John Hartquist

Software Engineer · AI · ML · Systems

Roseville, CA | john@hartquist.com | (805) 305-4168 | johnhartquist.com | github.com/jhartquist
linkedin.com/in/john-hartquist

Summary

Software engineer building systems from prototype to production. 14 years shipping software, primarily at startups, with the last 8+ focused on ML and AI. MS in Computer Science from Cal Poly, with a thesis on real-time polyphonic guitar transcription. Most recently CTO at TakeOne, shipping ML video pipeline features (segmentation, audio cleanup, transcription, LLM script editing) across Python, iOS, Android, and Next.js. Open-source work spans real-time audio, Rust ML, and AI agent tooling.

Selected Independent Research & Open Source

resonators — Rust crate for real-time spectral analysis – github.com/jhartquist/resonators

- Rust implementation of Alexandre François's Resonate algorithm: per-bin spectral analysis without FFTs or windowing.
- Cross-platform: native on macOS, Linux, iOS, plus browser (WASM). Python (PyO3) and JavaScript bindings.
- Outperforms reference implementation (C++, Apple-only) by 1.6x on M2 Max while remaining portable.
- Live WebGL streaming-spectrogram demo runs on desktop and mobile.
- Write-up: johnhartquist.com/resonators.

groxide — Rust docs CLI for AI coding agents – github.com/jhartquist/groxide

- Converts rustdoc JSON into markdown so AI coding agents can query crate docs without spending tokens parsing HTML.
- Works for local projects, external dependencies, and any crates on crates.io.
- Ships an installable skill for coding agents.

rloop — DAG-based task executor for Claude Code – github.com/jhartquist/rloop

- Runs tasks autonomously in isolated git worktrees with verification (tests, clippy, build) gating merges.
- Retry-with-learnings: failed attempts produce structured findings that feed forward into the next attempt's context.
- 267 tests, clippy-pedantic clean.

candle-pitch — pitch detection models ported to Rust – github.com/jhartquist/candle-pitch

- End-to-end Rust ports of CREPE, Swift-F0, and Spotify's Basic Pitch via Hugging Face's Candle framework.
- Numerical-parity tested against the original TensorFlow / PyTorch implementations.

Real-Time Musical Analysis (RTMA)

- Cal Poly MS thesis (2012): real-time polyphonic guitar transcription using Non-Negative Matrix Factorization.
- Full Rust rewrite, GuitarSet evaluation, neural experiments with HCQT features and Mamba-family architectures.
- Adapted Spotify's Basic Pitch to run in real-time in the browser via `resonators` front-end.

Experience

Senior SWE → Staff SWE → Chief Technology Officer, TakeOne – Remote Aug 2023 – Mar 2026

- Shipped features across TakeOne's ML video pipeline: real-time on-device tracking, video segmentation (SAM, Track-Anything, GroundingDINO), color correction and matting, audio cleanup (DeepFilterNet + FFmpeg loudness normalization), transcription (Whisper, Deepgram), and LLM-driven script editing.
- Worked across the stack: Python ML, Swift iOS, Kotlin Android, Django backend, Next.js user-facing applications, and AWS GPU inference (EC2, ECS, Batch).
- Built a transcript-editing pipeline with an LLM-as-judge: structured Pydantic rubrics, hallucination-catching validation gates, and retry-with-feedback up to 5 attempts.
- Wound the company down in March 2026.

Senior Software Engineer, Lowe's Innovation Labs – Kirkland, WA Sept 2021 – July 2023

- Prototyped AR measurement and virtual-object simulation in Unity for in-store and at-home use cases.
- Built React / Babylon.js inspector and ingestion tool for 3D home assets; published a custom open-source product viewer.
- Built a room-scanning prototype using Apple's RoomPlan API to capture 3D room geometry; researched data stores (EdgeDB, Solid) for structured home-asset modeling.

Senior Software Engineer — ML Platform, Sigma IQ – Seattle, WA Apr 2019 – June 2021

- Productionized a financial-reconciliation ML model (Python, XGBoost) and migrated the pipeline to Spark / Scala for $>100\times$ speedups over the prior Python implementation.
- Built tools for managing datasets, training models, and evaluating performance, enabling researchers to iterate rapidly on reconciliation accuracy.
- Built a full-stack platform for running Jupyter notebooks in production (TypeScript, React, Next.js, Docker, GCP).

Independent, Self-Directed ML Study – San Francisco, CA Jan 2018 – Apr 2019

- Took a year off to study classic ML and deep learning in depth (FastAI parts 1–2, Andrew Ng's Coursera specializations).
- Built `fastai_audio` (194★), a library for on-the-fly GPU spectrogram generation during training; write-up.

Senior Software Engineer, Ruvixx, Inc. – San Francisco, CA Apr 2015 – Dec 2017

- Built licensing and brand-protection platform (Ruby on Rails, Angular) for HDMI, Dolby, and Philips.
- Engineered data model, optimized SQL queries, and developed real-time dashboards in a multi-tenant architecture.

Software Engineer, Emjoyment – Mountain View, CA May 2014 – Apr 2015

- Built iOS app (Objective-C) and Python/Django backend services for an early-stage 'Tinder for Jobs' startup.

Software Development Engineer, Amazon.com – San Luis Obispo, CA July 2012 – May 2014

- Built internal services for the CreateSpace Royalties & Financials team (Java, Spring, AWS) using test-driven development.

Education

California Polytechnic State University, BS + MS in Computer Science – San Luis Obispo, CA Sept 2007 – June 2012

- Thesis: *Real-time Musical Analysis of Polyphonic Guitar Audio*

Skills

Languages: Rust, Python, TypeScript, Scala, Swift, Kotlin, C#, Java

ML & Audio: PyTorch, Candle, TensorFlow, FastAI, Weights & Biases, XGBoost, ONNX, torchaudio, librosa, DeepFilterNet, FFmpeg, Whisper, Deepgram, CREPE, Basic Pitch, NMF, FFT, CQT/HCQT

LLM Systems: Claude API, Anthropic Bedrock, MCP, Pydantic, LLM-as-judge, retry-with-feedback, agent orchestration

Systems & Web: Rust + WASM, Tauri, ratatui, Axum, FastAPI, Django, Next.js, React, Three.js

Infrastructure & Tools: AWS, Docker, GCP, Cloudflare, jj (Jujutsu), git, Claude Code