

Shanmukh Upadhyayula

312-805-4303 | shanmukh.upad@gmail.com | [linkedin.com/in/shanmukh-upadhyayula](https://www.linkedin.com/in/shanmukh-upadhyayula) | github.com/ShanmukhUpad | shanmukhupad.github.io/portfolio

EDUCATION

University of Illinois Urbana-Champaign

Urbana, IL

Bachelor of Science in Computer Science & Geographic Information Science

May 2027 (Expected)

- GPA: 3.95/4.0 | Involvement: ACM Project Lead in Special Interest Groups for HCI and AI
- Relevant Coursework: Applied Linear Algebra, Virtual Reality, Computer Systems, Data Structures, Algorithms & Models of Computation, Human-LLM Interaction, Spatial Analysis
- Upcoming: Computational Photography, Computer Vision, Applied Machine Learning
- Planned: Master of Computer Science (MCS), August 2027 – May 2028

EXPERIENCE

UIUC CyberGIS Lab

Urbana, IL

Machine Learning Engineering Intern

September 2026 – Present

- Benchmarking the lab's geospatial AI agent on all 50 GeoAnalystBench tasks, where it matches published LLMs on code but plans more steps than the expert on 43 of 48 tasks
- Running user tests on the live agent across dataset search, map projections, raster data, file uploads, and adversarial prompts, and reporting 4 bugs that break core features
- Improving the agent's architecture, including how it picks tools, keeps its search index current, and runs code

Semantic Segmentation Failure Analysis

Urbana, IL

Undergraduate Researcher | shanmukhupad.github.io/portfolio/seg.pdf

May 2026 – August 2026

- Trained a DeepLabV3-ResNet50 model on aerial imagery and built a tiled scanner that maps where it is likely wrong across full-size images, using only the model's own confidence and no labeled data
- Validated the error maps on 4 benchmark datasets, found the model is about half as accurate along object edges, and presented the work in an IEEE-format paper and a CyberGIS talk

Digital Diary (ACM SIG-CHI) | github.com/sigchi-uiuc/digital-diary

Urbana, IL

Project Lead

September 2025 – April 2026

- Built a real-time camera pipeline with OpenCV and MediaPipe that tracks face and hand landmarks from webcam video to detect mood and gesture controls, running as its own Docker service
- Led a 15-person team building a Next.js journaling app with a guided mode that pulls in Google Calendar, weather data, and Claude AI

UIC Electronic Visualization Laboratory | github.com/ShanmukhUpad/curio

Chicago, IL

Software Engineering Intern

June 2025 – August 2025

- Built a Pandas merge tool that cut processing time by about 30% across hundreds of Chicago urban data pipelines
- Added dataset merging and multi-tab data views to Curio, a visual data-pipeline tool used by researchers

PROJECTS

Multi-Agent LLM Sycophancy | shanmukhupad.github.io/portfolio/mas.pdf

March 2026 – May 2026

- Showed that one injected false claim causes a 100% cascade through hierarchical LLM agentic workflows by turn 3
- Measured the failure without an LLM judge, using stance flips, planning errors, and deference language
- Found that anti-sycophancy prompting does not stop the collapse, and co-authored an ACL-format paper on it

Off the Clock: Standalone VR Game | C#, Unity XR | Virtual Reality Course Project

April 2026 – May 2026

- Built motion controls and a combat system for sword fighting, laser guns, and mech flight, using controller tracking data and sphere-cast sweeps so fast sword swings register hits reliably
- Tested Android builds on Meta Quest 3 against a 90 FPS target, adding an in-headset overlay for frame time, draw calls, and memory use to catch performance drops on the device
- Debugged laser effects that rendered on PC but disappeared in the headset, traced the cause to shaders that do not support single-pass stereo rendering, and replaced them with a compatible fallback

TECHNICAL SKILLS

Languages: Python, C++, C, C#, Java, TypeScript, SQL, R

Computer Vision & Image Processing: OpenCV, PyTorch, torchvision, MediaPipe, NumPy, ModernGL/GLSL

Tools: Git, Docker, Android, Unity XR, React, Next.js, XGBoost, GeoPandas, ArcGIS Pro