# Chanseok **Kang**

AI Software Engineer · LG Electronics AI Lab

*Seoul, South Korea*

chanseok.kang@lge.com | kcsgoodboy@gmail.com | kcsgoodboy.github.io | goodboychan | chanseokk | Google Scholar


## Summary


AI software engineer at LG Electronics since 2014. Brought reinforcement-learning control into mass-produced home appliances; my team's offline RL spin-cycle model shipped in an LG washing machine. Now focused on deploying and validating quantized models and policies on edge accelerators (Jetson, Hailo-8, Coral Edge TPU, Qualcomm QNN).


## Experience

### LG Electronics

*Seoul, South Korea*

AI Software Engineer — *Dec 2014 – Present*

- Applied energy-efficient AI algorithms to home appliances and deployed them in mass-produced products:
  - Reinforcement-learning PID tuning for HVAC control in confined spaces
  - Data-driven statistical modeling for adaptive refrigerator compressor control
  - Reinforcement learning to reduce washing-machine vibration in the spin cycle, improved continually with real-world data
  - Reinforcement-learning drying-motion strategies for dryers under varying loads
- Offline RL spin (dehydration) model from my team launched in an LG washing machine; presented it at the LG SW Developer Conference.
- Built communication interfaces between appliances and PCs to form integrated RL environments, drawing on earlier platform software and device-driver work.
- Selected among LG Electronics' first 12 certified AI Specialists.

### Carnegie Mellon University

*Pittsburgh, PA, USA*

Visiting Researcher — *Apr 2019 – Aug 2019*

- Researched reinforcement learning for energy-efficient control with Prof. Bhiksha Ramakrishnan, sponsored by LG Electronics.

### Early Roles

*Seoul, South Korea*

*Feb 2012 – Dec 2012*

- Kinect hand-gesture recognition and bi-manual interaction UX at the Korea Institute of Science and Technology (wrote a 50-page Kinect C/C++ developer guide); taught NUI programming at the University of Seoul; built a Kinect rehabilitation prototype at Kerapist; led a team to the Imagine Cup semi-final round.

## Ongoing Research

### Quantization error on edge accelerators

Cross-backend testbed: GX10, Jetson, Coral Edge TPU, Hailo-8, Qualcomm QNN — *2026*

- Screening deployment risk of quantized RL policies by the direction of backend error relative to the policy's decision geometry, not just its size.
- How backend quantization error degrades frozen context encoders in meta-RL.
- Coverage guarantees for quantized models across edge backends, and where they drift.

## Selected Projects

### EdgeGate

*GitHub · Demo video*

Meta Global AI Developer Hackathon — *Oct – Nov 2026*

- A deployment gate that measures a perception skill on the edge chip it will run on. A YOLOv8n student distilled from SAM 3.1 labels runs as TensorRT on Jetson Orin Nano and INT8 on Raspberry Pi 5 + Hailo-8.
- Measured over 2 h 50 min of frozen recordings, the gate held all 12 deliberately broken builds; Muse Glimmer 30B audits the evidence through MCP tools as a veto-only reviewer.

### QuantGate

*GitHub · Demo video*

Nebius × NVIDIA Global AI Hackathon — *2026*

- A deployment gate for quantized robot policies on Jetson, setting on-device cost, offline audit scores and 300+ blinded real-arm trials side by side.
- Found FP16 cheaper per inference than INT8 on Jetson Orin Nano, and offline audit rankings that failed to predict which INT8 build did better on the arm.

### AREL (AutoResearch-Experiment-Loop)

*ASUS Ascent GX10*

Personal research framework — *2026*

- Multi-agent plan–review–experiment–analysis loop, SFT/DPO data built from human evaluation, and a Ray cluster with Jetson Orin Nano and Raspberry Pi 5 for staleness-aware learning experiments.

### LabLens

*GitHub · Demo video*

Google DeepMind AI Hackathon — *2026*

- A review tool for real SO-101 robot recordings (two cameras plus joint logs). Gemini 4 Argon on Vertex AI plans the checks; fixed tools compute every number from raw CSV and decode exact frames, and each claim links to its source frame or log window.

## Skills

| | |
|---|---|
| **Machine Learning** | Quantization and on-device deployment (TensorRT FP16/INT8, Hailo-8 INT8, Coral Edge TPU INT8, Qualcomm QNN HTP), reinforcement learning (offline, meta-RL, safe RL), statistics and data analysis, tinyML |
| **ML Frameworks** | JAX / Flax, PyTorch, TensorFlow / TFLite, ONNX Runtime, vLLM, Ray, LeRobot |
| **Systems** | Embedded systems, operating systems, Linux kernel, appliance–PC communication interfaces (UART, SRIO) |
| **Languages & Tools** | Python, C, R, Apache Spark, Jenkins, Parasoft, Bokeh |
| **Hardware** | NVIDIA Jetson Orin Nano, Raspberry Pi 5 + Hailo-8, Google Coral Edge TPU, Qualcomm QCM6490, ASUS Ascent GX10 |
| **Languages** | Korean (native), English |

## Education

**Seoul National University** — *Seoul, South Korea*

M.S. in Electrical Engineering & Computer Science — *Mar 2013 – Feb 2015*

- Advisor: Prof. Bernhard Egger. Thesis: *Hierarchical Power Management Framework on Manycore Systems Using OS Migration Techniques*, OS-level DVFS power management on the Intel Single-chip Cloud Computer.

