

Sangmin Woo

Applied Scientist @ AWS Agentic AI

2795 Augustine Dr, Santa Clara, CA 95054, United States

✉ sangminw@amazon.com | 🏠 sangminwoo.github.io | 📄 sangminwoo | 🎓 Sangmin Woo

Experience

AWS Agentic AI Applied Scientist	Santa Clara, CA, United States Jun. 2025 - Present
AWS Bedrock Applied Scientist Intern	Remote Sep. 2024 - Mar. 2025
AWS Bedrock Applied Scientist Intern	Santa Clara, CA, United States Jun. 2024 - Sep. 2024
NAVER LABS Research Intern	Suwon, Korea Apr. 2023 - Aug. 2023

Education

Korea Advanced Institute of Science and Technology (KAIST) Ph.D. in Electrical Engineering (Advisor: Changick Kim)	Daejeon, Korea Aug. 2021 - Aug. 2025
Gwangju Institute of Science and Technology (GIST) M.S. in Electrical Engineering and Computer Science (Advisor: Kangil Kim)	Gwangju, Korea Sep. 2019 - Aug. 2021
Kyungpook National University (KNU) B.S. in Electrical Engineering	Daegu, Korea Mar. 2013 - Aug. 2019

Project

2026

Multimodal Evaluators @ Strands-Evals

[Evals | Blog](#)

A multimodal evaluation framework for assessing image-to-text outputs, using MLLM-as-a-judge scoring to measure correctness, faithfulness, instruction following, and overall response quality with or without reference answers.

2025

Pre-built Evaluators @ AWS Bedrock AgentCore Evaluation

[AgentCore](#)

An automated evaluation service for continuously assessing AI agent quality. It measures end-to-end task success, tool-use accuracy, and other built-in or custom metrics across live interactions, enabling real-time monitoring.

Evaluation Agent @ AWS Labs & Strands-Agent-SOP

[Agent-EvalKit](#) | [Agent-SOP](#) | [Blog](#)

An AI assistant that automates end-to-end agent evaluation by using domain-specific evaluation skills to plan trace-based evaluations, generate executable metrics and code, run evaluations, and produce focused, reliable, and meaningful reports.

Publication

2026

An Empirical Study of Automating Agent Evaluation **Agent**

[Paper](#) | [Code](#) | [Blog](#)

Arxiv

Kang Zhou*, **Sangmin Woo***, Haibo Ding, Kiran Ramnath, Aosong Feng, Subramanian Chidambaram, Vivek Singh, Vinayak Arannil, Muhyun Kim, Ishan Singh, Darren Wang, Zhichao Xu, Megha Gandhi, Nirmal Prabhu, Soumya Smruti Mishra, Gouri Pandeshwar, Lin Lee Cheong (*Equal Contribution)

Beyond Correctness: Rewarding Faithful Reasoning in Retrieval-Augmented Generation **Agent**

[Paper](#)

TMLR

Zhichao Xu*, Zongyu Wu*, Yun Zhou, Aosong Feng, Kang Zhou, **Sangmin Woo**, Kiran Ramnath, Yijun Tian, Xuan Qi, Weikang Qiu, Lin Lee Cheong, Haibo Ding (*Equal Contribution)

Learning to Ideate for Machine Learning Engineering Agents **Agent**

[Paper](#)

Conference of the European Chapter of the Association for Computational Linguistics (EACL) (**Oral Presentation**)

Yunxiang Zhang, Kang Zhou, Zhichao Xu, Kiran Ramnath, Yun Zhou, **Sangmin Woo**, Haibo Ding, Lin Lee Cheong

Diffusion Language Model Inference with Monte Carlo Tree Search **Gen AI**

[Paper](#)

Findings of the Association for Computational Linguistics (EACL Findings)

Zheng Huang, Kiran Ramnath, Yueyan Chen, Aosong Feng, **Sangmin Woo**, Balasubramaniam Srinivasan, Zhichao Xu, Kang Zhou, Shuai Wang, Haibo Ding, Lin Lee Cheong

