

Gyeongmin Kang

Seoul, Republic of Korea
ttrex0965@gmail.com | +82-10-7330-3106

SUMMARY

I am an undergraduate researcher interested in how information should be represented so that models can faithfully use the knowledge humans have organized. My current work focuses on long-document retrieval: I analyze where visual page retrievers fail and complement them with lightweight text evidence that carries the missing context and structure. Building on this, I plan to study how visual and linguistic representations can be aligned to preserve fine visual detail and context, how knowledge can be structured as graphs and updated as new information arrives, and how these ideas extend to scene-graph-based robot behavior in real environments.

INTERESTS:

Multimodal Document Understanding, RAG, Knowledge Graphs

RESEARCH APPOINTMENT

Dec. 2025 - Present	SEOUL NATIONAL UNIVERSITY Lab. of Artificial Intelligence and Big Data for Agriculture <i>Undergraduate Research Intern (Advisor: Prof. Taehyeong Kim)</i>	Seoul, Korea
------------------------	--	--------------

EDUCATION

Mar. 2021 - Feb. 2027	SEOUL NATIONAL UNIVERSITY <i>Bachelor of Biosystems Engineering</i> <i>Interdisciplinary Major in Artificial Intelligence</i> ● Overall GPA 3.7/4.3 ● Undergraduate Thesis: <i>Field2Guide</i> (in progress) ● Relevant Coursework: Basics of Deep Learning, Introduction to Computer Vision, Fundamentals of Machine Learning and Applications in Biosystems Engineering, Introduction to Data Mining, Biosensor Application and Laboratory (English-taught), Precision Agriculture System Technology	Seoul, Korea
--------------------------	---	--------------

MANUSCRIPTS IN PREPARATION

CAPE: Context-Anchored Page Evidence for Visual Page Retrieval in Long Multimodal Documents

Gyeongmin Kang, Sanghyeok Choi, Taehyeong Kim

- *In preparation for submission to NAACL*

RESEARCH EXPERIENCE

Mar. 2026 - Present	CAPE: Context-Anchored Page Evidence for Visual Page Retrieval in Long Multimodal Documents <i>Research Project Lab. of Artificial Intelligence and Big Data for Agriculture, SNU</i> ● Analyzed ColPali's top-1 errors on held-out sets (MMDocIR, DUDE) and identified two blind spots of visual page retrievers: loss of context carried over from preceding pages, and insensitivity to how terms are bound together on a page ● Proposed CAPE, which generates Page Context Anchors and Evidence Units offline with a VLM and fuses them with a frozen visual retriever, without retraining or query-time model calls ● Outperformed or matched 10 baselines, including agentic methods, on MMLongBench-Doc and LongDocURL, improving multi-evidence PR@5 over ColPali by +9.8 and +6.7 points Keywords Visual Document Retrieval, Long Multimodal Documents, Multimodal RAG
------------------------	---

Jun. 2026 - Present	Field2Guide: Multi-Hop Reasoning Benchmark for Agricultural Identification and Regulatory Recommendation Retrieval <i>Undergraduate Thesis Seoul National University</i> ● Building a benchmark that requires linking pest identification guides from USDA and land-grant university extensions with EPA pesticide labels, two independently maintained sources with no shared index ● Designing a three-tier task taxonomy (single-document lookup, cross-label comparison, cross-agency multi-hop) with scientific-name-based entity linking and verified negative cases, since a confident but wrong answer can amount to off-label pesticide use Keywords Benchmark Construction, Multi-hop QA, Cross-Document Retrieval, Agricultural AI
------------------------	---

PROJECTS & COMPETITIONS

Jun. 2026 - Aug. 2026	SNU AI Challenge 2026: Caption-Guided Four-Frame Ordering Finalist (24 of 130 teams) <i>Graduate School of Data Science, Seoul National University</i> ● Fine-tuned 4-bit Qwen3.5-9B with LoRA to generate frame-order strings; scored all 24 valid permutations at inference with input-order test-time averaging Keywords Vision-Language Models, LoRA Fine-tuning, Structured Prediction
Jun. 2026	Hybrid Rule-LLM Verification for Security Protocol Traces <i>Hackathon-style Team Project Basics of Deep Learning, Seoul National University</i> ● Owned the LLM component of a rule-LLM hybrid verifier, designing confidence-based routing that defers only low-confidence cases to the LLM
Mar. 2026 - Jun. 2026	Content-Aware Tiling for Infographic Layout Detection <i>Course Project Introduction to Computer Vision, Seoul National University</i> ● Proposed whitespace-based tiling for DocLayout-YOLO on InfoVQA infographics, compared with whole-image and fixed tiling. Keywords Document Layout Analysis, Object Detection, Infographic VQA

AWARD & HONORSHIPS

DEAN'S LIST | Fall 2025

College of Agriculture and Life Science, Seoul National University

Selected for top academic achievement in the semester

FULL TUITION SCHOLARSHIP RECIPIENT | 2022, Spring 2025, Spring 2026

College of Agriculture and Life Science, Seoul National University

Awarded full-tuition scholarships for academic performance across multiple academic years

SKILLS & Language

Programming & Frameworks: Python, PyTorch, vLLM, Hugging Face Transformers

English: OPIc IH, New TEPS 437/600

EXTRACURRICULAR ACTIVITIES

Feb. 2025	GLOBAL SNUSR CORPS SNU Social Responsibility (SNUSR), Seoul National University <i>Vice Team Leader</i> ● Prepared arts and physical education classes for four months, then taught children in Uganda for two weeks
Mar. 2021 - Present	GWANAKSA SOCCER CLUB Seoul National University Residence Hall <i>Captain (2025), Vice-Captain (2022)</i> ● Active member for six years with weekly training