**Chung-Ang University** — *Seoul, South Korea*

B.S. in Electrical and Electronics Engineering — *2006 – 2012*

## Publications

[1] **C. Kang**. "Same Policy, Different Artifacts: Paired Task Audits for On-Device Neural Control." *NeurIPS 2026 Workshop on On-Device Intelligence: Foundation Models under Real-World Constraints (ODI)*, 2026.

[2] **C. Kang**. "When Reward and Safety Need Different Model Horizons: A Selection Theorem for Pessimistic Offline Planning." *RLC 2026 Workshop on Model-based RL in the Era of Generative World Models (WMW)*, 2026.

[3] J. Lee, **C. Kang**, G. Bae. "Sensor-Based Reward Learning from Video Labels for Tumble Motion Control in a Household Dryer." *IEEE 22nd International Conference on Automation Science and Engineering (CASE)*, 2026.

[4] **C. Kang**, G. Bae, D. Kim, K. Lee, D. Son, C. Lee, J. Lee, J. Lee, J. W. Yun. "Data-Driven Reinforcement Learning for Optimal Motor Control in Washing Machines." *IEEE Conference on Artificial Intelligence (CAI)*, pp. 418–424, 2024.

[5] **C. Kang**. "실물 환경에서의 데이터 기반 강화학습 적용 방안" [Applying data-driven reinforcement learning in real-world environments]. *The Magazine of the IEIE* (전자공학회지) 51(5), pp. 56–63, 2024.

[6] C. Ahn, S. Lee, **C. Kang**, B. Egger. "NuPow: Managing Power on NUMA Multiprocessors with Domain-Level Voltage and Frequency Control." *17th International Conference on the Economics of Grids, Clouds, Systems, and Services (GECON)*, Sep 2020.

[7] B. Kim, D. Lee, **C. Kang**, H. Ko, W. Kim. "Learning and Auto-Tuning PID Hyper-Parameters from Scratch without Human's Domain Knowledge by Reinforcement Learning." *SAREK Conference* (대한설비공학회 학술발표대회), pp. 297–299, 2018.

[8] **C. Kang**, S. Lee, Y.-J. Lee, J. Lee, B. Egger. "Scheduling for Better Energy Efficiency on Many-Core Chips." *19th Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP)*, Hyderabad, India, pp. 46–68, May 2015.

[9] 강종구, **강찬석**, 서은아, 이원주. "Windows 8 & Kinect for Windows." *KSCI Review* 20(1), pp. 39–43, 2012.

## Patents

**C. Kang**, D. Son, D. Kim, K. Lee, G. Bae. "Device for Collecting or Transmitting Data for Dehydration Control or Updating Dehydration Control Model Using Collected Data, or Method Therefor." WO 2025/216419 A1 (PCT/KR2025/001981), LG Electronics. Filed Feb 11, 2025; published Oct 16, 2025; pending. National and regional phase: CN 202580025081.7 (entered Sep 28, 2026), EP 25786891 (entered Oct 14, 2026), KR 10-2026-7033261, US Appl. No. 19/881,721.

## Talks

| | | |
|---|---|---|
| Sep 2025 | **Open Mentoring Session,** LG Software Developer Conference (SDC) | *LG Electronics* |
| Sep 2024 | **"실물 환경에서의 데이터 기반 강화학습 적용 방안",** SNU CSE ComSee [Applying data-driven RL in real-world environments] | *SNU* |
| Nov 2022 | **"Improving Model Performance with Real-World Data in Washing Machine",** LG SDC (press) | *LG Electronics* |
| Nov 2020 | **"Reinforcement Learning from Scratch (Revisited)",** LG Innotek SW Day | *LG Innotek* |
| Jul 2018 | **"Reinforcement Learning from Scratch",** Software Engineer's Energizing Day (SEED) (press) | *LG Electronics* |

## Honors & Books

| | | |
|---|---|---|
| 2026 | **25th Place,** NAS Unseen-Data Challenge, AutoML 2026 | |
| 2022 | **Excellent Academic Book,** National Academy of Sciences, for 그로킹 심층 강화학습 | |
| 2021 | **Translator,** 그로킹 심층 강화학습, Korean edition of *Grokking Deep Reinforcement Learning* (M. Morales) | *Hanbit Media* |
| 2019 | **AI Specialist,** LG Electronics, first cohort of 12 (press) | |
| 2019 | **Code Reviewer,** 텐서플로를 활용한 머신러닝 | *Hanbit Media* |