Pranay Gupta

Email: pranaygu@andrew.cmu.edu | +1 412-626-9880 | Google Scholar | LinkedIn | Github | Homepage

Education

Carnegie Mellon University — Pittsburgh, PA
Ph.D. in Robotics; Master of Science in Robotics; GPA ; 4.08/4 — *Aug. 2022 – Present*

International Institute of Information Technology — Hyderabad, India
Bachelor and Master of Science in Computer Science, GPA ; 8.03/10 — *Aug. 2016 – Jul 2021*

Research Interests

Human Robot Interaction, Assistive/Autonomous Driving, Human-AI Alignment, Learning for Robotics

Experience

Graduate Research Assistant — Oct 2022 – Present
The Robotics Institute, Carnegie Mellon University — *Pittsburgh, PA*

- Devised a novel approach for modeling a driver's situational awareness (SA) using their eye-gaze. CoRL 2024, VAM-HRI 2024 Workshop
- Curated a novel dataset and designed a counterfactual reasoning based approach to identify important objects in driving scenarios. IEEE RA-L 2024

Predoctoral Apprentice — May 2021 – July 2022
TCS research — *Delhi, India*

- Implemented 3D CNNs to approximate implicit functions for 3D Single View Reconstruction (SVR).
- Employed an energy based out-of-distribution (OOD) detection classifier to increase robustness for SVR.

Undergraduate Research Assistant — June 2018 – June 2021
Center for Visual Information Technology(CVIT), IIIT-H — *Hyderabad, Telengana*

- Explored into-the-wild and out-of-context human action recognition. IJCV-2021
- Devised a VAE backed approach which learned syntactically aware embeddings for zero shot skeleton action recognition. Achieved SOTA results on the NTU-60 and 120 datasets. 2021 IEEE ICIP.

Applied Scientist Intern — Jun. 2020 – Aug.2020
Amazon India — *Bengaluru, Karnataka*

- Trained LLMs within a siamese network to estimate semantic similarity between query and product description.
- Successfully adapted LLMs trained on data in English to European languages using transfer learning techniques.

Google Summer of Code Intern — Jun 2018 – Aug 2018
Purr-Data — *Remote*

- Successfully updated purr-data's core and the external libraries from single precision float to double precision.

Projects

Leveraging VLMs for Zero-Shot, Personalization of Multi-Object Rearrangement

- Devised an approach that leveraged pre-trained VLMs for preference-aligned object rearrangement. Under Submission at HRI 25, Human-LLM interaction workshop at HRI-24
- Demonstrated feasibility through successful implementation in a one-step table setting task. Presented at the Human-LLM interaction workshop at HRI-24

News-KVQA

- Curated a large scale video question answering dataset (12k videos, 1 million QA pairs). Automated question generation using videos, subtitles and knowledge base facts.
- Proposed a multi-modal, LLM backed approach that processed visual, textual and factual data for question answering. Published at PAKDD-22.

3D Scene Reconstruction using Monocular Image

- Developed a pipeline that systematically performed object segmentation and pose estimation using YOLOv3 and cube-RCNN and performed 3D reconstruction and localization of each detected object using pixelNERF and iNERF.

Improving PlanT

- Improved the results from the CORL 2022 paper "PlanT: Explainable Planning Transformers via Object-Level Representations" by ensuring consistency of the frame of reference across the inputs and maintaining a history of states.

Distributed Attendance System

- Automatic Attendance System based on Face Detection and Recognition via Facenet. Distributed System, worked simultaneously with multiple cameras.

Bash Shell

- Implemented as a part of a course project in Operating Systems course. Developed a unix shell in C. Implemented basic functionalities like killing a process, input/output redirection, piping and signal handling

Technical Skills

Additional Research Interests: Machine Learning/ Deep Learning for Computer Vision, Multi-modal learning, 3D Computer Vision and NLP
Languages: Python, Matlab, C/C++, HTML/CSS
Developer Tools: Git, Docker, Vim, VS Code, AWS
Libraries: Pytorch, Pytorch-3D, CARLA, PyBullet, Open3D, Opencv, Scikit-Learn, Pandas, NumPy, Matplotlib