2025

What and When to Look?: Temporal Span Proposal Network for Video Relation Detection **Video**

[Paper](#) | [Code](#)

Expert Systems With Applications (ESWA)

Sangmin Woo, Junhyug Noh, Kangil Kim

A Systematic Survey of Automatic Prompt Optimization Techniques **Gen AI**

[Paper](#)

Conference on Empirical Methods in Natural Language Processing (EMNLP)

Kiran Ramnath, Kang Zhou, Sheng Guan, Soumya Smruti Mishra, Xuan Qi, Zhengyuan Shen, Shuai Wang, **Sangmin Woo**, Sullam Jeoung, Yawei Wang, Haozhu Wang, Han Ding, Yuzhe Lu, Zhichao Xu, Yun Zhou, Balasubramaniam Srinivasan, Qiaojing Yan, Yueyan Chen, Haibo Ding, Panpan Xu, Lin Lee Cheong

Don't Miss the Forest for the Trees: Attentional Vision Calibration for Large Vision Language

[Paper](#) | [Code](#) | [Project](#)

Models **Gen AI** **Multi-modal**

Findings of the Association for Computational Linguistics (ACL Findings)

Sangmin Woo*, Donguk Kim*, Jaehyuk Jang*, Yubin Choi, Changick Kim (*Equal Contribution)

Modality Mixer Exploiting Complementary Information for Multi-modal Action Recognition **Multi-modal** **Video**

[Paper](#)

Computer Vision and Image Understanding (CVIU)

Sumin Lee, **Sangmin Woo**, Yeonju Park, Muhammad Adi Nugroho, Changick Kim

Parameter Efficient Mamba Tuning via Projector-targeted Diagonal-centric Linear Transformation **Learning**

[Paper](#)

IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)

Seokil Ham, Hee-Seon Kim, **Sangmin Woo**, Changick Kim

Black-Box Visual Prompt Engineering for Mitigating Object Hallucination in Large Vision Language

[Paper](#)

Models **Gen AI** **Multi-modal**

North American Chapter of the Association for Computational Linguistics (NAACL)

Sangmin Woo, Kang Zhou, Yun Zhou, Shuai Wang, Sheng Guan, Haibo Ding, Lin Lee Cheong

Diffusion Model Patching via Mixture-of-Prompts **Gen AI**

[Paper](#) | [Code](#) | [Project](#)

AAAI Conference on Artificial Intelligence (AAAI)

Seokil Ham*, **Sangmin Woo***, Jinyoung Kim, Hyojun Go, Byeongjun Park, Changick Kim (*Equal Contribution)

2024

RITUAL: Random Image Transformations as a Universal Anti-hallucination Lever in

LVLMS [Gen AI](#) [Multi-modal](#)

[Paper](#) | [Code](#) | [Project](#)

Arxiv

Sangmin Woo*, Jaehyuk Jang*, Donguk Kim*, Yubin Choi, Changick Kim (*Equal Contribution)

Flow-Assisted Motion Learning Network for Weakly-Supervised Group Activity Recognition [Multi-modal](#) [Video](#)

[Paper](#)

European Conference on Computer Vision (ECCV)

Muhammad Adi Nugroho, **Sangmin Woo**, Sumin Lee, Jinyoung Park, Yooseung Wang, Donguk Kim, Changick Kim

Spatio-Temporal Proximity-Aware Dual-Path Model for Panoramic Activity Recognition [Video](#)

[Paper](#)

European Conference on Computer Vision (ECCV)

Sumin Lee, Yooseung Wang, **Sangmin Woo**, Changick Kim

Switch Diffusion Transformer: Synergizing Denoising Tasks with Sparse Mixture-of-Experts [Gen AI](#)

[Paper](#) | [Code](#) | [Project](#)

European Conference on Computer Vision (ECCV)

Byeongjun Park, Hyojun Go, Jinyoung Kim, **Sangmin Woo**, Seokil Ham, Changick Kim

HarmonyView: Harmonizing Consistency and Diversity in One-Image-to-3D [Gen AI](#)

