

Shingo Takagi

www.linkedin.com/in/shingotakagi | scholar.google.com/citations?user=HxPqBrMAAAAJ
www.youtube.com/@shingo.takagi | aut.omat.ing | faviclone.com

SUMMARY

Founding and Senior Software Engineer specializing in **(1) photorealistic, real-time video AI agents**, and **(2) high-fidelity synthetic data generation** for ML model training. Expert in architecting ultra-low latency conversational clones utilizing linear rectified flow diffusion and robust multi-modal stacks (STT, TTS, LLM, VLM). Brings a unique technical perspective bridging physical sensor data with digital representations via **3D Gaussian Splatting (3DGS) and NeRFs**, built upon foundational experience developing industry-standard VFX tools (Katana) and publishing **cutting-edge generative AI research** (ECCV, WACV, FG).

EXPERIENCE

Automating | Mountain View, CA

Founding Software Engineer | March 2025 – Present

- **Photorealistic Video Agents:** Architected and deployed photorealistic video AI agents capable of seamless, multi-modal conversation via both text and real-time video interfaces.
- **Ultra-Low Latency Pipelines:** Engineered real-time interaction pipelines, utilizing linear rectified flow diffusion for ultra-low latency graphics generation and designing highly optimized streaming architectures for real-time voice chat.
- **Cross-Platform Deployment:** Spearheaded the end-to-end development and commercial launch of cross-platform AI applications (Android, iOS, Web) tailored for the education and automated retail (QSR) sectors.
- **Agentic Visualization Tools:** Developed custom agentic visualization tools that dynamically generate real-time diagrams and imagery, significantly enhancing the AI's ability to provide rich visual context during user interactions.
- **Multi-Modal Model Integration:** Integrated and fine-tuned a complex, multi-modal stack of Speech-to-Speech, STT, TTS, LLM, and VLM models to power robust conversational flows.

Meta | Menlo Park, CA

Senior Software Engineer | August 2017 – January 2025

- **Bridging the Sim-to-Real Gap:** Architected end-to-end PBR synthetic data pipelines for human-centric ML models, reducing training error by matching the statistical distribution of real-world sensor noise and lighting.
- **Sensor & Lens Fidelity:** Engineered a proprietary library for extracting ray-accurate sensor models from **CodeV/Zemax**, enabling the simulation of cameras through **pancake lens systems**.
- **Neural Ground Truth Generation:** Developed hybrid data workflows using **3DGS** and **NeRFs** to reconstruct high-fidelity digital twins, providing pixel-perfect ground truth for depth, occlusion, and semantic segmentation.
- **Scalable Rendering Architecture:** Optimized large-scale workflows (OpenCue, custom renderfarms) to accelerate synthetic dataset throughput, ensuring consistent multimodal metadata across RGB, Depth, and IR sensors.
- **Research Collaboration:** Partnered with research teams to support peer-reviewed submissions in human motion synthesis and generative 3D modeling, with accepted papers at **ECCV, WACV, and FG**.
- **Spatial Computing & Real-time Rendering:** Contributed to asset/shader loading and animation

components (OpenXR, Vulkan) for the Meta Spatial SDK, and implemented experimental algorithm prototypes for high-fidelity VR headset passthrough.

DreamWorks Animation | Glendale, CA

Senior Software Engineer | December 2011 – March 2016

- **Led development of Torch**, a renderer-agnostic node-graph lighting tool used across the studio's feature film slate.
- **Managed full-lifecycle feature delivery** as Scrum Lead, coordinating cross-functional scheduling and UI/UX enhancements for the lighting department.
- **Refactored core node-graph logic** to improve UI responsiveness and usability, significantly reducing artist iteration time on complex lighting setups.

Walt Disney Animation Studios | Burbank, CA

Senior Software Engineer | October 2005 – October 2010

- **Developed real-time GPU-based lighting tools**, utilizing deep pixel buffer caches to significantly reduce render-loop latency for real-time visual effects.
- **Optimized production performance** by migrating legacy Renderman shaders and C++ shade ops to modern GLSL, achieving substantial gains in interactive preview speeds.

Sony Pictures Imageworks | Culver City, CA

Software Engineer | August 2003 – September 2005

- **Collaborated as a core member of the team** to design and develop **Katana**, the industry-standard 3D Lighting & Look-Dev tool.
- **Engineered procedural scene resolution logic**, enabling artists to override complex 3D scene graphs at scale—a workflow now used globally in VFX.
- **Architected foundational APIs** for image and scene-graph processing, optimizing the pipeline for high-performance 3D asset evaluation.

TECHNICAL SKILLS

- **AI/ML:** 3DGS, NeRFs, Differentiable Rendering, Diffusion, ControlNet, LLM, VLM, LORA.
- **Graphics & Optics:** PBR, Ray Tracing, Camera/Sensor Calibration, CodeV, Zemax, GLSL, Vulkan, CUDA.
- **Software:** C++, Python, USD (Universal Scene Description), Maya, Blender, Arnold, Cycles, Unity, Katana.

EDUCATION

- **Master of Science, Computer Science** | University of British Columbia
 - Thesis: Photometric Optical Flow using RGB channels.
- **Bachelor of Science, Computer Science** | University of Toronto
 - Recipient of the Canada Scholarship.

SELECTED PUBLICATIONS

- [OASIS: Object-guided Attention for Text-conditional Diffusion Synthesis of Human Interaction Sequences](#), FG 2025
- [Generalizable Human Gaussians for Sparse View Synthesis](#), ECCV 2024
- [Towards Realistic Generative 3D Face Models](#), WACV 2024
- [MotionGPT: Human Motion Synthesis with Improved Diversity and Realism via GPT-3 Prompting](#), WACV 2024