

**CHO, SEONGLAE**  
Gower St, London WC1E 6BT  
+44 7356-161248 | seonglae.cho.24@ucl.ac.uk

## EDUCATION

### UNIVERSITY COLLEGE LONDON

*Artificial Intelligence for Sustainable Development MSc*

### YONSEI UNIVERSITY

*Computer Science BE*

London, England, United Kingdom  
September 2024 - Present  
Seoul, South Korea  
March 2017 - August 2024

## PUBLICATIONS

- Cho, S., Jang, M., Yeo, J., & Lee, D. (2023). *RTSUM: Relation Triple-based Interpretable Summarization with Multi-level Salience Visualization*. In Proceedings of the NAACL 2024 System Demonstrations Track. Association for Computational Linguistics. <https://arxiv.org/abs/2310.13895>
- Cho, S. (2025). SAE Training Dataset Influence in Feature Matching and a Hypothesis on Position Features. *LessWrong*. <https://www.lesswrong.com/posts/ATsvzF77ZsfWzyTak/dataset-sensitivity-in-feature-matching-and-a-hypothesis-on-1>

## EXPERIENCE

### HOLISTIC AI

#### *AI Research Engineer Intern*

London, United Kingdom  
May 2025 - Present

- Implemented an evaluation pipeline for Deep Research AI Agent using OpenSSF baseline metrics to assess method performance

Seoul, South Korea

### KAKAO MOBILITY

#### *Software Engineer, Digital Twin Team*

December 2021 - September 2022

- Led a 3-person team in developing a national-scale 3D mapping service as a part of the Autonomous Driving pointcloud pipeline
- Ported a C++ 3D-projection algorithm to Rust with Node.js bindings, making the library cross-platform

Seoul, South Korea

### STRYX

#### *Software Engineer, 3D Mapping Team*

Seoul, South Korea

- Reduced build time by 70% and simplified dependency management by merging multiple repositories into a monorepo

November 2019 - December 2021

## AWARDS

### HERMES, 1ST PLACE (£3,000) , HOLISTIC AI HACKATHON (2024)

London, England, United Kingdom  
November 2024 – November 2024

#### *Team Lead*

- Fine-tuned Sparse AutoEncoder (SAE) for GPT-2 to identify and steer correlated features for multiple biases for AI Safety
- Reduced stereotypical text generation by 20% from an initial 90% rate by applying a Steering Vector derived from the SAE

Seoul, South Korea

### MBTIGPT, 1ST PLACE (₩3,000,000), YONSEI GENAI COMPETITION (2023)

September 2023 - January 2024

#### *Team Lead*

- Built an end-user AI service that employs RAG on user chat history by an MBTI personality analyzer with Redis and Faiss

## PROJECTS

### UNIVERSITY COLLEGE LONDON

#### *ControlRL*

London, England, United Kingdom  
January 2025 – April 2025

#### *FaithfulSAE*

January 2025 – February 2025

- Increased Gemma-2 2B's MMLU score from 51.8% to 55.5% via a novel SAE-based automated RL control method

#### *YONSEI UNIVERSITY DATA & LANGUAGE INTELLIGENCE LAB*

Seoul, South Korea

#### *ReSRer*

September 2023 - January 2024

- Indexed 21M Wikipedia-scale corpus into a Milvus vector database and accelerated LLM training through distributed training

March 2023 - August 2023

#### *RTSum*

- Published as the first author, designed Knowledge Graph (KG)-based summarization experiment for Interpretable AI framework

Seoul, South Korea

#### *YONSEI UNIVERSITY*

June 2023 - July 2023

#### *LLaMa2GPTQ*

- Reduced LLaMa2 memory cost by 75% with 4-bit GPTQ quantization and integrated RAG vector search for Local LLM

November 2021 - June 2023

#### *Texonom*

- Built an ANN-based vector retrieval API by embedding all 30,000 content pages in service into Postgres pgVector database

## SKILLS

- Service: Full stack | TypeScript | Node.js | PostgreSQL | Redis | Rust | Vector Database | Faiss | Milvus | Vite | React | Streaming
- AI: Python | PyTorch | AI Agent | AI Cost Optimization | DDP | FSDP | ONNX | Pydantic AI | RAG | LangGraph | AI Evaluation
- Infra: Kubernetes | Linux | Git | CI/CD | Docker | Docker Compose | Ansible | ETL | Github Action | Hadoop | Distributed Systems