conference Reviewer: NeurIPS 2025, CVPR 2026, NeurIPS 2026

Workshop Reviewer: SCOPE @ ICLR 2025

REFERENCES

Prof. Eunho Yang, Associate Professor, KAIST | eunhoy@kaist.ac.kr