Jichu Mao

Graphics Performance Engineer

maojichu@gmail.com · jichu.art
linkedin.com/in/jichu-mao-a3a980226 · github.com/jichumao

Experience

TikTok, Graphics Performance Engineer — Jul 2026 – Present
Optimizing TikTok's real-time effect rendering engine (C++), used by hundreds of millions of creators daily: frame-rate and latency optimization across the GPU pipeline, shaders, memory and thread scheduling on Android, plus performance profiling and AI-assisted regression analysis across low-end to flagship devices.

TikTok, Software Engineer Intern — Jun 2025 – Aug 2025
Effect House team. Built a real-time performance analysis system for the Effects engine, instrumenting the C++ rendering core with Tracy, and an LLM-agent pipeline that reads renderer parameters and the scene graph to diagnose effects and generate semantic tags.

University of Pennsylvania, Teaching Assistant — Aug 2024 – May 2025
Held office hours and graded assignments for Interactive Computer Graphics (CIS 4600) and Advanced Rendering (CIS 4610).

Infold Games, Graphics Rendering Intern — Jun 2024 – Aug 2024
Worked on an Unreal Engine rendering pipeline matched to the project's art style: engine customization for lighting, rendering and resource management, and graphics performance and compatibility work across platforms.

Apple, Software Engineer Intern — Jan 2023 – May 2023
Software development for the Rel Engineering team in Shanghai.

University of Alberta, Software Research Intern — Jul 2022 – Oct 2022
Mixed reality development: improved DICOM model conversion and volume rendering on a dataset from more than 50 patients, with about 35% faster image loading. Funded by the China Scholarship Council.

Selected projects

World-Space ReSTIR in Vulkan, team of two — 2024
World-space global illumination in a Vulkan ray-tracing pipeline. About 40% more accurate than screen-space ReSTIR in large scenes, for roughly 5% more render time.

CUDA Path Tracer, individual project — 2024
A GPU path tracer in CUDA with glTF scenes, physically based materials, depth of field and environment lighting. Stream compaction cuts frame time from 25.7 to 18.2 ms in open scenes.

WebGPU Forward+ and Clustered Deferred Shading, individual project — 2024
Forward+ and clustered deferred shading in the browser. With 5,000 moving lights at 4K, clustered deferred runs above 100 fps where naive forward shading manages 3.

Vulkan Grass Rendering, individual project — 2024
Real-time grass in Vulkan. A compute shader simulates and culls every blade; with 32,768 blades, culling lifts the frame rate from 70 to 228 fps.

CUDA Flocking, individual project — 2024
Reynolds' boids simulated on the GPU with CUDA. A coherent uniform grid runs 100,000 boids at 2,677 fps where brute force manages 43.

Education

University of Pennsylvania — Aug 2023 – May 2026
Master of Science in Engineering, Computer Science & Computer Graphics

Beijing Jiaotong University — Sep 2019 – Jun 2023
Bachelor's degree, Computer Software Engineering

Skills

Languages: C++, CUDA, GLSL / HLSL, TypeScript, JavaScript, Python

Graphics APIs: OpenGL, Vulkan, WebGL

Graphics debugging: Nvidia Nsight Graphics, RenderDoc

Content tools: Maya, Houdini, Qt

Other: Perforce, Git, Unity, Unreal Engine