Omkar Vengurlekar

Tempe, AZ, USA | ovengurl@asu.edu | +1 (623) 283-0227
Google Scholar | omkarv23.github.io | github.com/OmkarV23 | linkedin.com/in/omkar-veng

SUMMARY

Ph.D. student in computational imaging (ASU, expected Fall 2028): physics-based differentiable rendering and neural scene representations (implicit neural fields, Gaussian splatting) for unconventional sensing modalities such as time-of-flight, sonar, and Schlieren imaging. Owns the full pipeline, from custom CUDA/OptiX renderers and adjoint solvers to calibrated multi-camera capture hardware. Six publications and two U.S. provisional patent applications, including co-first-author ACM TOG / SIGGRAPH Asia 2026, first-author 3DV 2026 oral, and second-author IEEE TPAMI. Seeking a Summer 2027 research internship in computational imaging or 3D vision.

EDUCATION

Arizona State University, Tempe, AZ, USA — May 2024 – Present
Ph.D., Electrical and Computer Engineering (expected Fall 2028), Imaging Lyceum Lab; advisor: Dr. Suren Jayasuriya

Arizona State University, Tempe, AZ, USA — Aug 2022 – May 2024
M.S., Robotics and Artificial Intelligence — GPA: 3.91/4.00

University of Pune, Pune, India — Aug 2016 – May 2020
B.Tech., Electronics and Telecommunication — GPA: 3.8/4.00

EXPERIENCE

Arizona State University, Imaging Lyceum Lab — Tempe, AZ, USA
Graduate Research Assistant — Jan 2024 – Present

- Co-first author (equal contribution) of **SDR-ToF** (ACM TOG / SIGGRAPH Asia 2026), an optical time-of-flight sensor using a commodity software-defined radio for RF-coherent quadrature demodulation: 10 million depth samples/s at 20–123 μm RMS depth noise (100 Hz–1 MHz bandwidth) on a ≈$10k prototype, versus 1–75 Hz per-point rates for prior coherent ToF systems.
- Led SDR-ToF signal-processing and experimental validation: recovered audio from direct and indirect surface vibrations using STFT spectral gating; measured tool vibrations and velocities at up to 300 Hz; and characterized scanned depth resolution using ISO 12233 slanted-edge MTF (0.71 mm slow-axis, 1.68 mm fast-axis).
- Led **SH-SAS** end to end (first author, 3DV 2026 oral): developed an implicit neural representation of complex acoustic scattering using spherical-harmonic coefficients, a hash-encoded MLP, and a custom CUDA autograd kernel for SH projection; trained directly from raw 1D time-of-flight measurements by analysis-by-synthesis without beamforming.
- Achieved 55% lower Chamfer distance than the prior neural baseline, improved IoU/F1, remained robust to 20% view subsampling, and trained ≈3× faster; validated on simulated, in-air, and underwater scenes.
- Co-developed **Z-Splat** (second author, IEEE TPAMI), camera–sonar fusion via differentiable Gaussian splatting; wrote the CUDA/PyTorch extension for differentiable sonar Gaussian projection. Addressed the missing-cone problem in 3D reconstruction, improving novel-view PSNR by +5 dB and reducing Chamfer distance by 60% across simulation, emulation, and hardware experiments.
- Developing a coupled adjoint for the refractive radiative transfer equation (RRTE) to jointly recover refractive index and scattering along curved rays; built a differentiable RRTE renderer in Mitsuba 3 / Dr.Jit with a symplectic eikonal integrator and cubic B-spline refractive fields, with gradients validated against finite differences.

Chiba University — Chiba, Japan
Visiting Ph.D. Student (ASU-sponsored research visit) — Jul 2026 – Aug 2026

- Developed a differentiable Background-Oriented Schlieren (BOS) tomography system for reconstructing time-varying 3D refractive-index fields using custom CUDA adjoint eikonal ray tracing and cubic B-spline field representations.
- Built a synthetic-to-real validation pipeline with GPU Boussinesq flow simulation (NVIDIA Warp), calibrated multi-camera BOS rendering, and measurement-only multi-resolution reconstruction. On synthetic plumes, achieved 39.1 dB PSNR and 0.28 relative L2 error versus 36.1 dB / 0.49 for the baseline, with 3× lower held-out flow error.
- Validated the system experimentally on a calibrated six-camera rig imaging time-varying density-driven flows in air and water.

Dartmouth College — Hanover, NH, USA
Visiting Ph.D. Student (host: Prof. Adithya Pediredla) — May 2024 – Jul 2024

- Prototyped acoustic non-line-of-sight (NLOS) imaging hardware and acousto-optic capture of occluded surfaces.