[Paper](#) | [Code](#) | [Project](#) | [Demo](#)

IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)

Sangmin Woo*, Byeongjun Park*, Hyojun Go, Jinyoung Kim, Changick Kim (*Equal Contribution)

Denosing Task Routing for Diffusion Models [Gen AI](#)

[Paper](#) | [Code](#) | [Project](#)

International Conference on Learning Representations (ICLR)

Byeongjun Park*, **Sangmin Woo***, Hyojun Go*, Jinyoung Kim*, Changick Kim (*Equal Contribution)

Sketch-based Video Object Localization [Multi-modal](#) [Video](#)

[Paper](#) | [Code](#)

Winter Conference on Applications of Computer Vision (WACV)

Sangmin Woo, Soyeong Jeon, Jinyoung Park, Minji Son, Sumin Lee, Changick Kim

2023

AHFu-Net: Align, Hallucinate, and Fuse Network for Missing Multimodal Action Recognition [Multi-modal](#) [Video](#)

[Paper](#)

International Conference on Visual Communications and Image Processing (VCIP) ([Oral Presentation](#))

Muhammad Adi Nugroho, **Sangmin Woo**, Sumin Lee, Changick Kim

Multi-modal Social Group Activity Recognition in Panoramic Scene [Multi-modal](#) [Video](#)

[Paper](#)

International Conference on Visual Communications and Image Processing (VCIP)

Donguk Kim, Sumin Lee, **Sangmin Woo**, Jinyoung Park, Muhammad Adi Nugroho, Changick Kim

Cross-Modal Alignment and Translation for Missing Modality Action Recognition [Multi-modal](#) [Video](#)

[Paper](#)

Computer Vision and Image Understanding (CVIU)

Yeonju Park, **Sangmin Woo**, Sumin Lee, Muhammad Adi Nugroho, Changick Kim

Audio-Visual Glance Network for Efficient Video Recognition [Multi-modal](#) [Video](#)

[Paper](#)

International Conference on Computer Vision (ICCV)

Muhammad Adi Nugroho, **Sangmin Woo**, Sumin Lee, Changick Kim

Towards Good Practices for Missing Modality Robust Action Recognition [Multi-modal](#) [Video](#)

[Paper](#) | [Code](#)

AAAI Conference on Artificial Intelligence (AAAI) ([Oral Presentation](#))

Sangmin Woo, Sumin Lee, Yeonju Park, Muhammad Adi Nugroho, Changick Kim

Modality Mixer for Multi-modal Action Recognition [Multi-modal](#) [Video](#)

[Paper](#)

Winter Conference on Applications of Computer Vision (WACV)

Sumin Lee, **Sangmin Woo**, Yeonju Park, Muhammad Adi Nugroho, Changick Kim

2022 and earlier

Explore-And-Match: Bridging Proposal-Based and Proposal-Free with Transformer for Sentence

Grounding in Videos [Multi-modal](#) [Video](#)

[Paper](#) | [Code](#)

Arxiv

Sangmin Woo, Jinyoung Park, Inyong Koo, Sumin Lee, Minki Jeong, Changick Kim

Tackling the Challenges in Scene Graph Generation with Local-to-Global Interactions [Multi-modal](#) [Image](#)

[Paper](#) | [Code](#)

IEEE Transactions on Neural Networks and Learning Systems (TNNLS)

Sangmin Woo, Junhyug Noh, Kangil Kim

Temporal Flow Mask Attention for Open-Set Long-Tailed Recognition of Wild Animals in Camera-Trap Images Image Paper
 International Conference on Image Processing (ICIP)
 Jeongsoo Kim, **Sangmin Woo**, Byeongjun Park, Changick Kim

Impact of Sentence Representation Matching in Neural Machine Translation Learning Paper
 Applied Sciences
 Heeseung Jung, Kangil Kim, Jong-Hun Shin, Seung-Hoon Na, SangKeun Jung, **Sangmin Woo**

