Kyle Marino

kmarino@alumni.usc.edu | kylemarino22 | kylemarino.me

Education

University of Southern California – Los Angeles, CA

- **M.S. Computer Engineering (2024) - GPA: 4.0**
- **B.S. Electrical and Computer Engineering (2023) - GPA: 3.93**
  - Honors: Presidential Scholarship, Summa Cum Laude
  - Relevant Coursework: Computational Intelligence and Neural Learning, Internet and Cloud Computing, Computer Systems Organization, Deep Learning, Probability Theory, Linear Algebra

Experience

Eta Compute – Sunnyvale, CA

Embedded ML Engineer | Jan 2025 – Current

- Developed demo for gesture and face detection models on NXP's MCXN947 in C++. Pipeline: DMA from camera, downsampling, inference on NXP's ML accelerator, video over UART to Python GUI
- Wrote low-latency image downsampling algorithm to process camera stripes without losing pixel data
- Built Python software for feature extraction on ECG data. Supported FFT, bandpass filters, min-max / z-score normalization. Models achieved a SOTA test F1 of 95% across 7 heart pathologies

Apple – Cupertino, CA

FPGA Architecture Intern | May 2024 – Nov 2024

- Developed Python driver and Verilog AXI interface for SPI DAC peripheral on MIX FPGA
- Integrated DisplayPort IP at 8k 60 Hz over CIO80 on Agilex7 FPGA with hardware testing platform
- Wrote SERDES link training module and PRBS generator in Verilog for 1.5 GHz loopback tests

AMD/Xilinx – San Jose, CA

RTL Integration Intern | May 2023 – Aug 2023

- Built a distributed design rule verification system (1000+ jobs) in Python to analyze full-chip RTL.

SOC Integration Intern | May 2022 – Aug 2022

- Designed and integrated Python and Perl tools for automated EMIR reporting and simulation processes.

USC Rocket Propulsion Laboratory – Los Angeles, CA

Lead FPGA Engineer | Sep 2019 – Jan 2023

- Built 13 channel data collection and processing platform on Intel MAX10 FPGA recording at 10 MB/s
- Designed microprocessor and Python compiler for real-time quaternion integration of IMU data
- Wrote C++ drivers to interface with FPGA and battery management boards
- Wrote Verilog modules for SD Bus, I2C, and SPI interfaces for ADCs, IMUs, and off-chip memories

Really, Inc. – Mountain View, CA

Backend Development Intern | Jun 2019 – Mar 2020

- Created data preprocessing pipelines for a machine learning platform analyzing social media and email data to detect early signs of Alzheimer's disease
- Built a linguistic analysis tool using BERT, spaCy, and NLTK to measure sentiment shifts and lexical complexity over time, storing results in PostgreSQL for long-term tracking

Research / Classwork

Directed Research – Los Angeles, CA

Advised by Dr. Viktor K. Prasanna | Feb 2023 – May 2024

- First Author on **ME-ViT: A Single-Load Memory-Efficient FPGA Accelerator for Vision Transformers** (HiPC 2023 – **Best Paper Award**) [link]
- Designed a Vision Transformer accelerator for FPGAs, achieving up to a 17.89x reduction in memory bandwidth and a 2.16x improvement in throughput per DSP over SOTA designs

Computational Intelligence and Neural Learning Final Project – EE 689 | Aug 2023 – Dec 2023

- Novel method for generating adversarial input tokens to force specific hallucinations for Llama 2 7B
- Implemented reversed back-propagation to calculate embedding vectors for targeted outputs and performed nearest-token searches to identify adversarial input sequences

Internet and Cloud Computing – EE 542 | Aug 2022 – Dec 2022

- Designed a custom file transport protocol with multithreaded UDP sockets that outperforms TCP over a high loss link, achieving over 75 Mbps out of 80 Mbps theoretical
- Modified TCP Linux kernel to improve throughput over a high loss link, from 480 Kbps to 10 Mbps
- Led hardware development of a low-power GPS tracking solution, designing power-switching circuits and detection algorithms for a highly efficient prototype with real-time web interface

Personal Projects

Pysystemtrade | Jul 2024 – Current

- Contributor to open source quant trading platform originally developed by Rob Carver
- Built server to record live data on 400 instruments (~4k reqs/s) using asyncio and multithreading
- Wrote portfolio weighting algorithm using bootstrapping and Markowitz mean-variance optimization

Live GPT | Jan 2025 – Current

- Fine-tuned GPT-4o for processing live transcriptions for real-time, conversational interactions
- Designed a retraining algorithm that critiques responses and generates corrected fine-tuning examples

Quadcopter | Dec 2016 – Jun 2019

- Designed from scratch self-stabilizing quadcopter with integrated camera streaming
- Wrote all embedded software, including flight controller with PID for stability, and sensor drivers for serial communication

3D Physics Engine | Nov 2019

- Built GPU-accelerated real-time physics and rendering engine in Java using LWJGL
- Designed custom collision detection algorithm suitable for arbitrary object models

Leadership

USC Climbing Team Captain | Aug 2021 – May 2024

- Led weekly team practices and develop group workouts for competition training
- Mentored members on technique, strength training, and injury rehab

Other Activities

Eagle Scout	**2018**
SM Hacks Best Web Application	**2017**
Code Day Hackathon Best Web Application	**2016**
Rock Climbing, Guitar, Kendama	