RAIN AI Remote, USA
AI Hardware and Systems Intern Aug 2023 – Dec 2023

- Implemented custom CUDA kernels for pre-trained ML model workloads and benchmarked GPU performance against the company's custom AI accelerator hardware.

Automaton AI Pune, India
Computer Vision Engineer Mar 2021 – Jul 2022

- Designed an end-to-end automatic data annotation pipeline for object detection, segmentation, and keypoint detection.
- Deployed efficient model inference on NVIDIA edge devices using post-training quantization (PTQ).

PUBLICATIONS

- **SDR-ToF: Million Optical Depth Samples per Second Using Software-Defined Radio** *paper*, *website*
  ACM Transactions on Graphics (SIGGRAPH Asia 2026, journal track), accepted
  Ramchander Rao Bhaskara*, Dhawal Sirikonda*, **Omkar Vengurlekar***, Juhyeon Kim, Joseph Lazarro, Suren Jayasuriya, Adithya Pediredla (*= equal contribution)
- **SH-SAS: An Implicit Neural Representation for Complex Spherical-Harmonic Scattering Fields for 3D Synthetic Aperture Sonar** *paper*, *code*, *website*
  International Conference on 3D Vision (3DV), 2026, oral presentation
  Omkar Vengurlekar, Adithya Pediredla, Suren Jayasuriya
- **Z-Splat: Z-Axis Gaussian Splatting for Camera-Sonar Fusion** *paper*, *code*, *website*
  IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), ICCP 2024 special issue
  Ziyuan Qu, **Omkar Vengurlekar**, Mohamad Qadri, Kevin Zhang, Michael Kaess, Christopher Metzler, Suren Jayasuriya, Adithya Pediredla
- **Reconstructing Microwave Synthetic Aperture Images Using Neural Fields** *paper*
  Antenna Measurement Techniques Association Symposium (AMTA), 2025
  Cecilio Obeso, Kavian Zirak, **Omkar Vengurlekar**, Suren Jayasuriya, Mohammadreza F. Imani
- **ToFGS: Temporally-Resolved Inverse Rendering with Gaussian Splatting for Time-of-Flight** *paper*
  CVPR 2025, 2nd Workshop on Neural Fields Beyond Conventional Cameras
  Omkar Vengurlekar, Aaron Saju Augustine, Suren Jayasuriya
- **A Hybrid ViT-CNN Architecture for Volumetric Target Classification with Sub-Bottom SAS Data** *paper*
  IEEE OCEANS, 2024
  Gregory D. Vetaw, **Omkar Vengurlekar**, Daniel C. Brown, Suren Jayasuriya

PATENTS

- **MHz-Rate Optical Depth Imaging Using Software-Defined Radios**. U.S. Provisional Patent Application No. 64/113,861, filed Jul. 17, 2026.
- **Method for 3D Synthetic Aperture Sonar Reconstruction using Complex Spherical-Harmonic Scattering Fields**. U.S. Provisional Patent Application No. 64/012,111, filed Mar. 20, 2026.

PROJECTS

Transient Time-of-Flight Renderer for 3D Time-Domain Back-Projection *code*

- Built a physics-accurate CUDA/OptiX single-bounce transient ToF simulator (Principled BSDF, 14,360-bin transient histograms) for time-domain back-projection; generated the simulated data for SH-SAS (3DV 2026).
- Sustained ≈1,500 sensor positions/s at 1M rays per position: a 54k-position cylindrical synthetic aperture (≈54B rays) rendered in 35 s.

FastBeamformer *code*

- Wrote a custom CUDA beamforming kernel for time-of-flight image reconstruction: ≈175× faster than NumPy and ≈8× faster than PyTorch CUDA tensors (111 s / 4.95 s / 0.63 s).

SKILLS

- **Research:** Computational imaging; differentiable and inverse rendering; neural rendering; 3D reconstruction; unconventional cameras and sensing; time-of-flight, sonar, and tomography
- **Programming:** C++, CUDA, Python, MATLAB
- **ML / graphics:** PyTorch, JAX, Mitsuba 3 / Dr.Jit, OptiX, NVIDIA Warp, OpenGL, OpenCV, TensorRT
- **Systems / hardware:** Linux, Git, Slurm/HPC, GPU kernel optimization, calibrated multi-camera capture

AWARDS

Ira A. Fulton Fellowship, ASU (competitive award to top incoming Ph.D. students) Aug 2024 – Aug 2028