Revisiting Dropout: Escaping Pressure for Training Neural Networks with Multiple Costs Learning Paper | Code
 Electronics
Sangmin Woo, Kangil Kim, Junhyug Noh, Jong-Hun Shin, Seung-Hoon Na

Korean Domestic

Mitigating Hallucinations in Vision Language Models with High-pass filter Gen AI Multi-modal
 Institute of Electronics and Information Engineers 2024 (IEIE)
 Donguk Kim, **Sangmin Woo**, Jaehyuk Jang, Yubin Choi, Changick Kim

Contrastive Decoding with Image Processing Filters for Mitigating Object Hallucinations in Large-scale Vision-Language Models Gen AI Multi-modal
 Institute of Electronics and Information Engineers 2024 (IEIE)
 Yubin Choi, Donguk Kim, **Sangmin Woo**, Jaehyuk Jang

Light-Weighted Korean Speech Recognition System for Edge Devices Learning
 Institute of Electronics and Information Engineers 2023 (IEIE)
 Yooseung Wang, **Sangmin Woo**, Changick Kim

On Learning Relations between Objects in Images Multi-modal Image
 Korea Institute of Military Service and Technology 2022 (KIMST)
Sangmin Woo, Changick Kim

Effective Trash Classification using Attentional Learning Image Code
 Korea Software Congress 2018 (KSC)
Sangmin Woo, Soon Ki Jung

Patent

Model Adaptive Visual Prompt Routing System US Patent
Sangmin Woo, Kang Zhou, Yun Zhou, Shuai Wang, Sheng Guan, Haibo Ding, Lin Lee Cheong

Group Activity Recognition Apparatus and Method using RGB Videos and LiDAR Data KR Patent 10-2741168
 Changick Kim, Jinyoung Park, Donguk Kim, Sumin Lee, Muhammad Adi Nugroho, **Sangmin Woo**, Yooseung Wang

Method and Device for Inferring Dynamic Relationship between Objects in Video KR Patent 10-2590916
Sangmin Woo, Kangil Kim

Human Activity Recognition Apparatus, Human Activity Recognition Method and Computer Program for Classifying Activity Categories of User KR Patent: 10-2515065
 Changick Kim, Inyong Koo, Yeonju Park, Minki Jeong, Sumin Lee, **Sangmin Woo**

Scene Graph Generation Apparatus KR Patent 10-2254768
Sangmin Woo, Kangil Kim

Honors & Awards

Jun 2025	Outstanding Reviewer (711/12,593), IEEE Conference on Computer Vision and Pattern Recognition
Dec 2024	Finalist, Qualcomm Innovation Fellowship 2024 Korea
Oct 2023	Invited Paper Talk, Center for Applied Research in Artificial Intelligence (CARAI) Workshop
Dec 2022	Finalist, 29th HumanTech Paper Award @ Samsung Electronics Co., Ltd.
Dec 2021	Top Award (\$ 10,000), LG Electronics Robot Contest @ LG Electronics Co., Ltd.
Nov 2019	Excellence Award (\$ 500), Creative Space G A.&IoT Makerthon @ GIST

Academic Activity

2025 Oct Area Chair, Association for Computational Linguistics Rolling Review (ACL ARR)
2024 ~ Reviewer, IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)
2025 ~ Reviewer, International Conference on Computer Vision (ICCV)
2024 ~ Reviewer, European Conference on Computer Vision (ECCV)
2026 Jan ~ Reviewer, Association for Computational Linguistics Rolling Review (ACL ARR)
2024 ~ Reviewer, Annual Conference on Neural Information Processing Systems (NeurIPS)
2024 ~ Reviewer, International Conference on Learning Representations (ICLR)
2025 ~ Reviewer, International Conference on Machine Learning (ICML)
2023 ~ Reviewer, AAAI Conference on Artificial Intelligence (AAAI)
2023 ~ Reviewer, IEEE Transactions on Neural Networks and Learning Systems (TNNLS)
2025 ~ Reviewer, Transactions on Machine Learning Research (TMLR)
2025 ~ Reviewer, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)
2025 ~ Reviewer, IEEE Transactions on Image Processing (TIP)