

Seil Kang

Ph.D. Student in Computer Science

seil@yonsei.ac.kr · Seoul, KR · seilk.github.io

LinkedIn · Google Scholar · GitHub

A Ph.D. student in Computer Science focused on building better multimodal AI systems. My research centers on rigorously understanding how large-scale transformers internally represent and process cross-modal information, and applying those insights to improve both model performance and downstream user experience.

EXPERIENCE

Ph.D. Student · MICV Lab, Yonsei University, College of Computing Mar 2023 – Present

Undergraduate Researcher · MICV Lab, Yonsei University, College of Computing Mar 2022 – Mar 2023

EDUCATION

Yonsei University · Ph.D. Student in Computer Science Mar 2023 – Present

Yonsei University · B.S. in Chemical Engineering Mar 2020 – Mar 2023

PUBLICATIONS

Conference

Interpretable Motion-Attentive Maps: Spatio-Temporally Localizing Concepts in Video Diffusion Transformers ★ **Highlight Paper** CVPR | Feb 2026

<https://arxiv.org/abs/2603.02919>

Authors: Youngjun Jun, [Seil Kang](#), Woojung Han, Seong Jae Hwang.

ViKey: Enhancing Temporal Understanding in Videos via Visual Prompting CVPR | Feb 2026

<https://arxiv.org/abs/2603.23186>

Authors: Yeonkyung Lee, Dayun Ju, Youngmin Kim, [Seil Kang](#), Seong Jae Hwang.

Rare Text Semantics Were Always There in Your Diffusion Transformer NeurIPS | Oct 2025

<https://arxiv.org/abs/2510.03886>

Authors: [Seil Kang](#)*, Woojung Han*, Dayun Ju, Seong Jae Hwang.

Your Large Vision-Language Model Only Needs A Few Attention Heads for Visual Grounding CVPR | Mar 2025

★ **Highlight Paper (< 3%)**

<https://arxiv.org/abs/2503.06287>

Authors: [Seil Kang](#), Jinyeong Kim, Junhyeok Kim, Seong Jae Hwang.

See What You Are Told: Visual Attention Sink in Large Multimodal Models ICLR | Mar 2025

<https://arxiv.org/abs/2503.03321>

Authors: [Seil Kang](#)*, Jinyeong Kim*, Junhyeok Kim, Seong Jae Hwang.

Journal

Complementary branch fusing class and semantic knowledge for robust weakly supervised semantic segmentation Pattern Recognition | Nov 2024

<https://www.sciencedirect.com/science/article/pii/S0031320324006733>

Authors: Woojung Han, [Seil Kang](#), Kyobin Choo, Seong Jae Hwang. IF: 7.5.

Workshop

Interpreting Attention Heads for Image-to-Text Information Flow in Large Vision-Language Models ★ **Spotlight Paper (< 13%)** NeurIPS 2025 Workshop on Mechanistic Interpretability | Oct 2025

<https://arxiv.org/abs/2509.17588>

Authors: Jinyeong Kim, [Seil Kang](#), Jiwoo Park, Junhyeok Kim, Seong Jae Hwang.

FALCON: Frequency Adjoint Link with CONTinuous Density Mask for Fast Single Image Dehazing CVPR 2025 Workshop on Image Restoration and Enhancement | Mar 2025

<https://arxiv.org/abs/2407.00972>

Authors: Donghyun Kim, [Seil Kang](#), Seong Jae Hwang.

Pre-print

Real-Time Visual Attribution Streaming in Thinking Model

Pre-print | Jan 2026

<https://arxiv.org/abs/2604.16587>

Authors: [Seil Kang](#), Woojung Han, Junhyeok Kim, Jinyeong Kim, Youngeun Kim, Seong Jae Hwang.

Physics in 2-Steps: Locking Motion Priors Before Visual Refinement Erases Them

Pre-print | Jan 2026

Authors: Woojung Han, [Seil Kang](#), Youngjun Jun, Min-Hung Chen, Fu-En Yang, Seong Jae Hwang.

Neuron-Level Approach for Multi-Hop Reasoning in Large Vision-Language Models

Pre-print | Nov 2025

Authors: [Seil Kang](#), Jinyeong Kim, Seong Jae Hwang.

WoLF: Wide-scope Large Language Model Framework for CXR Understanding

Pre-print | Mar 2024

<https://arxiv.org/abs/2403.15456>

Authors: [Seil Kang](#), Junhyeok Kim, Donghyun Kim, Seong Jae Hwang.

OPEN-SOURCE CONTRIBUTIONS

Opensmi

Feb 2026

<https://github.com/seilk/opensmi>

Open-sourced opensmi, a zero-dependency GPU allocation enforcement tool spanning multiple SSH nodes with an interactive TUI and policy-based violation detection.

- Zero-dependency architecture
- Multi-SSH-node policy enforcement
- Interactive TUI with violation detection

Opencode

Feb 2026

<https://github.com/anomalyco/opencode/pull/8248>

Contributed a merged PR to Opencode main branch.

Openclaw

Feb 2026

<https://github.com/openclaw/openclaw/releases/tag/v2026.2.22>

Contributing to Openclaw since v2026.2.22, with ongoing contributions.

SCHOLARLY ACTIVITIES

Invited Speaker · Samsung Electronics & S-oil

Aug 2025 – Present

Delivered technical lectures on Large Language Models, Transformers, and advanced machine learning to engineering and data science teams.

Reviewer · Academic Conferences

Jan 2025 – Present

Conference reviewer: (2026) ICLR, CVPR, ICML, ECCV; (2025) ICCV, ICML, AAAI, CVPR.

AWARDS

South Korea AI Graduate School Symposium - 2nd Prize · Ministry of Science and ICT (South Korea)

Aug 2025

SCHOLARSHIPS

Seoul Future Foundation Scholarship · 10M KRW (\$7,000 USD)

Nov 2025

SKILLS

Programming Languages: Python, C++, C, JavaScript, TypeScript

Frameworks: PyTorch, TensorFlow, MCP, Node.js

LANGUAGES

Korean (Native), English (Professional)