Aniket Agarwal

aniketa@andrew.cmu.edu aniket-agarwal1999 aniket-agarwal1999 aniket-agarwal1999.github.io

Education

Carnegie Mellon University - School of Computer Science Pittsburgh, PA
Master of Science in Robotics | ***CGPA:*** *3.89/4.0* *Aug. 2023 - Present*

Indian Institute of Technology Roorkee Roorkee, India
Integrated Masters of Science in Applied Mathematics | ***CGPA:*** *7.82/10* *July 2017 – May 2022*

Experience

Carnegie Mellon University Nov 2023 – Present
Research Assistant | Advisor: Prof Laszlo Jeni *Pittsburgh, Pennsylvania*

- Worked on the **DARPA Triage Challenge**, focusing on remotely assessing casualty physiological signs using multimodal data (RGB imagery, thermal imaging and speech inputs), securing 4th place out of 12 teams from academia and industry.
- Fine-tuned LLM and TTS models for injury assessment and verbal alertness, quantized them for efficiency, and developed real-time ROS inference pipelines deployed on NVIDIA Jetson systems.
- **MSR Thesis Project:** Developing a video-based Human Mesh Recovery (HMR) system leveraging SMPL representation to reconstruct human meshes across frames with enhanced temporal smoothness, resulting in significant improvements in 3D pose trajectory stability through advanced trajectory space analysis

Microsoft July 2022 – Aug 2023
Data & Applied Scientist | MSAI Team *Hyderabad, India*

- Developed Learning to Rank (LTR) systems to improve email and message relevance for the MSAI team
- Designed and implemented an ML-driven enhancement to the FastRank method, incorporating dimensionality reduction based on feature importance, resulting in a 10% improvement in existing NDCG performance
- Collaborated on migrating the LTR ML pipeline from Aether to Azure ML, ensuring smooth integration and assisting in validating parity in NDCG metrics between platforms to maintain performance consistency

Princeton University Oct. 2021 – Sept 2022
Research Collaboration | Advisor: Prof Karthik Narasimhan *New Jersey, USA*

- Developed a novel automated pipeline for creating long form benchmarks for Long Video Understanding (LVU) using publicly available sports videos
- Conducted comprehensive human evaluations and benchmarking using advanced Video Understanding models on a newly developed dataset, identifying significant performance differences ($\approx 70\%$), thereby validating the benchmark's quality and impact on advancing video understanding research.

King Abdullah University of Science and Technology (KAUST) May 2020 – April 2021
Research Collaboration | Advisor: Prof Mohamed Elhoseiny *Saudi Arabia*

- Developed two new benchmarks specifically for long-tailed detection in the Visual Relationship Recognition (VRR) task. Also proposed a novel loss and augmentation technique to help the classification of tail classes
- Subsequently, proposed a novel Transformer based architecture for long-tailed VRR task with a memory module, improving the model performance by $\approx 20\%$ on tail classes

Publications

1. **Aniket Agarwal***, Alex Zhang*, Karthik Narasimhan, Igor Gilitschenski, Vishvak Murahari, Yash Kant. "Building Scalable Video Understanding Benchmarks through Sports." *Data-Centric ML Research Workshop at International Conference on Learning Representations (**DMLR@ICLR**), 2024.* [Project Page]
2. Jun Chen, **Aniket Agarwal**, Sherif Abdelkarim, Deyao Zhu, Mohamed Elhoseiny. "RelTransformer: A Transformer-Based Long-Tail Visual Relationship Recognition." *Computer Vision and Pattern Recognition Conference*(**CVPR**), 2022. [Paper Link]
3. **Aniket Agarwal***, Sherif Abdelkarim*, Panos Achlioptas, Jun Chen, Jiaji Huang, Boyang Li, Kenneth Church, Mohamed Elhoseiny. "Exploring Long Tail Visual Relationship Recognition with Large Vocabulary." *International Conference on Computer Vision* (**ICCV**), 2021. [Project Page]

Technical Skills

Graduate Coursework: Computer Vision, Geometry-based Methods in Vision, Learning for 3D vision, Machine Learning
Frameworks/Libraries: PyTorch, PyTorch Lightning, NumPy, ROS, TensorFlow, OpenCV
Developer Tools: Linux, Vim, Git, Slurm, AWS, Google Cloud Platform, Microsoft Azure
Programming Languages: Python, C++