# Chanyoung Jung

RESEARCH ASSISTANT

50 Yonsei-ro, Seodaemun-gu, Seoul, 03722, Republic of Korea

(+82) 10-5188-9516  |  march11381@gmail.com  |  lovelyzzkei.github.io  |  lovelyzzkei  |  chanyoungjung

## Summary

Highly motivated Master's student in Computer Science, **recipient of the MobiSys 2025 Best Paper Award**, specializing in on-device AI systems optimized for mobile/edge platforms. Expertise spans ML inference optimization through heterogeneous/parallel computing and efficient utilization of mobile processors for Vision Foundation Models and 3D object detection on mobile/edge devices. Seeking to leverage advanced skills in **optimizing complex AI applications for resource-constrained environments.**

## Research Areas and Interests

| | |
|---|---|
| **Mobile & Edge Systems** | Real-time Systems, **Heterogeneous Computing**, Energy-efficient Systems |
| **On-device AI** | **Neural Processing Units (NPUs)-based Optimization**, DNN Inference Optimization, HW-SW Co-design |
| **Artificial Intelligence** | Vision Transformer, Vision Foundation Models (VFMs), Efficient AI |
| **Computer Vision** | Depth Estimation, Semantic Segmentation, 3D Object Detection, Gaussian Splatting |

## Experience

### Mobile Embedded Systems Lab.
*Seoul, South Korea*

*Research Assistant*
*Jan. 2024 - Present*

- Participated in 3 projects that were accepted at top-tier mobile systems conferences, with 1 ongoing project
  - Designed and co-developed a heterogeneous computing system (ARIA) for real-time vision AI on Qualcomm Snapdragon SoCs, sustaining 30 FPS while boosting model accuracy up to +72%
  - Conducted research on a mobile streaming system (Vega) utilizing Gaussian Splatting, an emerging technique in immersive computing
  - Designed and developed a high-performance, multi-view 3D object detection pipeline (Panopticus) for autonomous systems on NVIDIA Jetson SoCs, achieving a 2.1× speedup and +62% accuracy
  - Currently exploring the design of a heterogeneous computing system for real-time multi-view vision transformer inference on mobile SoCs
- Awarded **Best Paper Award** at MobiSys 2025 by ARIA
- Authored 3 top-tier mobile systems papers, including 1 first-author paper (MobiSys '25) and 2 co-authored papers (MobiCom '24, MobiCom '25)
- Awarded a government fellowship, receiving approximately $9,000 for one year
- Served as a teaching assistant for Operating Systems and Systems Programming courses

### Mobile Embedded Systems Lab.
*Seoul, South Korea*

*Undergraduate Research Assistant*
*Dec. 2022 - Dec. 2023*

- Conducted two semesters of graduation research on "Real-time Omnidirectional 3D Object Detection on Edge Devices"
  - Developed an end-to-end baseline system for real-time omnidirectional 3D object detection during the first semester
  - Optimized critical bottleneck in image preprocessing using CUDA, achieving a 37.0x reduction in processing time (vs. naive Python baseline) and an inference rate of 17 FPS (57% of 30 FPS target)
  - Enhanced system performance by 47% to reach 25 FPS in the second semester, integrating additional CUDA-accelerated modules and adaptive logic
  - Contributed foundational work that was subsequently expanded into the Panopticus project, accepted at MobiCom 2024

## Education

### Yonsei University
*Seoul, South Korea*

*M.S. in Computer Science and Engineering*
*Mar. 2024 - Now*

- Mobile Embedded Systems Group, Advised by Prof. Hojung Cha
- Expected graduation: Feb. 2026

