

Hangeol Chang

PH.D. STUDENT, KAIST KIM JAECHUL GRADUATE SCHOOL OF AI

✉ hangeol@kaist.ac.kr | 🏠 Born 2001 | 📧 hangeol | 🌐 hangeol-chang | 🎓 Google Scholar

Profile

Ph.D. student at KAIST working on language model reasoning, retrieval, reinforcement learning, and reward-guided model behavior. My recent work focuses on developing empirical pipelines that steer and structure LLM behavior and translate advanced AI methods into practical impact.

Education

Ph.D in KAIST

2024.08 – Present

KIM JAECHUL GRADUATE SCHOOL OF AI, ADVISOR: JONGCHUL YE

large language models, reasoning, retrieval-augmented generation, reinforcement learning for language models

M.S. in KAIST

Aug. 2022 – 2024

KIM JAECHUL GRADUATE SCHOOL OF AI, ADVISOR: JONGCHUL YE

diffusion models, Image editing, editing system using LLM.

B.S. in Yonsei university

Mar. 2018 – Feb. 2021

SCHOOL OF INTEGRATED TECHNOLOGY,

Full ride scholarship, 1 year early graduation, Grade: 3.95/4.3

Research publications

* denotes equal contribution.

Universal Reasoner: A Single, Composable Plug-and-Play Reasoner for Frozen LLMs

ICML 2026

JAEMIN KIM*, HANGEOL CHANG*, HYUNMIN HWANG*, CHOONGHAN KIM, JONG CHUL YE

arXiv, May 2025

Reward-trained plug-and-play guidance module that steers frozen LLMs at inference time and composes across tasks.

Hypothesis-Conditioned Query Rewriting for Decision-Useful Retrieval

Under review

HANGEOL CHANG, CHANGSUN LEE, SEUNGJOON RHO, JUNHO YEO, JONG CHUL YE

arXiv, Mar. 2026

Training-free retrieval planning framework that uses working hypotheses to retrieve decision-useful evidence for LLM-based QA.

Dementia-R1: Reinforced Pretraining and Reasoning from Unstructured Clinical Notes for Real-World Dementia Prognosis

EMNLP 2026

CHOONGHAN KIM*, HYUNMIN HWANG*, HANGEOL CHANG*, JAEMIN KIM*, JINSE PARK, JAE-SUNG LIM, JONG CHUL YE

arXiv, Jan. 2026

RL-based longitudinal reasoning framework using verifiable intermediate clinical rewards for dementia prognosis from clinical notes.

AgentPSO: Evolving Agent Reasoning Skill via Multi-agent Particle Swarm Optimization

ICML workshop 2026

HYUNMIN HWANG*, JAEMIN KIM*, CHOONGHAN KIM, HANGEOL CHANG, JONG CHUL YE

arXiv, May. 2026

PSO-based multi-agent framework for evolving transferable reasoning skills.

Ground-A-Score: Scaling Up the Score Distillation for Multi-Attribute Editing

IPIU 2026

HANGEOL CHANG*, JINHO CHANG*, JONG CHUL YE

arXiv, Mar. 2024

Grounded divide-and-conquer pipeline for multi-attribute diffusion editing using multimodal LLMs and zero-shot grounding.

Honors

Silver Prize, 32nd Samsung Humantech paper award

Feb, 2026

UNIVERSAL REASONER: A SINGLE, COMPOSABLE PLUG-AND-PLAY REASONER FOR FROZEN LLMs

2nd prize in Signal Processing

Silver Prize, Best Paper Award at IPIU 2026

Feb, 2026

GROUND-A-SCORE: SCALING UP THE SCORE DISTILLATION FOR MULTI-ATTRIBUTE EDITING

Top 2% paper in IPIU

Additional Research & Project Experience

RAG Pipeline Development for Clinical Decision Support

ON-SITE PROJECT EXPERIENCE AT NAVER

Jul. 2025 – Oct. 2025

Developed a retrieval-augmented generation pipeline for a physician-support system through a government-funded project.

LLM-based Dementia Data Processing and Clinical Support Platform

GOVERNMENT-FUNDED RESEARCH PROJECT

Jun. 2024 – Present

Contributed to patient-note preprocessing, conversion of clinical notes into eCRF format, LLM module selection, integration of language and imaging diagnosis modules, demo development, and an automated patient retrieval module for dementia-related data processing.

Graduate Coursework and Research Experience in Solid-State Physics

YONSEI UNIVERSITY

Mar. 2021 – Jun. 2022

Conducted solid-state physics research and completed 1.5 years of graduate coursework before transitioning to AI research.

Patent Applications

A Single, Composable Plug-and-Play Reasoner for Frozen LLMs

Sep. 2025

CO-INVENTOR

Korean Patent Application

Method and Apparatus for Providing an Integrated Chatbot for Medical AI Models

Jul. 2025

CO-INVENTOR

Korean Patent Application

Technical Skills

Programming	Python, PyTorch
LLM	RAG pipelines, vLLM inference, Hugging Face Transformers
RL / Reasoning	GRPO/RL training, reward-based reasoning, LLM evaluation
Systems	Experiment pipeline development, Linux servers, Git