

EMMANUEL AZUH MENSAH

4029 7th Ave NE, Apt 313, Seattle, WA 98105

☎ 617-949-9732 ✉ emmanuel.azuh@gmail.com 🌐 <https://emazuh.github.io/>

Research Interests & Skills

Edge AI, Federated Learning, Speech & Multimodal Learning, Representation Learning, Hardware-Algorithm Co-Design

Deep Learning: PyTorch, Transformers, CNNs, Mixture-of-Experts

Learning Paradigms: Self-supervised, Semi-supervised, Few-shot, Online Learning

Multimodal: Speech, Audio, Vision, Time-Series

Edge & Systems: Model compression, Quantization, Knowledge Distillation, On-device inference, Metaheuristics

Programming: Python (Pytorch, Tensorflow, Jax), Java, C/C++, Linux, AWS

Signal Processing: Speech and time-series fundamentals

Research & Industry Experience

Graduate Research Assistant (Edge AI & Multimodal Learning)

University of Washington

Seattle, WA

Jan 2021 – Present

- Developed **edge-deployable multimodal deep learning models** combining vision and sensor data for wildlife monitoring under compute and power constraints.
- Designed **mobile-efficient transformer architectures** inspired by MobileViT and Mixture-of-Experts for conditional computation on edge devices.
- Conducted accuracy-latency tradeoff analyses for embedded deployment; research supported by the **U.S. NAIRR Pilot** and the **Azure AI for Earth** grant.
- Led end-to-end research projects from problem formulation and model design through large-scale experimentation, analysis, and manuscript preparation. Ongoing collaboration with **Ian Quah** on **neuro-inspired efficient deployable models**.
- Ongoing collaboration with Iyer lab on Computer Vision for battery free insect scale robots (late-breaking poster at IEEE ICRA).

Research Scientist Intern (Applied ML)

Amazon

June 2023 – Sept 2023

- Led an applied ML research project designing **novel learned metrics** for product quality and customer experience.
- Developed custom **loss functions and regularization strategies** to ensure temporal stability and robustness across domains.
- Authored internal technical reports and experimental analyses that guided cross-team design decisions and informed model deployment strategies.

Student Researcher (Multimodal Perception)

Google

June 2022 – Sept 2022

- Proposed approaches for **multimodal representation learning** from video, audio, and text.
- Investigated **on-device vs server-side learning** tradeoffs, including model compression and inference constraints.
- Evaluated cross-modal alignment under **limited labeled data** regimes.

Graduate Research Assistant (Speech & Vision)

MIT Spoken Language Systems Group

Sept 2018 – Sept 2019

- Developed CNN-based models for **bilingual speech representation learning** using images as an interlingua.
- Studied **weakly supervised and cross-lingual transfer** learning for speech translation.
- Published at **Interspeech 2019**.

Undergraduate Research Intern (Online Learning)

June 2018 – Aug 2018

Microsoft Research – FarmBeats

- Applied **online deep learning** to agricultural time-series data, achieving a 50% MAE reduction in forecasting.
- Used **representation learning** to synthesize virtual sensor signals from low-cost hardware.

Software Development Engineer I

Sept 2019 – Dec 2020

Amazon Web Services

- Built high-throughput, low-latency networking systems supporting large-scale ML services.

Selected Research Projects (Edge AI & Systems)

Adaptive Federated Learning in a 2G Context

UW CSE 561

w/ Prof. Ratul Mahajan

- Simulated **2G network conditions** using AWS containers configured as edge nodes.
- Studied convergence behavior and communication bottlenecks in federated learning under extreme bandwidth constraints.
- Implemented **model weight compression** and analyzed accuracy, convergence speed, and network utilization tradeoffs.

FPGA-Based Matrix Multiplication Profiling

UW CSE 548

w/ Prof. Michael B. Taylor

- Built a profiling framework on a **PYNQ FPGA** with the Black Parrot processor.
- Analyzed effects of data reuse strategies on **pipeline stalls and cache misses**.

Hardware Acceleration for Graph Convolutions

MIT 6.888

w/ Prof. Vivienne Sze & Prof. Joel Emer

- Co-designed a **weight-stationary dataflow** for graph convolution workloads.
- Extended a simulator to quantify **MAC count and energy savings**.

Ongoing & Applied Research

- Deploying efficient neural models on **Arduino-class ultra-low-power devices** using **quantization and knowledge distillation** in battery-free solar-powered setups.

Leadership & Research Outreach

Founder – African Youth Research Program (AYRP)

- Founded and registered an **undergraduate student organization** at UW, introducing African students to **AI research methodology**.
- Ran a recent workshop on **using LLMs as undergraduate research assistants** for students at KNUST.
- Designed curricula for research in **low-resource and limited-supervision settings**.