

Joel Villarino

903-348-9535 | JoelAVillarino@gmail.com | linkedin.com/in/joel-villarino | github.com/javillarino | Houston, TX

Education

Rice University

Bachelor of Science in Computer Science and Bachelor of Science in Statistics
Minor in Data Science

Houston, Texas

Aug 2022 – May 2026

GPA: 3.95/4.0

Experience

Nextdoor

Software Engineering Intern

- Shipped a production multimodal classification pipeline (OCR + Google Vision + LLM routing) processing 15k+ articles/day; improved classification accuracy by 30% (offline eval + A/B validation) and reduced manual deletion incidents.
- Built safe rollout + ops controls: sitevar-gated deployments with a killswitch, metrics for Vision/LLM failure modes, and end-to-end dashboards for filtering outcomes and processing completion.
- Eliminated a major SQL bottleneck by rewriting a high-cost query path; cut compute from **5 CPU-hours/day** (~5% of cluster load) to under **1 minute** and reduced end-to-end article latency.
- Deployed an embedding-based RSS deduplication and canonicalization service that self-heals 301 redirects and verifies URL shifts; improved feed reliability and removed recurring manual triage.
- Implemented real-time intent classification for Nextdoor's RAG chat; routed tasks via GraphQL + workers and surfaced contextual actions.

San Francisco, California

May 2025 – August 2025

Paycom

Software Development Intern

- Built an internal **iOS** prototype in **Swift** integrating ATS APIs; implemented job search and candidate detail flows.
- Developed reusable **SwiftUI/UIKit** components and resume parsing/validation for recruiter-facing views.
- Produced wireframes and interaction specs to align iOS/Android behavior and speed iteration with stakeholders.

Grapevine, Texas

May 2024 – Aug 2024

Rice University Teaching Assistant

COMP 412 – Compiler Construction, COMP 215 – Object-Oriented Programming

Houston, Texas

May 2024 – Present

- Held office hours, graded assignments, and guided students through register allocation and instruction scheduling.

RiceApps

Software Developer

Houston, Texas

June 2023 – Present

- Built a React-based impact calculator for United Way Houston that generates personalized reports; automated a previously manual reporting workflow.
- Led development of a React Native app for UT Health enabling clinicians to assign and monitor physical therapy workouts.
- Mentored 12 developers on React architecture and Git/PR workflow; improved review quality and standardized design patterns.

Projects

Autonomous Drone Herding System - *Algorithmic Robotics & Control*

- Engineered a **real-time multi-agent simulation pipeline** capable of processing thousands of geospatial updates per frame; vectorized core computations with **Numba JIT** compilation to achieve **50–130x speedup** over baseline.
- Designed a **hybrid neighbor-caching strategy** with movement thresholds, amortizing k-NN search cost and yielding 17–30% additional throughput for large-scale flock scenarios.
- Implemented **Server-Sent Events** architecture for real-time state synchronization, replacing HTTP polling and enabling continuous multi-drone coordination with sub-second latency.
- Led architectural separation of simulator and application layers; added 70+ tests enabling safe refactoring across backend codebase.

djprep - *Open-Source DJ Audio Analysis*

- Built a **Rust** CLI for BPM/key detection and stem separation; exports Rekordbox XML as an open-source alternative to Mixed In Key.
- Integrated **HTDemucs** via **ONNX Runtime** with automatic execution provider selection across CoreML, CUDA, and DirectML.
- Designed bounded concurrency isolating **rayon** CPU work from GPU inference to prevent memory exhaustion during batch processing.

MCP Agent Attestation Extension - *Security Research*

- Identified an authentication gap in MCP (no agent/model provenance); wrote threat analysis and spoofing scenarios.
- Designed a **JWT**-based attestation layer with **Ed25519** signatures, SPIFFE-compatible identifiers, and JWKS distribution.
- Built a Python proof-of-concept with replay protection and an attack harness (8 vectors); integrated with MCP SDK (140 tests).

Activities

Mariachi Luna Llena (Guitar) · **SHPE** (Socials Chair) · **HACER** (College Rep) · **Data Science Club** (Mentor) · **Edinburgh Abroad** (Salsa/Bachata)

Courses

Algorithmic Robotics, Compilers, Concurrent Programming, Computer Systems, Algorithms, Optimization, Machine Learning, Bayesian Statistics

Skills

Languages: Go, C++, Python, SQL, Rust, Swift, TypeScript, Java, C, R

Frameworks/Tools: Docker, Git, GraphQL, React, SwiftUI, scikit-learn, TensorFlow, Django