Manoj Srinivasan

New York, NY | (201) 526-5591 | ms14845@nyu.edu | manoj-152.github.io | linkedin.com/in/manoj-s-nyu

EDUCATION

New York University (NYU), Tandon School of Engineering — Sep 2023 - May 2025
M.S. Computer Engineering, Cumulative GPA: 3.93/4.0 — New York, NY, USA

Indian Institute of Technology Madras (IITM) — Aug 2018 - Jul 2023
B.Tech - Mechanical Engineering, M.Tech - Robotics, Cumulative GPA: 3.51/4.0 — Chennai, TN, India
Minors in Computing, Artificial Intelligence & Machine Learning

TECHNICAL SKILLS

Programming Languages: Python, C/C++, SQL, Java, Bash/Shell scripting
Frameworks: PyTorch, Keras, Tensorflow, OpenCV, Scikit-Learn, HuggingFace, Numpy, SciPy, Pandas, W&B
Tools: Git/Github, Docker, Unix/Linux, CLI, AWS, HPC, Slurm, Singularity, Azure, Spark, CUDA
Domains: Machine Learning, Deep Learning, Computer Vision, NLP, Generative AI, LLMs, Data Analytics

PROFESSIONAL EXPERIENCE

Graduate Research Assistant - NYU Video Lab (New York City, New York, US) — Aug 2025 – Present

- Implemented a pipeline integrating 2D **video diffusion** priors with **3D Gaussian Splatting** (3DGS) using **PyTorch**, **HPC** clusters, and **Slurm**, achieving a **19% LPIPS improvement** (0.476 → 0.384) in novel-view synthesis under sparse input conditions.
- Enhanced 3D view consistency by formulating view generation as a temporal continuity task, integrating **camera-pose embeddings** with **diffusion-guided latent features** across viewpoints.
- Exploring **mesh registration** for human poses using learned skinning methods such as **SMPL**, to improve 3D consistency.

Graduate Engineering Intern - Intel Corporation (Santa Clara, California, US) — Jun 2024 – Aug 2024

- Developed a **computer vision** framework for automated detection of IC package design violations, emulating manually-performed inspection heuristics for 8 key checks.
- Optimized computational efficiency using **HuggingFace**'s **segmentation models** and **OpenCV**'s **morphological algorithms**, reducing detection pipeline runtime by **85%** (from over 4 hours to under 30 minutes).
- Honored with Intel's **Impact Award** for demonstrating strong productivity and delivering high-quality solutions in a short timeframe.

Machine Learning Intern - Preimage (Bangalore, Karnataka, India) — Sep 2022 – Dec 2022

- Adapted a **transformer-based multi-view stereo (MVS) pipeline** for dense 3D reconstruction from aerial drone imagery, improving accuracy by ~10% in sparse-view and large-scale outdoor scenes.
- Deployed large-scale reconstruction experiments on **Azure VMs**, leveraging **AWS S3** for dataset management and **PyTorch Lightning** with **CUDA** for efficient distributed training.

Image Processing Intern - GalaxEye Space (Chennai, Tamil Nadu, India) — Dec 2021 – Jan 2022

- Built a **super-resolution** neural network to upsample low-quality remote-sensing data in the form of SAR images, along with a **generative model** to predict RGB optical images from the super-resolved SAR output.
- Conducted experiments on cross-public datasets, leading to increased super-resolution quality even at **scales up to 16x**.

RESEARCH PROJECTS

3D Reconstruction and Restoration of Underwater Images (Computational Imaging Lab, IITM) [*Link*] — *Dr. Kaushik Mitra*

- Proposed **U2NeRF**, a fully self-supervised transformer-based framework for joint **underwater image restoration** and **neural 3D reconstruction**, embedding physics-informed light modeling into the NeRF pipeline.
- Achieved state-of-the-art results on the newly curated Underwater View Synthesis (UVS) benchmark across 12 calibrated scenes, with **+11% perceptual similarity** and **+4% restoration quality** over prior methods.
- Developed as part of **Master's thesis** at IIT Madras and later published at **ICLR 2024** (Tiny Papers Track).

Efficient Dense Video Captioning using Multimodal Transformers (NYU CDS) [*Link*] — *Dr. Kyunghyun Cho*

- Built a two-stage **dense video captioning** pipeline for untrimmed videos, leveraging **multimodal transformers** to jointly process audio-visual inputs for temporal event localization and caption generation.
- Achieved competitive performance with SOTA-generated caption quality, while reducing memory usage by **~45%** on lengthy videos.

Adversarial Prompt Engineering for Large Language Model (LLM) Jailbreaking (NYU) [*Link*] — *Dr. Siddharth Garg*

- Designed adversarial prompts using KOV-MCTS, successfully jailbreaking **40%** of harmful queries passed to GPT-3.5 Turbo.
- Delivered an **8.4×** improvement in bypass success rate through surrogate model refinement and heuristic optimization techniques.

Counterfactual Image Generation Using Text Guidance (IITM) [*Link*] — *Dr. Sutanu Chakraborti*

- Produced **text-conditioned counterfactuals** from opposite-class input samples while maintaining >90% content preservation.
- Performed **targeted inpainting** on ~10% high-confidence regions using **Stable Diffusion** and **CLIP** for localized attribute control.

LEADERSHIP AND ACHIEVEMENTS

- **Competitions**: "Inter-IIT Tech Meet" (3rd out of 23 other IITs) at IITK, "NTSE Scholar" (top 0.4%) - organized by NCERT, India.
- **Authored and published a paper**: "U2NeRF: Unifying Unsupervised Underwater Image Restoration and Neural Radiance Fields" at International Conference on Learning Representations (ICLR), May 2024, Tiny Papers Track. [*Link*]
- **Certifications**: Machine Learning in Production, Python for Computer Vision with OpenCV and Deep learning, LLMOps.