### Yonsei University
*Seoul, South Korea*

*B.S. in Computer Science and Engineering*
*Mar. 2018 - Feb. 2024*

- GPA: 4.24/4.50, Rank, 6/106, Graduated with High Honors

# Skills

| | | | |
|---|---|---|---|
| **Programming** | Python, C/C++, Java | **GPU Frameworks** | CUDA, OpenCL, OpenGL |
| **ML Inference Frameworks** | ONNX, LiteRT(Tensorflow Lite), TensorRT | **NPU Frameworks** | Qualcomm AI Engine Direct (QNN) SDK |
| **ML Model Frameworks** | PyTorch, LibTorch, MMDetection3D | **Edge Devices** | Nvidia Jetson Series (Orin, Xavier, Nano) |
| **Build tools** | Bazel, CMake, Gradle | **Mobile Devices** | Qualcomm Snapdragon SoCs |

# Projects

### ARIA                                                                                        *2024*
*Project Leader*
- Skills: Android, Qualcomm AI Engine Direct (QNN) SDK, LiteRT, ONNX
- Designed a **heterogeneous computing pipeline** that distributes VFM inference workloads across GPU and NPU based on processors' performance characteristics.
- Developed an inference engine for Snapdragon SoCs that schedules VFM inference based on scene dynamics, device motion, and GPU performance variations due to thermal throttling.
- Achieved sustained performance of 30 FPS with up to 72% improvement of prediction accuracy compared to baselines.

### Vega                                                                                        *2024 - 2025*
*Project Member*
- Skills: Android, OpenGL, LiteRT, PyTorch
- Developed a system for next-generation mobile immersive video streaming using Gaussian Splatting, enabling more immersive experiences than traditional volumetric video.
- Conducted comparative analyses against various baselines to validate Vega's superior performance.

### Panopticus                                                                                  *2024*
*Project Member*
- Skills: Linux, PyTorch, LibTorch, TensorRT, CUDA
- Developed a BEV 3D object detection system featuring dynamic selection of optimal inference paths per camera view, deployed on Jetson edge GPUs.
- Acheived 2.1x speed up with up to 62% improvement of prediction accuracy compared to baselines.

# Publications

### ARIA: Optimizing Vision Foundation Model Inference on Heterogeneous Mobile Processors for Augmented Reality                                                 *MobiSys '25*
*Chanyoung Jung*, Jeho Lee*, Gunjoong Kim, Jiwon Kim, Seonghoon Park, and Hojung Cha*
- ACM Annual International Conference on Mobile Systems, Applications, and Services
- Awarded **Best Paper Award**

### Vega: Fully Immersive Mobile Volumetric Video Streaming with 3D Gaussian Splatting          *MobiCom '25*
*Gunjoong Kim*, Seonghoon Park*, Jeho Lee, Chanyoung Jung, Hyungchol Jun, and Hojung Cha*
- ACM Annual International Conference on Mobile Computing and Networking

### Panopticus: Omnidirectional 3D Object Detection on Resource-constrained Edge Devices        *MobiCom '24*
*Jeho Lee, Chanyoung Jung, Jiwon Kim, and Hojung Cha*
- ACM Annual International Conference on Mobile Computing and Networking

# Honors & Awards

| | | |
|---|---|---|
| 2025 | **Best Paper Award**, MobiSys 2025 | *Anaheim, California, US* |
| 2024 | **Academic Prize**, Graduated with High Honors, Yonsei University | *Seoul, South Korea* |
| 2023 | **Grand Prize**, Software Capstone Design, Yonsei University | *Seoul, South Korea* |
| 2023 | **Academic Prize**, Honors, Yonsei University | *Seoul, South Korea* |
| 2019 | **Academic Prize**, Honors, Yonsei University | *Seoul, South Korea* |
| 2019 | **Academic Prize**, High Honors, Yonsei University | *Seoul, South Korea* |
| 2018 | **Academic Prize**, Highest Honors, Yonsei University | *Seoul, South Korea* |

| | | |
|---|---|---|
| 2025 | **Best Paper Award**, MobiSys 2025 | *Anaheim, California, US* |
| 2024 | **Academic Prize**, Graduated with High Honors, Yonsei University | *Seoul, South Korea* |