

Joshua Urban Davis

www.joshurbandavis.com | San Francisco, CA 94110 | joshurbandavis@gmail.com

Computer Vision/Machine Learning Engineer & CS PhD (Dartmouth) specializing in the intersection of **Real-Time Perception and Human-Centric AI**. 6 years of experience building real-time end-to-end CV/ML systems at small startups and big tech (Meta Reality Labs) with a track record of translating high-level research into production-ready systems (3 Patents, Meta Reality Labs, Adobe Research). Pre-PhD, worked at several startups including Bregmann Media Labs, Hex Plexus between 2013 - 2017.

Languages: Python, C++, Swift, JavaScript, C#, MATLAB

Frameworks & Libraries: PyTorch, NumPy, OpenCV, MediaPipe, OpenGL, TensorFlow, scikit-learn, Pandas, ROS2, Reality Kit

Tools: Git, Docker, Unity, Jira, CI/CD pipelines

Specialties: 3D reconstruction, SLAM, lighting estimation, semantic scene understanding, on-device inference, generative AI, AR/VR

OVERLAY ROBOTS (pre-seed, stealth startup)

Founding Computer Vision and Machine Learning Engineer (Lead)

Oakland, CA

Sept 2025 - Present

- Architected and built the **end-to-end perception stack** as the **founding CV/ML engineer** for the world's first intelligent makeup robot
- Drove sensor selection, sensor fusion, and camera calibration for real-time high precision **3D reconstruction, pose estimation, eye tracking, alignment**, and skin surface mapping deployed on **embedded hardware** (Nvidia [Jetson/Orin]) under strict latency and safety constraints.
- Built **data collection, labeling, and evaluation infrastructure** from scratch: automated capture pipelines for face scans
- Defined the **perception roadmap and prioritized technical investments** across 3D reconstruction, tracking, color perception, and safety monitoring, making architectural decisions that balanced prototype speed with production scalability as the system **moved from lab demos to manufacturing-ready hardware**.

META | REALITY LAB

Research Scientist and Software Engineer

Menlo Park, CA

March 2024 - Sept 2025

- Designed real-time **CV/ML lighting estimation** pipeline achieving 6-9ms inference **on device**, enabling photorealistic virtual object compositing at 72/90 FPS across 25+ indoor/outdoor scene categories in MR passthrough mode.
- Research, design, and develop applied color perception algorithms for camera and AR applications including novel **metric development** and **key performance indicator evaluation** pipelines
- Developed camera calibration and algorithms for **3D surface reconstruction and semantic scene understanding** for multi-sensor perception pipeline, fusing dToF and wide-angle RGB sensors to enable depth completion, sensor calibration, and cross-modal feature alignment
- Implement and support **large video, audio, and depth dataset generation** (~15,000 samples) and processing for algorithm development (AI/ML) in new technology explorations.
- Integrate **SLAM, pix2pix, and SIFT** algorithms for novel lighting estimation and developed evaluation pipelines for nightly benchmark runs in **C++, Python, OpenCV, and internal libraries**.
- Managed **full-lifecycle deployment**, from model training in PyTorch to high-performance C++ implementation on embedded targets.



ADOBE RESEARCH

Research Assistant

San Jose, CA

2021 - 2024 Extended Collaboration

2021 Summer Internship

- Led the development of media augmented video conferencing **creativity tools** that leveraged **body pose recognition, gesture and speech detection** to deliver immersive experiences using **Python, Javascript, MediaPipe, and OpenCV**.
- Designed and executed comprehensive **mixed-method qualitative and quantitative studies** involving 46 users and translated research findings into actionable recommendations.
- Optimized models for **on-device object detection and 3D human understanding** running at 60/72 FPS on consumer laptops combining MediaPipe pose estimation, custom gesture classifier (94% accuracy on 18-class vocabulary), and speech-triggered augmentation with <150ms end-to-end latency.
- Moved work to publication and incorporated into thesis research.



MICROSOFT RESEARCH

Software Engineering and Research Intern

Redmond, WA

Summer 2020



- Developed and implemented an interactive AR mask to display real-time video of the user's mouth on the mask's surface.
- Developed perspective-corrected UV mapping algorithm compensating for mask curvature and viewer angle, achieving <3 px distortion error at 72/90 FPS; optimized inference pipeline memory footprint by 40% for sustained on-device operation.
- Led perception software integration with hardware and accessibility research teams to transition from research prototype to commercial-ready product, addressing thermal constraints, power budget, and sustained inference stability for continuous operation.
- Led a mixed-methods user study with **18** participants utilizing a collaborative charades-style protocol, demonstrating a **40%** increase in communication accuracy and a **35%** reduction in task completion time compared to traditional opaque masks.

AUTODESK RESEARCH

User Interface and Graphics Research Intern

Toronto, ON

Winter 2019



- Spearheaded collaboration research across teams to pioneer cutting-edge techniques for **authoring 3D objects using generative AI** in VR. Leveraged **GANs, Transformers, and Autoencoders** to develop novel techniques for **human-AI collaboration**
- Prototyped and deployed interactive generative adversarial networks (GANs) specifically tailored for 3D design tasks in virtual reality. Employed tech stack including **pyTorch, Pandas, Numpy, and Unity** to create user-friendly generative AI authoring experiences.
- Securing a patent for AI virtual reality interaction techniques and system design and published technical paper in peer-reviewed proceedings. (doi.org/10.1145/3450741.3465260)

EDUCATION

DARTMOUTH COLLEGE

Doctorate of Computer Sciences

Hanover, NH

2024

UNIVERSITY OF TEXAS

Bachelors of Computer Sciences

Austin, TX

2013

PATENTS

- (Calliope) A System for Supporting Human-AI Collaboration in Virtual Environments. Josh Urban Davis, Fraser Anderson, George Fitzmaurice.
- (TangibleCircuits) A System for Converting Circuit Diagrams to Tangible and Audio Interfaces. Josh Urban Davis, Xing-Dong Yang
- (Circuit Style) A System for Peripherally Reinforcing Best Practices in Hardware Computing. Josh Urban Davis, Jun Gong, Xing-Dong Yang