

Arnav Samal

arnavsm@cmu.edu | (412) 812-7576 | arnavsm.github.io | github.com/arnavsm | linkedin.com/in/arnavsamal

EDUCATION

Carnegie Mellon University - School of Computer Science

Master of Science in Computer Vision (Robotics Institute)

Relevant Coursework: Advanced Computer Vision, Visual Learning & Recognition, Learning for 3D Vision

Pittsburgh, PA

Dec 2027

National Institute of Technology, Rourkela

Bachelor of Technology, Computer Science & Engineering | GPA: 9.13/10.0 | Dept Rank: 7/153

Relevant Coursework: Computer Vision, Deep Learning, Machine Learning, Natural Language Processing

Rourkela, India

May 2026

SKILLS

Programming Languages: Python, C++, C, SQL

Tools: AWS (EC2, EBS), Linux, Git, Conda, Weights & Biases, Docker, Jupyter, LaTeX

Libraries/Frameworks: PyTorch, OpenCV, HuggingFace, Sci-kit Learn, SciPy, NumPy, Matplotlib, Flask

Computer Vision & ML: Vision-Language Models (VLMs), Diffusion Models, 3D Vision, Vision Transformers, Self-Supervised Learning, Explainability (Concept Bottleneck Models/CBMs, Post-hoc XAI)

RESEARCH EXPERIENCE

IISc Bangalore | Visual Computing Lab, Dept. of CDS

Research Intern (Advisor: Dr. Anirban Chakraborty)

Bengaluru, India

Jan 2026 – Jun 2026

- Benchmarked multiple SOTA federated learning algorithms (FedAvg, FedDyn, SCAFFOLD, etc.) on CIFAR-10/100 and TinyImageNet across ResNet (18-101), finding overparameterization did not consistently mitigate non-IID client drift.
- Proposed and led evaluation of the generator-critic loop behind SurfScaff3D, a VLM-based agentic text/image-to-3D framework, achieving 0% generation/editing failure, 65% prompt alignment, and up to 33% lower geometric distortion than the strongest baseline.

IIT Hyderabad | Machine Learning & Vision Group, Dept. of CSE

Research Intern (Advisor: Prof. Vineeth N Balasubramanian, with Microsoft Research)

Hyderabad, India

May 2025 – Dec 2025

- Investigated concept quality in CBMs by analyzing spatial concept locality to diagnose concept leakage; processed 2TB+ of medical imaging data using locality-to-concept attribution metrics.
- Found that CBM models attend 45% more to spurious background regions than diagnostically relevant features, exposing a key failure mode in visual concept faithfulness for medical-imaging models.

IIT Hyderabad | Data-Driven Intelligence & Learning Lab, Dept. of AI

Research Intern (Advisor: Dr. Konda Reddy Mopuri)

Hyderabad, India

May 2024 – Aug 2024

- Evaluated spatial representations across Vision Transformer architectures by benchmarking post-hoc explainability techniques, establishing clear tradeoffs between model interpretability and efficiency scaling.
- Reduced inference compute by 66% while maintaining baseline classification accuracy within 1-3% by implementing token pruning strategies that leverage top-K token overlay and patch importance metrics.

PUBLICATIONS

- Anay Loya*, Rishi Gupta*, **Arnav Samal**, and Anirban Chakraborty. "SurfScaff3D: Bridging Geometric and Semantic Representations for LLM-based 3D Object Generation." *WACV 2027 (Under Review)*

PROJECTS

Explainable Sketch to Face Generation – NIT Rourkela, Research Project | Code

Aug 2025 – Nov 2025

- Fine-tuned Stable Diffusion v1.5 using ControlNet and LoRA, achieving an FSIM of 0.7386 to outperform baseline architectures in structural fidelity.
- Engineered a temporal XAI framework adapting Grad-CAM to regression U-Nets, mapping gradient attribution across DDIM sampling timesteps to visualize feature-anchored sketch-to-face reasoning.

Capsule Vision Challenge 2024 – Global Competition | Preprint, Code

Aug 2024 – Oct 2024

- Designed and built a multi-model CNN and Vision Transformer ensemble pipeline, mitigating severe class imbalance via focal loss optimization and weighted sampling strategies.
- Secured global Rank 5 out of 500+ participating teams.