Diogo Silva

+351 918 978 337 | Braga, Portugal | diogo.coelho.silva@gmail.com | github.com/diogocsilva12 | linkedin.com/in/diogocsilva12 | diogocsilva12.dev

Profile

MSc student in Advanced Computing at the University of Minho specializing in **High-Performance Computing, parallel programming, distributed systems, and performance optimization**. Experienced with C/C++, Python, MPI, OpenMP, CUDA, SLURM, and Linux. Currently developing a Master's thesis on disaggregated memory for heterogeneous systems, with previous research in CPU-based LLM inference and industry experience at Bosch developing AI-assisted security tooling.

Experience

Software Developer Intern — Bosch Portugal Jul 2026 — Sep 2026

- Developed a TypeScript VS Code extension for real-time security vulnerability detection.
- Integrated Generative AI APIs for context-aware code remediation and automated security-fix suggestions.

EUMaster4HPC Summer School Jul 2026 — Jul 2026

- Explored emerging HPC and AI topics while collaborating with students, researchers, and HPC professionals across Europe.

Education

University of Minho Braga, Portugal
MSc in Advanced Computing *Sep 2025 — Present*

- **Focus:** High-Performance Computing, Parallel Computing, Distributed Systems and Performance Optimization

University of Minho Braga, Portugal
BSc in Computer Science *Sep 2021 — Jul 2025*

Selected Projects

Master's Thesis — Disaggregated Memory in Heterogeneous Systems 2026 — Present

- **Ongoing Master's thesis research.** Developing a memory-disaggregation framework for heterogeneous HPC and data-center environments.
- Investigating remote-memory access, communication overhead, data placement, and performance-aware memory management across distributed architectures.
- **Tech:** C++, Python, Bash, MPI, Distributed Systems, Memory Management

LLM Inference on CPU (github.com/diogocsilva12/LLM-Inference-CPU) Feb 2026 — Jun 2026

- **Grade: 20/20.** Benchmarked and optimized local LLM inference on CPU-based HPC architectures.
- Evaluated quantization and runtime configurations using llama.cpp, GGUF, Linux, and SLURM.
- **Tech:** C/C++, Python, Bash, llama.cpp, GGUF, SLURM

Real-Time 3D Rendering Engine (github.com/diogocsilva12/CG-Project) 2025 — 2025

- **Grade: 18/20.** Developed a real-time 3D graphics engine with rendering pipelines, shaders, lighting, and scene management.
- **Tech:** C++, OpenGL, GLSL

Technical Skills

- **Programming:** C, C++, Python, TypeScript, Java, SQL
- **HPC & Parallel:** MPI, OpenMP, CUDA, SYCL, SLURM, performance optimization
- **Systems & Tools:** Linux, Bash, Git, Docker, Apache Spark, Milvus, MySQL
- **Areas:** HPC, Distributed Systems, Scientific Computing, Heterogeneous Systems, AI Systems

Additional

Languages: Portuguese — Native · English — Professional · Spanish — Basic

Activities & Certifications: EUMaster4HPC Summer School · EPICURE Hackathon · ENEI · English B2 — Babelium