

Seil Kang

Mail: seil@yonsei.ac.kr · [Web: <https://seilk.github.io>] · [LinkedIn: Seil Kang] · [Google Scholar: Seil Kang]

Professional Summary

A Ph.D. Student in Computer Science with a domain of study in investigating the behavior and inner workings of large-scale multimodal transformers. My research focuses on interpretability-driven model improvement and alignment of large-scale multimodal transformers such as Large Vision-Language Models (LVLMs) and Diffusion Transformers (DiT).

Education

Yonsei University Ph.D. Student in Computer Science	Seoul, South Korea Mar. 2023 – Present
Yonsei University B.S. in Chemical Engineering	Seoul, South Korea Mar. 2020 – Mar. 2023

Publications (*: Equal Contribution)

Conference Publications

- Interpretable Motion-Attentive Maps: Spatio-Temporally Localizing Concepts in Video Diffusion Transformers** [paper]
Youngjun Jun, Seil Kang, Woojung Han, Seong Jae Hwang
Conference on Computer Vision and Pattern Recognition (CVPR), 2026
- ViKey: Enhancing Temporal Understanding in Videos via Visual Prompting** [paper]
Yeonkyung Lee, Dayun Ju, Youngmin Kim, Seil Kang, Seong Jae Hwang
Conference on Computer Vision and Pattern Recognition (CVPR), 2026
- Rare Text Semantics Were Always There in Your Diffusion Transformer** [paper]
Seil Kang*, Woojung Han*, Dayun Ju, Seong Jae Hwang
Neural Information Processing Systems (NeurIPS), 2025
- Your Large Vision-Language Model Only Needs A Few Attention Heads for Visual Grounding** [paper]
Seil Kang, Jinyeong Kim, Junhyeok Kim, Seong Jae Hwang
Conference on Computer Vision and Pattern Recognition (CVPR), 2025
Selected as Highlight Paper (< 3%)
- See What You Are Told: Visual Attention Sink in Large Multimodal Models** [paper]
Seil Kang*, Jinyeong Kim*, Junhyeok Kim, Seong Jae Hwang
International Conference on Learning Representations (ICLR), 2025

Journal Publications

- Complementary branch fusing class and semantic knowledge for robust weakly supervised semantic segmentation** [paper]
Woojung Han, Seil Kang, Kyobin Choo, Seong Jae Hwang
Pattern Recognition (PR), 2025, IF: 7.5

Workshop Papers and Pre-prints

- Real-Time Visual Attribution Streaming in Thinking Model** [paper]
Seil Kang, Woojung Han, Junhyeok Kim, Jinyeong Kim, Youngeun Kim, Seong Jae Hwang
Pre-print
- Interpreting Attention Heads for Image-to-Text Information Flow in Large Vision-Language Models** [paper]
Jinyeong Kim, Seil Kang, Jiwoo Park, Junhyeok Kim, Seong Jae Hwang
NeurIPS 2025 Workshop on Mechanistic Interpretability
Selected as Spotlight Paper (< 13%)

3. **Neuron-Level Approach for Multi-Hop Reasoning in Large Vision-Language Models** [paper]
Seil Kang, Jinyeong Kim, Seong Jae Hwang
Pre-print
4. **FALCON: Frequency Adjoint Link with CONTinuous Density Mask for Fast Single Image Dehazing** [paper]
Donghyun Kim, Seil Kang, Seong Jae Hwang
CVPR 2025 Workshop on Image Restoration and Enhancement
5. **WoLF: Wide-scope Large Language Model Framework for CXR Understanding** [paper]
Seil Kang, Junhyeok Kim, Donghyun Kim, Seong Jae Hwang
Pre-print

Experience

MICV Lab, Yonsei, College of Computing

Ph.D. Student Mar. 2023 – Present

- Developed novel methodologies to enhance Transformer-based multimodal models (e.g., Large Vision-Language Models and Multi-Modal Diffusion Transformer)

Undergraduate Researcher Mar. 2022 – Mar. 2023

- Co-developed FALCON, a frequency-domain approach for image dehazing, resulting in a second-author publication at CVPRW 2025.
- Contributed to a complementary-branch framework for weakly supervised segmentation, improving model robustness and achieving a second-author publication in Pattern Recognition (IF: 7.5).

Awards

South Korea 2025 AI Graduate School Symposium Aug. 2025
2nd prize, hosted by South Korea's Ministry of Science and ICT

Scholarly Activities

Invited Speaker

- **Samsung Electronics and S-oil** Aug. 2024
Delivered a series of technical lectures on Large Language Models, Transformers, and Advanced Machine Learning to engineering and data science teams.

Conference Reviewer

- (2026) ICLR, CVPR
- (2025) ICCV, ICML, AAI, CVPR

Open-source Contributions

Opensmi Jan. 2026
Open-sourced opensmi, a zero-dependency GPU allocation enforcement tool spanning multiple SSH nodes with an interactive TUI and policy-based violation detection: <https://github.com/seilk/opensmi>

Opencode Jan. 2026
Contributing to Opencode (merged to main branch) : <https://github.com/anomalyco/opencode/pull/8248>

Skills

Programming Languages: Python, C++, C
 Frameworks: PyTorch, TensorFlow, MCP
 Spoken Languages: Korean, English