

Sonika Vuyyuru

svuyyuru7@gatech.edu | (703) 884-7028 | [linkedin.com/in/sonikavuyyuru](https://www.linkedin.com/in/sonikavuyyuru) | sonikavuyyuru.com

EDUCATION

Georgia Institute of Technology | *M.S. in Computer Science, Machine Learning (Thesis Track)* | **GPA: 4.00** *Expected May 2027*

- Thesis: Equivalence and Differences Between Contrastive Learning and Generative Models in Neural Representation Learning
- Relevant coursework: Machine Learning for Neural and Behavioral Data Analysis [A], Machine Learning [A]
- Awards: NSF Graduate Research Fellowship Program (GRFP)

University of California, Berkeley | *B.A. in Computer Science (Honors) and Cognitive Science* | **GPA: 3.86** *May 2024*

- Relevant coursework: Deep Learning [A], Machine Learning [A], Artificial Intelligence [A], Optimization Models [A-], ML for Hardware Design [A], Algorithms [A-], Discrete Math and Probability [A], Digital Design and Integrated Circuits [A, A+]
- Awards: EECS Honors Program (110-person cohort), Cal Alumni Association Leadership Merit Scholarship

INDUSTRY EXPERIENCE

WHOOOP

Boston, MA

Machine Learning Engineer Co-op

Jan. 2026 - August 2026

- Developed and benchmarked self-supervised learning objectives for transformer-based physiological signal foundation models
- Conducted representation analysis and ablation studies on a masked autoencoder trained on multi-channel PPG data, informing optical sensing hardware design at WHOOOP

NASA, Glenn Research Center | tinyurl.com/sonika-nasa-ml

Cleveland, OH

Machine Learning Engineer Intern

May - Aug. 2023

- Authored paper on ML-driven engine model optimization published at AIAA SciTech Forum 2024
- Designed and trained neural network using **TensorFlow** to optimize control laws, reducing transient stall margin in novel hybrid-electrified gas turbine airplane engine, improving energy efficiency and **lowering operational costs by a projected 20%**

Multiply Labs

San Francisco, CA

Software Engineering Intern

May - Aug. 2022

- Led development of an automated ‘one-click’ IoT device setup, **reducing deployment time by 50%** and enhancing security and reliability across production systems through custom **AWS CloudFormation** stacks for Secrets, Alarms, and EC2 instances
- Implemented **Pytest**-based testing framework, leveraging **fixtures, context managers, mocks** to **reduce testing time by 30%**

ACADEMIC RESEARCH

Brain Intelligence and Machine Learning Lab, Georgia Institute of Technology

Atlanta, GA

Graduate Research Assistant, advised by Prof. Anqi Wu

Aug. 2025 - Present

- Designing and implementing experiments comparing contrastive learning methods (e.g., CEBRA) and generative models (e.g., pi-VAE, Swap-VAE) for neural representation learning, with focus on identifiability, disentanglement, and performance

Neural Systems and Machine Learning Lab, Lawrence Berkeley National Laboratory

Berkeley, CA

Undergraduate Research Assistant, advised by Dr. Kristofer Bouchard

Aug. 2023 - May 2024

- Developed sDCA, statistical machine learning method for identifying behaviorally relevant neural subspaces from high-dim. data
- Implemented and benchmarked sDCA on high-dimensional speech and synthetic datasets, evaluating performance against PCA
- Scaled large-scale benchmarking experiments using **Slurm**-based distributed HPC workflows to support reproducible evaluation

PROJECTS

Personalized Educational AI Assistant | *Python, HuggingFace, PEFT, LoRA* | sonikavuyyuru.com/ed

- Fine-tuned **LLaMA2 LLM** on course forum data using **Parameter-Efficient Fine-Tuning (LoRA)**, **reducing trainable parameters by 50%** while enabling context-aware Q&A and maintaining accuracy for resource-constrained deployment

Audio Convolutions & Probabilistic Audio Generation | *Python, Google Magenta, Max/MSP, Orchidea*

- Engineered generative audio synthesis pipeline combining Markov Chain transition matrices and Google Magenta to algorithmically compose Bach-style chorales and polyphonic structures
- Implemented 1D/2D audio convolutions and spectral feature extraction in Max/MSP and Orchidea to blend timbre and perform real-time sound hybridization across multi-channel audio streams

FPGA RISC-V CPU, Apple Hardware Outstanding Project Award Winner | *Verilog* | sonikavuyyuru.com/riscv

- Built 3-stage pipelined RISC-V processor on **FPGA** with full instruction support, **memory-mapped I/O**, and UART interface
- Implemented an optimized global history branch predictor (5-bit register) achieving **96.3% accuracy** and **1.051 CPI**
- Selected for Outstanding Project Award by a panel of Apple hardware engineers from a class of 130 students

TECHNICAL SKILLS

- Languages: Python, C, Java, JavaScript, SQL, Verilog, MATLAB, RISC-V
- Tools & Frameworks: TensorFlow, PyTorch, Scikit-learn, SciPy, HuggingFace, NumPy, Simulink, React, Django, AWS, Git, Jira