

Jaret Eduardo González Carrasco

High-Performance Computing & AI | Software Engineering Student

+52 871-903-1056 | contact@jareteduardo.dev | linkedin.com/in/jaret-eduardo | github.com/JaretEduardo
jareteduardo.dev

SUMMARY

Software Engineering student and Software Engineer focused on high-performance computing and AI systems. Hands-on experience with C, C++, Rust, CUDA, parallel computing, compilers, AutoML, Linux, and performance optimization, with measured project results including 4x-8x CPU speedups and 3x faster ML execution.

TECHNICAL SKILLS

Languages: C, C++, Rust, Python, Bash, Assembly (x86/ARM)

HPC & AI: CUDA, GPU acceleration, parallel computing, multi-threading, multiprocessing, joblib, AutoML, numerical verification

Performance & Systems: Linux, POSIX APIs, memory profiling, CPU cache optimization, GCC, LLVM, CMake, GNU Make, GDB, Valgrind, cross-compilation

SOFT SKILLS & LANGUAGES

Soft Skills: Analytical Problem Solving, Technical Communication, Team Collaboration, Adaptability, Ownership, Continuous Learning

Languages: Spanish, English

WORK EXPERIENCE

API INMOBILIARIA | Software Engineer | Dec 2025 - Present

- Optimized data pipelines and legacy code in a Linux environment using performance-oriented data structures and hardware-level optimization, improving system throughput by 20%.
- Refactored performance-critical code with an emphasis on efficient resource usage, maintainability, and system-level behavior.

SELECTED PROJECTS

ParallelPilot - AI-Assisted Automatic Parallelization | Python, multiprocessing, joblib, Codex, GPT-5.6

- Built a CLI that analyzes serial Python code, generates parallel implementations, verifies numerical correctness, and benchmarks real speedup across CPU core counts.
- Achieved 4x-8x speedups on tested simulation scripts without manual code modification.

KAIML - High-Performance Computing Library | C++20, Rust, CUDA

- Developed a multi-threaded, GPU-accelerated computing library that improved machine-learning execution speed by 3x.

KAI-CC - Custom Compiler Architecture | C++17, Valgrind

- Architected a compiler pipeline from lexical analysis and parsing to AST construction and native code emission; reduced execution time and memory footprint by 30% through IR profiling and optimization.

KAI - AutoML Platform Core | C++, concurrency, CPU optimization

- Architected a high-throughput AutoML core using concurrency and CPU-cache optimization to process 50 GB of data daily without memory bottlenecks or leaks.

EDUCATION

B.S. Software Engineering | Universidad Tecnológica de Torreón | Expected Dec 2027

TSU Software Development | Universidad Tecnológica de Torreón | Apr 2026