

Alejandro Amat

Electrical & Computer Science Engineer

alejandroamatp@gmail.com — alejandroamat.me — GitHub: AlejandroAmat — LinkedIn: alex-amat-paya-

Professional Summary

- 3+ years of experience in Computer Vision, 3D Graphics, Machine Learning, and differentiable rendering
- Specialized in translating cutting-edge computer vision, machine learning, and graphics research into real-time functioning systems
- Proficient in ML/Computer Vision using C++/CUDA and Python/PyTorch, combined with graphics programming experience in Vulkan/OpenGL
- Expert in bridging theoretical research and practical real-time applications, creating systems that make advanced graphics and ML techniques accessible for real-world deployment

Education

Carnegie Mellon University, Pittsburgh, PA 2024–Present
Master of Science in Computer Science

Carnegie Mellon University, Robotics Institute, Pittsburgh, PA 01/2024–07/2024
Bachelor's Thesis: Real-Time Differentiable Rendering with Slang

CFIS-UPC, Barcelona, Spain 2019–2024
Dual Bachelor of Engineering in Electrical Engineering & Computer Science
GPA: 8.43/10 (Rank: 6/237) & 8.1/10 (Top 5%)

Professional Experience

Research Associate 02/2024–Present
Carnegie Mellon University, Pittsburgh, PA

Conducting research at the intersection of machine learning, computer graphics, and high-performance computing, specializing in real-time systems and optimization for 3D computer vision applications.

- Contributed to static rendering component for accepted SIGGRAPH 2025 journal paper on 3D live streaming using Gaussian Splatting, achieving 3-4× speedup through load balancing kernels and CUDA graph pipelines with 15ms/iteration for 1.5M Gaussians
- Built high-performance computer vision systems with PyTorch prototyping and translation to real-time ML/graphics applications using advanced C++ optimizations including SIMD vectorization, OpenMP parallelization, Eigen library achieving 2-6ms/frame performance with cross-platform deployment using CMake and Python bindings
- Developed differentiable rendering pipeline using Slang and gfx API with custom GPU optimizers delivering 116× speedup over NVIDIA Falcor and 4× speedup in 2D Gaussian Splatting with profiling using Nsight Systems, Nsight Compute, and RenderDoc
- Train MLP architectures with PyTorch prototyping and TensorRT deployment, and develop custom differentiable optimizations with second-order nonlinear optimization algorithms across distributed systems
- Supervise visiting researchers and interns in 3D Gaussian Splatting and computer vision projects, conducting technical interviews and project coordination

AI Researcher 05/2023–12/2023
I2cat Foundation, Barcelona, Spain

- Pioneered real-time 3D imaging algorithms with MIMO FMCW mmWave radar for reflective surface optimization
- Published 2 papers on location-based utilization of reconfigurable intelligent surfaces for mmWave sensing

Unity VR Developer 10/2023–04/2024

IESE Business School, Barcelona, Spain

- Developed cross-platform VR application for IESE/MIT/LBS collaboration, deployed to Android, iOS, and Web/Desktop with Unity and Google Cardboard SDK
- Implemented Firebase real-time synchronization for multi-user collaborative sessions and progress tracking

Student Research Lead

02/2023–07/2023

HP, Barcelona, Spain

- Led university team on collaborative project with HP, developing Python package using Intel RealSense cameras
- Implemented DBSCAN and RANSAC algorithms for pointcloud segmentation and wall plane detection

Big Data Engineer Intern

05/2022–09/2022

Mango, Barcelona, Spain

- Built data pipelines using Jenkins CI/CD, Airflow, MySQL, AWS, Python & Docker for retail analytics
- Developed automated data processing workflows for large-scale retail data analysis

Publications

- Junkai Huang*, Saswat Subhajyoti Mallick*, **Alejandro Amat**, Marc Ruiz Olle, Albert Mosella-Montoro, Bernhard Kerbl, Francisco Vicente Carrasco, Fernando de la Torre. "Echoes of the Coliseum: Towards 3D Live Streaming of Sports Events." *ACM Transactions on Graphics (SIGGRAPH 2025)*.
- Filip Lemic, Jalal Jalali, Gerard Calvo Bartra, **Alejandro Amat**, Jakob Struye, Jeroen Famaey, Xavier Costa Perez. "Location-based real-time utilization of reconfigurable intelligent surfaces for mmWave integrated communication and sensing in full-immersive multiuser Metaverse scenarios." In *IET Advanced Metaverse Wireless Communication Systems*, pp. 101-136, 2025.

Selected Projects

Vulkan 3D Gaussian Splatting Renderer - Cross-platform C++17 neural rendering engine

- Built complete 3DGS implementation breaking CUDA vendor lock-in with Vulkan and GLSL compute shaders, featuring tile-based rasterization with multi-stage radix sort (histogram computation, work-efficient prefix sum using up-sweep/down-sweep phases across ping-pong buffers)
- Created ImGui keyframe animation system for real-time smooth sequence generation with interpolated camera movements, FOV control, culling, wireframe rendering, performance monitoring, and cross-platform deployment with optimized memory layouts
- Achieved stable 60 fps on Nvidia 3060 Ti and 24-45 fps on Apple Silicon (MoltenVk translation) on 2M Gaussian pointcloud with bilinear upsampling for Retina display support and cache-performant data structures

High-Performance CUDA Optimization Suite - GPU-accelerated algorithms: Sobel Edge Detection (224,733× speedup over NumPy, 10.8 billion pixels/sec), 3×3 SVD (120× speedup vs PyTorch), Fused MS-SSIM, and GPU-JPEG (125× speedup).

Differentiable Rendering with Slang - Real-time inverse rendering with 116× speedup over NVIDIA Falcor and 4× boost in 2D Gaussian Splatting using automatic differentiation and custom GPU optimizations.

Operating System Kernel - Custom x86 kernel with process scheduling, memory management (virtual memory, paging), system calls, and hardware interrupt handling.

Custom Rendering Engines - Vulkan/OpenGL implementations with advanced lighting, shadow mapping, and optimized shader pipelines.

Technical Skills

Programming Languages: C++, Python, CUDA, PyTorch, Unity/C#, Assembly

Graphics & Rendering: OpenGL/Vulkan, HLSL/GLSL, Differentiable Rendering, 3D Gaussian Splatting

Machine Learning: PyTorch, Neural Networks, TensorRT, MLP Architectures, Computer Vision

HPC & Optimization: SIMD, OpenMP, Eigen, Nsight Systems/Compute, RenderDoc

Systems: Git, Docker, CI/CD, Linux/Windows, CMake, Python Bindings