SU PARK

suinthelab@gmail.com | https://sunotsue.github.io/ | U.S. Citizen

Education

Carnegie Mellon University, School of Computer Science, Language Technologies Institute Pittsburgh, PA
M.S. in Artificial Intelligence 2021 – 2023

Columbia University, Columbia College New York, NY
B.A. in Economics-Statistics & Linguistics 2013 – 2017

Research

- **SFT-GO: Supervised Fine-Tuning with Group Optimization for Large Language Models** Developed a token importance-aware post-training method for improving training signal utilization in large language models. *arXiv, 2025.*
- **Small Language Models: Architecture, Evolution, and the Future of Artificial Intelligence** Proposed a multi-axis taxonomy for classifying small language models and synthesized emerging approaches to capability-efficiency tradeoffs. *arXiv, 2025.*
- **Harnessing Business and Media Insights with Large Language Models.** Pre- and post-trained a domain-specialized LLM for business intelligence and media analysis for Fortune Magazine. *arXiv, 2024. Video.*
- **Model Probing and Capability Attribution.** Developed a probing framework for identifying which latent linguistic signals drive model predictions, enabling capability attribution, systematic error analysis, and instance-level failure diagnosis. *Report, 2022.*
- **Behavioral Effects of Model Compression.** Investigated how pruning and quantization alter learned representations and downstream model behavior, revealing compression-induced shifts in calibration, prediction dynamics, and output distributions across diverse task settings. *Report, 2022.*
- **Low-Resource Machine Translation.** Investigated data augmentation and transfer learning strategies for low-resource machine translation. Combined target-side monolingual data augmentation with LangRank-guided transfer language selection, improving BLEU scores by up to 345% over baseline systems for Belarusian-English and Azerbaijani-English translation. *Report, 2022.*
- **Low-Resource Multilingual ASR.** Investigated tokenization and self-supervised representation learning for low-resource speech recognition, evaluating HuBERT, wav2vec 2.0, language model integration, and Byte Pair Encoding vocabulary design in African-accented French. *Report, 2022.*
- **Temporal Action Localization.** Investigated architectural approaches for long-range temporal reasoning in video understanding, developing extensions to Boundary-Matching Network that improved ActivityNet-1.3 localization performance by 0.9 AUC through temporal feature propagation and global context aggregation. *Report, 2021.*

Work Experience

Turing San Francisco, CA
Member of Technical Staff, Post-Training & Evals February 2026 – Present

- Leading Terminal Bench. Designing reward functions and human/synthetic data generation pipelines for post-training and env quality evaluation. Researching data attribution and data efficiency via failure modes and feedback signals.

Accenture Mountain View, CA
Tech Lead, AI Research Scientist November 2023 – January 2026

- Led end-to-end post-training research for various domain-specific language models, spanning distributed training, data curation, and evaluation to study the effects of data quality and training signals on downstream behavior.
- Developed synthetic data generation and pre/post-training pipelines for a domain-specialized LLM built on 50+ years of proprietary Fortune Magazine content. Designed benchmark suites and evaluation methodologies spanning temporal reasoning, factuality, retrieval, knowledge boundaries, model personality, and safety alignment.
- Authored SFT-GO, a token importance-aware post-training method that improves training signal utilization. Designed novel optimization objectives and conducted large-scale training and hyperparameter studies across Llama and Qwen model families.

Nearpod Brooklyn, NY
Data Scientist January 2020 – April 2021

- First and only data science hire. Built API pipelines, predictive modeling, metric definition, and dashboard development. Played a key role in $650M acquisition, providing data-driven insights to support the process.

IBM New York, NY
Data Scientist August 2017 – January 2020

- Led a 30-person team across marketing, product, and design. Launched 200+ personalization experiments, led 30+ user interviews, and leveraged funnel-based A/B testing to drive 230% engagement growth and $10M in revenue gains. Earned the 2019 Q1 CMO Award and the Manager's Choice Award.