Last updated: May 28, 2026

Yumin Choi

yuminchoi@kaist.ac.kr ⋄ Homepage ⋄ GitHub ⋄ Google Scholar

RESEARCH INTERESTS

My research aims to build Self-Evolving AI systems that continuously improve by accumulating, organizing, and reusing knowledge from external sources and their own experience. I study LLM-driven optimization as a mechanism toward this goal, where LLMs generate high-quality candidate solutions, prioritize promising ones before expensive evaluations, and use execution outcomes to guide future search. My recent work studies experience-driven improvement for long-horizon LLM agents, including pre-task memory construction and hindsight self-distillation from validated trajectories.

EDUCATION

KAIST, Seoul, South Korea — *Feb. 2027 (expected)*
M.S. in Artificial Intelligence

Korea University, Seoul, South Korea — *Feb. 2025*
B.S. in Mechanical Engineering & Artificial Intelligence
GPA: 4.18 / 4.50, GPA of Mechanical Engineering and Artificial Intelligence: 4.28 / 4.50

CONFERENCE PUBLICATIONS & PREPRINTS

[8] EvolveTrade: Experience-Driven Policy Refinement for Self-Evolving LLM Trading Agents
Sehee Kim*, **Yumin Choi***, Minki Kang, Sung Ju Hwang
Under review, 2026

[7] It Takes Two: Complementary Self-Distillation for Contextual Integrity in LLMs
Sangwoo Park*, Woongyeong Yeo*, Seanie Lee, **Yumin Choi**, Hyomin Lee, Kangsan Kim, Jinheon Baek†, Seong Joon Oh†, Sung Ju Hwang†
arXiv Preprint, 2026

[6] HINT-SD: Targeted Hindsight Self-Distillation for Long-Horizon Agents
Woongyeng Yeo*, **Yumin Choi***, Taekyung Ki, Sung Ju Hwang
arXiv Preprint, 2026

[5] PREPING: Building Agent Memory without Tasks
Yumin Choi, Sangwoo Park, Minki Kang, Jinheon Baek†, Sung Ju Hwang†
arXiv Preprint, 2026

[4] T-MAP: Red-Teaming LLM Agents with Trajectory-aware Evolutionary Search
Hyomin Lee, Sangwoo Park, **Yumin Choi**, Sohyun An, Seanie Lee, Sung Ju Hwang
arXiv Preprint, 2026

[3] THINKSAFE: Self-Generated Safety Alignment for Reasoning Models
Seanie Lee*, Sangwoo Park*, **Yumin Choi**, Gyeongman Kim, Minki Kang, Jihun Yun, Dongmin Park, Jongho Park, Sung Ju Hwang
arXiv Preprint, 2026

[2] Multimodal Prompt Optimization: Why Not Leverage Multiple Modalities for MLLMs
Yumin Choi*, Dongki Kim*, Jinheon Baek†, Sung Ju Hwang†
International Conference on Learning Representations (**ICLR**), 2026

[1] System Prompt Optimization with Meta-Learning
Yumin Choi*, Jinheon Baek*, and Sung Ju Hwang
Conference on Neural Information Processing Systems (**NeurIPS**), 2025

*: equal contribution; †: equal advising

JOURNAL PUBLICATIONS

[4] Impact of platelet transfusion and bleeding risk stratification in patients with immune thrombocytopenia before procedures
Ka-Won Kang, **Yumin Choi**, Hyung-Jun Lim, Kunye Kwak, Yoon Seok Choi, Yong Park, Byung Soo Kim, Kwang-Sig Lee, Ki Hoon Ahn
Scientific Reports, 2025

[3] Blood Transfusion Utilization in Patients with Severe Coronavirus Disease 2019 in the Republic of Korea: A Nationwide Population-Based Study
Young Joo Oh, Jeong Yeon Kim, Jin Woong Suh, Yujin Jeong, **Yumin Choi**, Hyung-Jun Lim, Jang Wook Sohn, Ki Hoon Ahn, Young Kyung Yoon
Journal of Clinical Medicine, 2024

[2] Explainable artificial intelligence for predicting red blood cell transfusion in geriatric patients undergoing hip arthroplasty: Machine learning analysis using national health insurance data
Hyunyoung Seong, Kwang-Sig Lee, **Yumin Choi**, Donghyun Na, Jaewoo Kim, Hyeon Ju Shin, Ki Hoon Ahn
Medicine, 2024

[1] Explainable Model Using Shapley Additive Explanations Approach on Wound Infection after Wide Soft Tissue Sarcoma Resection: "Big Data" Analysis Based on Health Insurance Review and Assessment Service Hub
Ji-Hye Choi, **Yumin Choi**, Kwang-Sig Lee, Ki-Hoon Ahn, Woo Young Jang
Medicina, 2024

EXPERIENCE

KAIST MLAI Lab *Jun. 2024 - Present*
Research Assistant (Advisor: Prof. Sung Ju Hwang)
Lead research on self-evolving LLM agents, including agent memory construction, context optimization, and self-distillation, alongside projects on system prompt optimization, multimodal prompt optimization, and task-free agent memory construction.

Korea University MLV Lab *Nov. 2023 - May 2024*
Undergraduate Research Assistant (Advisor: Prof. Hyunwoo J. Kim)
Conducted research on multimodal large language models and diffusion models, spanning visual question answering, subject-driven generation, and multimodal video understanding.

Korea University Data Science & AI Society KUBIG *Jul. 2023 - Jun. 2024*
Project Team Leader (AsKU / AI Researcher Agent)
Led a project team building AsKU, an LLM agent for AI researchers with Semantic Scholar API-based paper search and summarization. Delivered a project presentation and released the implementation repository.

Korea University Anam Hospital *Aug. 2022 - Jan. 2024*
Research Assistant (Clinical ML with National Health Insurance Data)
Co-authored clinical machine learning studies with specialists across departments using national health insurance data. Led technical discussions on disease-specific research questions, candidate risk factors, preprocessing, modeling, and SHAP/XAI analysis.

AWARDS

K-ds hackathon, President's Award of the National Research Foundation of Korea (2023): Population Modeling using DL.

ACADEMIC SERVICE

Reviewer: Conference on Neural Information Processing Systems (NeurIPS), 2026

LANGUAGES

Korean: Native **English**: Fluent

TECHNICAL SKILLS

Machine Learning: LLM agents, Optimization problem with LLM as optimizer, self-play for agent memory, self-distillation for long-horizon agents
Agentic AI Tools: Codex, Claude Code, and Hermes Agent for agent-assisted research and coding workflows; practical experience with skill, memory, and harness-based agent design.