Jin-Ho Ju

A researcher who bridges people and problems through data.

wlsgh20728@naver.com
+82 10-4919-1534 · Seongnam, South Korea

github.com/Wlsghdh · wlsghdh.github.io · edu-data.tistory.com

PROFILE

Data Science undergraduate at Suwon University, focused on **computer vision, deep learning theory, and applied ML systems**. First-author on research explaining why deep ResNets actually work (BMVC, under review) and on generative augmentation for small-sample industrial defect detection (KIISE). Hands-on builder of end-to-end ML products — from 3D Gaussian Splatting pipelines for a used-car marketplace to LambdaRank-based ETF trading systems. Driven by an obsession with the *why* behind every result, and a track record of leading teams that ship.

EDUCATION

Suwon University — B.S. in Data Science *Mar 2021 — Feb 2027 (Expected)*
GPA 4.33 / 4.5

Sungil High School *Mar 2019 — Feb 2021*

PUBLICATIONS

Why Deep ResNets Train Successfully: Self-Selection of Effective Depth Enabled by Skip Connections [1st Author]
arXiv · BMVC, under review (2026)
Jin-Ho Ju, Hong-Ryul Ahn

Proposed **Learnable Residual Scaling (LRS)** — $y = \alpha \cdot F(x) + (1-\alpha) \cdot x$ with a single learnable scalar per block — as a direct measurement tool for per-block usage. Showed that deep ResNets *self-select* a far smaller effective depth than their nominal one (a 200-layer network effectively uses only 5–6 blocks). Validated with α-based pruning that removes blocks with $\alpha < 0.03$ at **zero accuracy loss without retraining**, across CIFAR-10/100 and ImageNet on ResNet-50/101/152/200.

github.com/Wlsghdh/Learnable_Residual_Scaling

Effects of Generative-AI Augmentation for Small-Sample Industrial Defect Detection [1st Author]
KIISE 2026
Jin-Ho Ju, Dae-Yoon Lim, Jin-Woo Yang, Hong-Ryul Ahn

Tackled the structural data scarcity of industrial defect detection (−20 samples per class). Demonstrated that **semantic-level diversity matters more than quantity**: with Mask R-CNN, traditional augmentation *hurt* mAP (−1.76), while Gemini-2.0-based generative augmentation *improved* it (+1.90) on the same volume. Found an **8× sweet spot** when combining both (+5.0 mAP), consistent across 2-stage CNN, 1-stage CNN, and Transformer detector families.

github.com/Wlsghdh/VISION-Instance-Seg

An Integrated Preprocessing Pipeline for Model Performance Comparison on a Multimodal Gas Sensor Dataset [2nd Author]
KIISE 2025
Y. Maeng, Jin-Ho Ju, J. Yoon, W. Jung, Hong-Ryul Ahn

Designed a 4-stage standardized pipeline (anchor detection → outlier removal → min-max normalization → reproducible 20-rep CV split) for the MultimodalGasData benchmark. Under this fair-comparison protocol, benchmarked **8 models** and revealed a 22.2 %p accuracy gap between best (MT-Fusion E+D+C, 97.8 %) and worst (Random Forest, 75.6 %) — proving that *fusion strategy outweighs model depth*.

github.com/Ahn-Laboratory/Gas-Leakage-Detection

SELECTED PROJECTS

CarNeRF — AI-powered used-car dealer platform *Dec 2025 — Present*
Personal / Team Project · In active competition

End-to-end web & mobile platform where a 1-minute seller video triggers an automated pipeline: **FastGS-based 3D reconstruction (CVPR 2026)** → YOLOv8 exterior-damage detection → LightGBM price prediction → fake-listing classification → LLM-driven natural-language search. Built the entire 3D pipeline from scratch (frames → COLMAP → rembg + depth → Gaussian Splatting → SPLAT export → web viewer), reaching **PSNR 31.88 on vanilla 3DGS** and a **~15× speed-up after migrating to FastGS** (60K → 30K iter ≈ 2–3 min on A100). Resolved pycolmap segfaults via OPENBLAS thread tuning and shipped a React Native (Expo) beta client.

› *FastGS · COLMAP · SAM · YOLOv8 · LightGBM · FastAPI · SQLAlchemy · RAG · React Native*

github.com/Wlsghdh/CarNeRF · *lifeai.suwon.ac.kr:5199*

ETF with AI — LambdaRank-driven ETF / stock trading system *Feb 2026 — Present*
Personal / Team Project · In active competition

Integrated trading system combining a **LightGBM LambdaRank (Learning-to-Rank)** model over 85 features (technical indicators + macroeconomic signals + z-scores), an automated TradingView scraper (Playwright + SSH-tunneled MySQL), and a Next.js monitoring dashboard. Connected to the KIS API for real-money execution with cron-driven daily prediction and monthly retraining. Currently scaled to 1,000 tickers with a 3×3 multi-AI fusion grid (Technical / Fundamental / Market).

