Joseph Carl Santos

(760) 916-3175 | jsanto78@calpoly.edu | https://arcsin-santo.tech

Education

California Polytechnic State University, San Luis Obispo — San Luis Obispo, CA
Bachelor of Science in Computer Science — *Expected May 2027*

- Cumulative Transfer GPA: 3.806/4.0 (from California State University, San Marcos)
- Relevant Coursework: Data Structures and Algorithms, Computer Architecture, Linear Algebra, Calculus I–III, Discrete Math, Physics of Mechanics and Sound, Physics of Electromagnetism and Optics, Assembly and Digital Circuits

Experience

Software Engineer Intern — June 2025 – August 2025
Pure Holdings Inc. d/b/a One Take Audio — *Los Angeles, CA (Remote)*

- Built a model precaching system (C++/ONNX Runtime) that preloaded multiple ONNX models in memory, enabling quicker model switching at 150 milliseconds.
- Optimized system performance by profiling CPU/memory load, GPU utilization, and audio latency; proposed fixes for jitter and buffering issues.
- Improved CUDA cross-platform compatibility on Windows and documented environment requirements to enable reliable builds and future performance optimizations.

Learning Assistant — August 2024 – December 2024
California State University San Marcos Corporation — *San Marcos, CA*

- In partnership with faculty, supported and implemented active and collaborative learning activities that may lead to more equitable and effective learning outcomes
- Facilitated small group interactions, such as group activities and tutorials, to address student challenges with course content effectively
- Provide feedback to faculty about student perspectives and misconceptions
- Demonstrated leadership and effective communication skills to diverse student population

Projects

HiFi-GAN Architecture | *Python (NumPy, PyTorch), arXiv.org (research reference)* — March 2025 – June 2025

- Reimplemented the HiFi-GAN generator architecture for high-fidelity speech synthesis; the focus is on the Multi-Receptive Field (MRF) module
- Implemented custom residual blocks (ResBlocks) with varying kernel sizes and dilation rates, as described in the NVIDIA paper.
- Loaded `.npy` mel spectrograms from an external TTS model and successfully reconstructed audio waveforms, validating model output.

Voice Conversion Software | *C++ (Qt), Python* — August 2023 – May 2024

- Built a desktop application for voice conversion using an open-source Python model, with a C++ Qt-based frontend; C++ inference using ONNX Runtime
- Implemented audio waveform visualization for frontend by analyzing real-time input; develop a threading algorithm.
- Implemented C++ preprocessing to read and stream audio, while matching input dimensions and data formats to maintain compatibility with the original PyTorch-based model during inference
- Prototyped alternative real-time voice conversion backend on Python by generating a synthetic parallel dataset from non-parallel audio; trained on runpod.io cloud GPUs and managed a sub-100GB dataset on Google Cloud Storage.

Technical Skills

Programming Languages: C++, Python
Frameworks & APIs: PyTorch, Qt6
Libraries: NumPy, pandas, Matplotlib, librosa, torchaudio
Developer & Debug Tools: Git, GitHub, Docker, WinDBG, NVIDIA Visual Profiler, Intel VTune Profiler
Cloud & Platforms: Google Cloud Platform (GCP), NVIDIA Triton, Vast.ai, RunPod, Lambda Labs, GPU-based training