STELLA HWANG

510-703-6305 | hwang020612@gmail.com | Berkeley, CA | linkedin.com/in/inseonhwang | github.com/inthree3

PROFESSIONAL EXPERIENCE

Center for Human-compatible AI, UC Berkeley **June 2026 - Present**
Research Engineer *Berkeley, CA*

- Engineered infrastructure for MySky, a customizable Bluesky recommendation feed algorithm and control UI serving **747K MAU**
- Built an automated incident response bot routing on-call alerts from **GCP** and surfacing contextual runbooks with Git, **Discord API** slashing **MTTA to < 15 minutes** across incidents ranging from Web Vitals regressions to recommendation feed pipeline degradation
- Implemented a user prompt-driven personalized feed with LLM API, and a multi-key API auth system
- Led the migration of Elasticsearch index with parallelized **Go** workers, fixing a **~0.6% of 2B+** post misclassification bug
- Shipped a PostHog-gated fail-fast pipeline, UTM redirector enabling failure modes to be rolled out to internal users only

National Assembly of Korea, Office of Representative Junseok Lee **August 2024 - December 2025**
Software Engineer *Seoul, Korea*

- Built and operated an end-to-end digital platform for a National Assembly member using **TypeScript, NestJS, Node.js, TypeORM, and PostgreSQL**; developed constituent case management workflows to process **270+ resident inquiries**, alongside an aggregated news feed.
- Built and operated multiple services on **Vultr** using **Docker, Kubernetes, Helm**, and **Nginx**, with **GitOps**-based deployments via **Argo CD** and CI/CD automation through **GitHub Actions**
- Built identity-verified API enabled political voting with Redis, collecting **21,694** verified responses in 2 days, later used as legal evidence in court
- Developed a **React** congratulatory video request platform and a **.NET** application for automated video editing and delivery, reaching approximately **10% of K–12 schools** across South Korea within two weeks
- Built a donation platform integrating payment APIs with OBS for real-time on-stream donor reactions and donor-record management compliant with Korea's Political Funds Act; raised **$2.5M+** from **50,000+** donors, ending the election period with a budget surplus

Visual AI Lab, Gwangju Institute of Science and Technology **April 2023 - June 2023**
Research Student *Gwangju, Korea*

- Built a video-to-3D reconstruction pipeline for clothed subjects (frame extraction, MiVOS video-object segmentation, COLMAP structure-from-motion pose estimation) feeding NeRF-based reconstruction, applied to 20 videos of a clothed mannequin
- Led a team of 11 researchers in reproducing SoTA NeRF and Instant-NGP using PyTorch
- Selected for two technology entrepreneurship programs, receiving a sponsored UC Berkeley opportunity and ~$22K (₩30M) funding

GroupbyHR **April 2021 - August 2021**
Software Engineer *Gwangju, Korea*

PROJECTS & OUTSIDE EXPERIENCE

Checky: Real-Time Voice Shopping Assistant for Smart Glasses **July 2026**

- Built a real-time shopping assistant that streams camera and audio from a **Next.js** client to a **Python** LiveKit worker running the **Gemini Live API**
- Replaced slow live scraping with a **Moss** semantic cache seeded offline from Bright Data snapshots, giving sub-second price lookups across Amazon, Walmart, and Target
- Made the assistant report a catalog miss and suggest alternatives instead of fabricating a price

MoatGraph: Technical Moat Recommendations from Papers and Patents **July 2026**

- Built a **TypeScript** pipeline that splits a business problem into sub-problems, maps each to academic concepts, searches and scores papers and patents, and generates one moat recommendation per sub-problem
- Visualized results as an interactive graph in a **React** frontend, backed by a graph schema and an **OpenAPI** (Swagger) spec

Seeya: Accessible E-Commerce Browser Extension for Visually Impaired Users **March 2024 - May 2024**
Gwangju, Korea

- Built a local Chrome extension for shopping, chaining quantized Whisper STT, a 4-bit vision-language model, and TTS
- Quantized the VLM, cutting weight memory ~3.5× while retaining 96.5% accuracy on a 200-item product-question benchmark
- Designed a layered cache (pHash, semantic similarity, and KV-cache reuse), lowering p50 latency from 2.8s to 0.6s at a 72% hit rate

EDUCATION

University of California - Berkeley **January 2026 - December 2026**
Global Computing and Data Science

Gwangju Institute of Science and Technology
Bachelor of Science (BS), Electrical Engineering and Computer Science

- **Relevant Coursework:** Machine Learning, Deep Learning, Data Structure, Linear Algebra, UI Design, Algorithms, Programming Languages and Compilers, Signals and Systems, Computer Systems (double TA), Network, OOP (double TA)

SKILLS

Skills: Python, PyTorch, SQL, Docker, Kubernetes, Google Cloud Platform (GCP), Git, Elasticsearch, OpenTelemetry (OTel), Go, Rust, JavaScript, TypeScript, C++, Argo CD, Helm, Nginx, CI/CD, React
Languages: Korean, English