

Seongmin Park

seongminpark.com | github.com/seongminp | llov0708@gmail.com | +821047349158 | Google scholar

EDUCATION

University of Pittsburgh	M.S. in Information Science (Dec. 2019)
Carnegie Mellon University	B.S. in Electrical and Computer Engineering (Dec. 2017)

WORK EXPERIENCE

Goodnotes Remote Jun. 2025 – Present goodnotes.com	Machine Learning Engineer Built and shipped LLM-powered meeting assistant features (live summaries, agenda notes, context-aware chat) in Goodnotes, the #1 global note-taking app used by tens of millions worldwide. Deployed on-device ML models to iOS, macOS, Windows, Android, and Web clients, using ONNX, OpenVINO, and QNN to run production speech-to-text and language models on-device. Fine-tuned and prompt-tuned models for real-time summarization of speech , improving quality and responsiveness for live meeting use cases. Built and scaled ML evaluation pipelines to improve accuracy and latency , enabling faster iteration and more reliable production releases. Partnered cross-functionally with product, design, QA, and engineering to move quickly from research prototypes to production features at scale.
ActionPower Seoul, Korea Jun. 2020 – Jun. 2025 actionpower.kr	Language Team Lead Led a team of three in developing language AI for the "Daglo" speech-to-text service. Employed from post-seed to post-Series B, contributing to a 125-fold MAU increase during tenure. Used Python and large language models to implement transcript summarization, retrieval-augmented generation (RAG), model context protocol, instruction tuning, alignment tuning, keyword extraction, error correction, text normalization, and knowledge distillation . Collaborated with Seoul National University Hospital to develop a model for suicide risk assessment. Developed proof-of-concepts for over 20 B2B and B2G projects (React, FastAPI, Gradio, Streamlit). Contributed to frontend development of Daglo's web application.
OkeyDokey Seoul, Korea May 2017 – Aug. 2018 white paper	Smart Contract Developer Developed open-source smart contracts on the Ethereum Blockchain with Solidity to control Internet of Things devices. Developed Android (Java) and iOS (Swift) mobile clients. Modified the go-ethereum (Go) Ethereum client to run on Android NDK . Acquired by Colong E&C after tenure.
SK Techx Seoul, Korea May 2016 – Aug. 2016	Intern Mobile Developer (Volo team) Maintained an iOS (Objective-C) travel journal social network application with approximately 80,000 users. Aided in porting an existing iOS app into Android.

PUBLICATIONS AND AWARDS

Awards 2023, 2024	1st place, Korean AI Competition 2023 (1st out of 233 teams) 1st place, Korean AI Competition 2024 (1st out of 280 teams) LLM instruction tuning, DPO, RLHF, ensemble, summarization, and classification.
-----------------------------	---

Selected publications

[JMIR Med Inform 2025] *Automated Speech Analysis for Screening and Monitoring Bipolar Depression*
[ICASSP 2025] *Prompt Crossing: Evaluating Whether LLM Response Stems from Jailbreak or Normal Prompt*
[COLING 2024] *RT-VQ2A2: Real-Time Vector Quantized Question Answering with ASR*
[COLING 2024] *Non-Essential is NEcessary: Order-agnostic Multi-hop QG*
[ICASSP 2024] *Unsup. Extractive Dialogue Summarization In Hyperdimensional Space*
[Interspeech 2023] *Unsup. Dialogue Topic Segmentation in Hyperdimensional Space*
[ACL Findings 2023] *Cross-task Knowledge Transfer for Extremely Weakly Supervised Text Classification*
[COLING 2022] *LIME: Weakly Supervised Text Classification Without Seeds*
[ACL 2022 WS] *Unsup. Abstractive Dialogue Summarization with Word Graphs and POV Conversion*
[EMNLP 2021 WS] *Finetuning Pretrained Transformers into Variational Autoencoders*

Patents

[US 11,749,260] *Method for speech recognition with grapheme information*
[US 11,887,002] *Method of generating data by using artificial neural network model having encoder-decoder structure*
[US 11,721,325] *Method for generating learning data for speech recognition error detection*
[US 17/962,011] *Method of generating summary based on main speaker*
[US 12,197,862] *Method for identify a word corresponding to a target word in text information*
[US 11,640,493] *Method for dialogue summarization with word graphs*
[KR 10-2444457] 단어 그래프를 이용한 대화 요약 생성 방법 (KR 10-2444457)
[KR 10-2478076] 음성 인식 오류 검출을 위해 학습 데이터를 생성하기 위한 방법
[KR 10-2478763] 자소 정보를 이용한 음성 인식 방법
[KR 10-2480423] 인코더-디코더 구조를 가진 인공 신경망 모델을 활용한 데이터 생성 방법
[KR 10-2517661] 텍스트 정보에서 타겟 단어에 대응하는 단어를 식별하는 방법
[KR 10-2520167] 비-대화형 텍스트 정보를 활용하여 대화 요약을 위한 학습 데이터를 생성하는 방법
[KR 10-2537173] 주 화자에 기초하여 요약문을 생성하는 방법
[KR 10-2545892] 텍스트 분류를 위한 학습 데이터를 생성하는 방법
[KR 10-2647510] 초차원 컴퓨팅을 이용하여 텍스트를 분할하는 방법
[JP 7,475,592] テキスト情報の中からターゲット単語と対応する単語を識別する方法

Hama NLP hamanlp.org

Open-source **Zig**, **Python** and **Javascript** library for client-side (sub-megabyte) natural language processing. Features include **grapheme-to-phoneme conversion**, **part-of-speech tagging**, and **automatic speech recognition on the edge**.

ENGLISH EXAM SCORES

PTE Academic Sep. 2025

90/90 (Listening 90, Reading 90, Speaking 90, Writing 90)