

Hyeongu Yun

(+82)10-8742-9670 | hyeongu.yun.1989@gmail.com | hyeongu.yun@lgresearch.ai

SUMMARY

An AI Scientist at LG AI Research with over 3 years of hands-on experience in developing and fine-tuning large language models (LLMs) using cutting-edge techniques.

My expertise spans the entire pipeline of generative language model development, from curating massive corpora for pre-trained LLMs to aligning with human preference for instruction fine-tuned LLMs, enabling models to optimally meet user expectations.

Since 2024, I have led a compact squad specializing in LLM post-training (including instruction fine-tuning, preference optimization, and reinforcement learning).

EDUCATION

Seoul National University

B.S. in Department of Electrical and Computer Engineering

Seoul, Republic of Korea

Feb 2015

Seoul National University

Ph.D. in Department of Electrical and Computer Engineering

Seoul, Republic of Korea

Aug 2022

- Advisor: Kyomin Jung
- Dissertation Title: Hierarchical Context Encoder for Natural Language Processing via Leveraging Contextual Information and Memory Attention

WORK EXPERIENCE

EXAONE Lab. @ LG AI Research

Dec. 2023 - Current

Lead of a compact squad focused on LLM post training

- Forging raw data collection to instruction fine-tuning data mixture
- Enhancing the diversity and quality of post-training data mixture
- Validating and developing cutting-edge algorithms in supervised fine-tuning, preference optimization, and reinforcement learning
- Contributing as a core team member: EXAONE 3.0, EXAONE 3.5, EXAONE Deep, and EXAONE 4.0

EXAONE Lab. @ LG AI Research

Jul. 2022 - Dec. 2023

Develop a bi-lingual LLM

- Curating extensively massive corpora from various sources for LLMs
- Fine-tuning LLMs toward instruction following ability
- Designing suitable benchmarks for LLMs with specific objectives including human-preference

PROJECTS

EXAONE-4.0; 1.2B, 32B

Mar. 2025 - Jul. 2025

- Huggingface Model: <https://huggingface.co/LGAI-EXAONE/EXAONE-4.0-32B>
- High performance on STEM and Coding tasks with toggleable thinking mode
- Led the overall process of LLM post-training, including data curation, model training, and evaluation.

EXAONE-Deep; 2.4B, 7.8B, 32B

Feb. 2025 - Mar. 2025

- Huggingface Model: <https://huggingface.co/LGAI-EXAONE/EXAONE-Deep-32B>
- Long CoT enhanced model family based on EXAONE-3.5-Instruct
- Led the overall process of LLM post-training, including data curation, model training, and evaluation.

EXAONE-3.5-Instruct; 2.4B, 7.8B, 32B

Aug. 2024 - Dec. 2024

- Huggingface Model: <https://huggingface.co/LGAI-EXAONE/EXAONE-3.5-32B-Instruct>
- High performance on instruction-following tasks in both English and Korean; Supporting up to 32k context length
- Led the overall process of instruction fine-tuning, including data curation, model training, and evaluation.

EXAONE-3.0-7.8B-Instruct

Dec. 2023 - Aug. 2024

- Huggingface Model: <https://huggingface.co/LGAI-EXAONE/EXAONE-3.5-32B-Instruct>
- Competitive benchmark performance against other state-of-the-art open models of similar size
- Led the overall process of instruction fine-tuning, including data curation, model training, and evaluation.

EXAONE Universe

Jan. 2023 - Nov. 2023

- a RAG-like system designed for specific expert domain
- Curated data for pre-training and instruction fine-tuning.

EXAONE

Jul. 2022 - Dec. 2022

- pre-trained bi-lingual Large Language Model
- Conducted model evaluation and handled prompt engineering.

PUBLICATIONS

- EXAONE 4.0: Unified Large Language Models Integrating Non-reasoning and Reasoning Modes arxiv preprint
LG AI Research
<https://arxiv.org/abs/2507.11407>
- EXAONE Deep: Reasoning Enhanced Language Models arxiv preprint
LG AI Research
<https://arxiv.org/abs/2503.12524>
- EXAONE 3.5: Series of Large Language Models for Real-world Use Cases arxiv preprint
LG AI Research
<https://arxiv.org/abs/2412.04862>
- EXAONE 3.0 7.8B Instruction Tuned Language Model arxiv preprint
LG AI Research
<https://arxiv.org/abs/2408.03541>
- ListT5: Listwise Reranking with Fusion-in-Decoder Improves Zero-shot Retrieval ACL 2024
Soyoung Yoon, Eunbi Choi, Jiyeon Kim, Yireun Kim, Hyeongu Yun, and Seung-won Hwang
<https://arxiv.org/abs/2402.15838>
- Investigating the effectiveness of task-agnostic prefix prompt for instruction following AAAI 2024
Seonghyeon Ye, Hyeonbin Hwang, Sohee Yang, Hyeongu Yun, Yireun Kim, and Minjoon Seo
<https://arxiv.org/abs/2302.14691>
- PR-MCS: Perturbation Robust Metric for MultiLingual Image Captioning Findings of EMNLP 2023
Yongil Kim, Yerin Hwang, Hyeongu Yun, Seunghyun Yoon, Trung Bui, and Kyomin Jung
<https://arxiv.org/abs/2303.08389>

TECHNICAL SKILLS

Languages: Python, C/C++**Preferred Frameworks:** PyTorch, DeepSpeed, TRL, Accelerator, NeMo-Aligner