› *LightGBM LambdaRank · FastAPI · Next.js · Playwright · MySQL · Docker Compose · KIS API*

github.com/Wlsghdh/etf-trading-projects · *ahnbi2.suwon.ac.kr/trading*

JUMP AI 2025 — 3rd Drug Discovery Competition (solo) [Top 4% · Rank 20 / 1,134] *Jul — Aug 2025*
DACON

Built a **MAP3K5 IC50 activity predictor** from SMILES inputs for drug-candidate screening. Engineered 140-dim features (RDKit descriptors + PCA-compressed Morgan ECFP4 fingerprints), tuned 5 base learners (LightGBM / XGBoost / RF / Extra Trees / MLP) with Optuna (30 trials × 5-fold CV), then combined them via a **two-track ensemble blend** (SLSQP-weighted + quantile-matching).

› *RDKit · LightGBM · XGBoost · Optuna · scikit-learn*

github.com/Wlsghdh/Jump-AI-2025

Personal Color Diagnosis System [Hwaseong City AI Forum, Invited Talk] *Jun 2024 — Mar 2025*
Personal / Team Project

Interactive single-booth web experience that diagnoses a user's personal color from a face photo and plays matching video content. A curiosity-driven study of whether CNNs — built for pattern recognition — could perform well on a pure color-composition task. Pipeline: collected ~50K celebrity reference images, applied white-balancing, segmented skin tone via OpenCV (eyes / nose / mouth removal), then routed the model output through an OpenAI prompt to a curated video / sound booth in real time.

› *PyTorch · OpenCV · OpenAI API · AWS · Docker · React*

github.com/Woochang6862/personal-color-app

EXPERIENCE

Teaching Assistant — DSML Vibe Coding *Sep 2025 — Present*
Suwon University, DSML
Supported coding lectures for undergraduate students and mentored hands-on practice sessions.

Research Intern — Hadd Science *Sep 2025 — Dec 2025*
Sungkyunkwan University (SKKU)
Managed the Hadd Science website and research data, authored articles, and produced AutoCAD deliverables.

Undergraduate Research Assistant *Sep 2024 — Present*
Suwon University, Ahn Laboratory
Computer-vision research that led to the first- and co-authored publications listed above. Built applied React / Node.js / MongoDB prototypes alongside the research work.

Data Engineer — Autonomous Driving Data Team *Sep 2024 — Dec 2024*
AIMMO Inc.
Cleaned vision-fail data, fine-tuned models on the refined sets, and collected additional training data.

Military Service — Republic of Korea Marine Corps *Jan 2022 — Jul 2023*
Completed mandatory military service.

AWARDS & HONORS

Academic Excellence Scholarship — *Suwon University · Merit-based, awarded 5 semesters (Fall 2021, Fall 2024 – Spring 2026)* *2021 — 2026*

Encouragement Award — *Korean Institute of Information Scientists and Engineers (KIISE)* *Dec 2025*

Top 4% · Rank 20 / 1,134 — *JUMP AI 2025 Drug Discovery Competition, DACON* *Aug 2025*

1st Place — *Stock Chart Data Collection #2, Suwon University DSML* *Nov 2025*

1st Place — *Stock Chart Data Collection #1, Suwon University DSML* *Oct 2025*

Excellence Award — *AI/Develops, Suwon University DSML* *Aug 2024*

TECHNICAL SKILLS

Deep Learning	PyTorch, PyTorch Lightning, LLM fine-tuning (LoRA / PEFT), DDP, Optuna, Weights & Biases
Computer Vision	YOLO, detectron2, mmdetection, OpenCV, Gaussian Splatting, COLMAP, SAM
NLP & LLM	HuggingFace, RAG, prompt engineering
Engineering	FastAPI, SQLAlchemy, Next.js, Docker, MySQL, Selenium, Playwright
Workflow	Git (PR / branch / issue / commit), Linux, AWS

COMPLETED PROGRAMS

Google Data Analytics Professional Certificate *Feb 2025*
Coursera × Google

Digital Healthcare (Beginner) *Nov — Dec 2025*
GUIP Biohealth Platform

WHY ME

Curiosity-driven research

I refuse to stop at shallow understanding. When a question lands, I design small experiments to chase it down myself — and the failures along the way have always been the most valuable part of the process.

Root-cause problem solving

With a habit of startup-minded tinkering, I rarely walk past an inconvenience. I ask *why* repeatedly until I'm looking at the underlying mechanism — whether it's an AI model or a paper I'm reading for the first time.

Growth as a team

I have led most of my collaborative projects. What I value most is matching each member's strengths to the right role — real synergy comes from communication, not from a sum of skills.