

# Shaoxuan Yin

## GRAPHICS, RENDERING & AI SYSTEMS ENGINEER

djmax96945147@gmail.com | [qervas.github.io](https://qervas.github.io) | [github.com/Qervas](https://github.com/Qervas) | [www.linkedin.com/in/sxyin/](https://www.linkedin.com/in/sxyin/)

Engineer spanning real-time graphics, neural-rendering research, and AI platform work. Built a standalone Vulkan engine with hybrid deferred and hardware ray-traced paths, a 12-camera capture system and a 209GB public dataset for 3D Gaussian Splatting research, and a serverless GPU platform serving seven generative models behind one API and an MCP server. Comfortable from GPU and shader code up through backend orchestration and developer-facing tooling.

### TECHNICAL SKILLS

**Languages:** Modern C++ (C++23 / C++20), CUDA, Python, TypeScript, Rust, C#, GLSL, HLSL

**Graphics / Rendering:** Vulkan, OpenGL 4.3, DirectX 12 (learning), deferred + hardware ray tracing, path tracing, PBR, global illumination, OIDN / OptiX / NRD denoising

**AI / Gen-media:** Serverless GPU inference, image / 3D-mesh / texture / motion / music models, MCP tool-calling, agentic workflows, SDK design

**Research / 3D:** NeRF, 3D Gaussian Splatting, neural rendering, Nerfstudio, differentiable rendering (Taichi, NVIDIA Warp, nvdiffrast)

**Systems / Tooling:** CMake, Git, Docker, Linux, async job orchestration, profiling with Nsight, RenderDoc, PIX

### EXPERIENCE

#### AI Startup

2025 - Present

AI Platform Engineer | Remote

- Building NIUA, an agentic game-dev tool: one API over seven self-hosted open-source models (image, 3D mesh, PBR texture, music, motion) with agent-driven, multi-step asset workflows.
- Built a serverless GPU pipeline with independently deployable services and async job orchestration; cut model cold-start from ~10 min to ~20 s.
- Shipped developer-facing surfaces: an SDK, an MCP server for tool-calling agents, and Unreal/Blender plugins.

#### Freedo Technology

Feb 2026 - Present

AI Algorithm Intern | Remote

- Working on fluid simulation and differentiable rendering systems.
- Using Taichi, NVIDIA Warp, and nvdiffrast in graphics- and simulation-related workflows.

#### Beijing Guoyao Xintiandi Information Technology

Aug 2022 - Sep 2022

C++ Development Engineer Intern | Beijing, China

- Built a 3D visualization tool in Qt + OpenGL, improving rendering performance by 40%; also wrote a custom C++ network protocol for the data layer.

### PROJECTS

#### Ohao Engine

Nov 2024 - Present

C++, Vulkan, GLSL -- ~50K C++ + ~15K GLSL across 121 shaders, no editor dependency

- Standalone Vulkan rendering engine. Hybrid deferred + KHR ray-tracing paths share one scene representation, bindless materials/textures, and a TLAS/BLAS lifecycle (rebuild on geometry change, refit on transform-only updates).

- Path tracer: multiple importance sampling (balance + power heuristic), environment-map importance sampling via precomputed marginal/conditional CDFs, Owen-scrambled Sobol QMC, variance-driven adaptive sampling, and three denoiser backends (Intel OIDN, NVIDIA OptiX, NRD).
- Deferred path: Cook-Torrance GGX, cascaded shadow maps, SSAO, screen-space reflections, subsurface scattering, TAA, bloom. Composited with ray-traced 1-bounce indirect (temporal blending) and ray-traced shadows on the same scene buffers.

### **Master's Thesis: Neural Rendering Dataset Collection**

Mar 2025 - Dec 2025

Linköping University -- C++, C#, WinUI, Sapera SDK, GenICam

- Built CamMatrixCapture, a 12-camera synchronized capture system with staggered triggering, Bluetooth turntable automation, and state-machine based workflow control.
- Created and released a 209GB public dataset for NeRF and 3D Gaussian Splatting with challenging reflective, transparent, and translucent materials.
- Benchmarked NeRF vs. 3DGS and found an 11dB PSNR advantage for 3DGS on reflective surfaces.

### **NocturneAI (open source)**

Mar 2025 - Oct 2025

TypeScript -- ~30K LOC multi-agent CLI framework, [github.com/Qervas/NocturneAI](https://github.com/Qervas/NocturneAI)

- Multi-agent CLI framework using GitHub Copilot's API as LLM backend. ReAct loops, tool registration/execution, workflow engine.
- Clean architecture (core / application / infrastructure / presentation), AgentCoordinator for multi-agent orchestration, three modes (Agent / Ask / Edit) via ModeManager.
- Iteration loop with per-step task analysis, workflow validator for multi-step plans, plugin manager, metrics collector, background-task manager.

### **Path Tracer**

2024

C++, CUDA

- Built a Monte Carlo path tracer with BVH acceleration, Russian roulette termination, and multiple material models.
- Reached 40 FPS on RTX 4060 at 1280x720 through CUDA acceleration.

## **EDUCATION**

### **Linköping University**

2023 - 2026

MSc, Computer Science -- Graphics and Visualization | Linköping, Sweden

- Graduated Feb 2026. Top grades (5/5) in Advanced Global Illumination, Modelling & Animation, Computer Games Design.

### **Beijing Information Science and Technology University**

2018 - 2023

BEng, Computer Science | Beijing, China

## **ADDITIONAL**

**Publication:** "Skiing Action Analysis Based on OpenPose," Intelligent Computer and Applications, 2022

**Languages:** English (fluent), Chinese (native), French (basic)