

Anuva Banwasi

anuva.banwasi@gmail.com - [anuvabanwasi.github.io](https://github.com/anuvabanwasi) - [linkedin.com/in/anuva-banwasi-941617212](https://www.linkedin.com/in/anuva-banwasi-941617212)

EDUCATION

Columbia University, School of Engineering and Applied Sciences B.S. in Computer Science, Concentration: Intelligent Systems (GPA: 3.996 / 4.0) <i>cum laude</i> , C.P. Davis Scholar	Graduated May 2024
■ Machine Learning, AI, NLP, Deep Learning for Computer Vision, Computational Robotics, Robot Learning, Spoken Lang Processing	
Stanford University Masters of Computer Science	In Progress
■ CS231a - 3D Perception & Reconstruction, CS227 - Robot Perception, CS234 - Reinforcement Learning, CS224R Deep RL	

RESEARCH EXPERIENCE

Researcher, Interactive Perception and Robot Learning Lab, Stanford University Mentors: Prof. Jeannette Bohg, Dr. Cherie Ho	April 2025 – present Palo Alto, CA
■ <i>Mobile Manipulation & Interaction Memory</i> . Integrating multimodal perception with a persistent memory for adaptive task planning. Grounded VLM affordances in physical feasibility and captured object properties inferred from real-time interaction in a lightweight scene graph. Demonstrated failure-aware re-planning in IsaacSim & real-world experiments with TidyBot++.	
■ <i>Dexterous manipulation</i> . Built grasp evaluation model for multi-finger grasping leveraging 3D Gaussian Splatting for object representation. Encoded 3DGS with custom VQ-VAE model implemented in PyTorch, concatenated with Allegro grasp config and optional DINO-based image embedding. Trained grasp evaluator and refiner on 2 TB of data (3.5M grasps, 4.3K objects). Achieved SoTA 95% grasp success rate in Isaac Gym, outperforming NeRF and BPS representations.	
Self Evaluation Using Zero-shot Learning for Long-Horizon Tasks, Columbia University Mentor: Prof. Shuran Song	Nov 2022 – Sept 2023 New York, NY
■ <i>Robot learning & Computer Vision</i> . Led a team of 4 researchers to build zero-shot learning framework that assesses step-by-step task execution for long-horizon tasks. Algorithm describes predicted step outcome with LLM, evaluates video data against language outcome description with VLM, and assesses incremental progress with custom scoring metric.	
■ First author on paper published in International Conference for Robotics and Computer Vision (ICRCV 2023) .	
Undergraduate Researcher, Speech Lab, Columbia University Mentor: Prof. Julia Hirschberg	May 2023 – Dec 2023 New York, NY
■ <i>Multimodal ML & Conversational AI</i> . Built multimodal system integrating speech processing, text + audio analysis, and fine-tuned LLMs to detect emotion shifts such as from neutral or positive to frustration or anger. Trained on 150 videos and 10GB of data. Improved accuracy on IEMOCAP benchmark from 40% to 78% via LLM fine-tuning and few-shot learning.	
Student Researcher, Creative Machines Lab, Columbia University Mentor: Prof. Hod Lipson	Sept 2022 – Dec 2022 New York, NY
■ <i>Deep learning & robotics</i> . Proposed model to learn and represent the kinematics of a robot using only 2d single-view video using NeRF-inspired approach. Designed, trained, and tested neural network in PyTorch that predicts space occupancy given input of 3D coordinates and joint angles. Obtained projection matrix from PyBullet and computed field of view and focal length for modified NERF model. System generates a fuzzy 3D model of robot that similarly matches ground truth.	

PROJECTS

Adv Topics in Deep Learning Project, Columbia University Mentor: Prof. Peter Belhumeur	Jan 2024 – Jun 2024 New York, NY
■ <i>LLM Guardrails & controllability</i> . Led a team of four evaluating LLM guardrail frameworks (Meta Llama Guard, NVIDIA NeMo Guardrails) for custom-defined datasets. Performed supervised fine-tuning with QLoRA and PEFT. Built ensemble model combining guardrail frameworks with MLP and random forest, enhanced prompt-classification accuracy from 89% to 99%.	
Machine Learning Research, Columbia University Mentor: Prof. Michelle Levine	Oct 2023 – Jul 2024 New York, NY
■ <i>ML & statistical forecasting</i> . Built forecasting models using vector autoregression to predict electric vehicle (EV) adoption based on policy incentive data and charging infrastructure growth over the past decade. Culminated in first-author publication at the International Conference on Time Series and Forecasting (ITISE 2024) .	

HONORS & AWARDS

Best Presentation, Artificial Intelligence & Pattern Recognition International Conference	Sep 2023
Tau Beta Pi Engineering Honor Society Member (awarded to top 10% at Columbia Engineering)	Nov 2023
PearX Summer 2024 Cohort - top-tier startup accelerator by Pear VC (<1% acceptance, 4.5k applications)	March 2024
Columbia Venture Competition Winner - 3rd place, awarded 10k seed funding	Apr 2021
C.P. Davis Scholar, Columbia Undergraduate Scholars Program - 1 of 25 selected of Columbia Engineering class	Sep 2020

SKILLS

Python, Java, C, PyTorch, TensorFlow, ROS, IsaacGym, OpenCV, scikit-learn, Git, Linux/Unix, Docker, SQL, TypeScript, HTML

WORK EXPERIENCE

Founder, Self Eval

Jun 2023 - Aug 2024

Building the future of real-time verification & quality assurance for high-stakes manufacturing.

Palo Alto, CA

- Smart co-pilot uses zero-shot multimodal ML to evaluate the operator's actions in real-time. Funded by Pear VC, Menlo Park. Partnered with Johnson & Johnson to verify workmanship on their line for medical robots with over 100 operators.

Software Engineer, MongoDB

Sept 2024 - present

Search & AI team

San Francisco, CA

- Developed drill-sideways ("hold one out") faceting for lexical search leveraging Apache Lucene. Estimated to increase usage by top customers by 30%. Implemented faceting for multi-value queries in Atlas Search. Established storage-optimized search nodes with AWS, saving customers 50% on storage costs. Backend development in Java, Go, and TypeScript.

Machine Learning Intern, Altair Engineering

May 2021 - Aug 2021, Apr 2022 - Jun 2022

3D Modeling and HyperWorks Simulation team

Troy, MI

- *Computer Vision*. Developed model for 3D object detection on CAD model. Trained & tested YOLOv5 object detection model in Python/PyTorch to locate and classify connectors in images of 3D CAD model. Implemented back-projection algorithm to project 2D coordinates of feature in multi-view images to 3D bounding box around the feature in CAD model via view matrix.
- Recognized by CTO, reduced labeling time for CAD engineers from 2 weeks to < 2 hours.

Software Engineer Intern, National Institutes of Health

Jun 2021 - Aug 2021

Coding it Forward Fellowship

Bethesda, MA

- *Natural Language Processing*. Built supervised ML model in Python to classify grant applications based on their title and abstract text ([GitHub code](#)). Model correctly classified 95% of proposals with an area under ROC curve of 0.90. Deployed app with Docker, YAML & Streamlit - continued use by 75+ analysts at Office of Data Science and Emerging Technologies (ODSET).

PEER-REVIEWED PUBLICATIONS

- MessyMem: Learning-from-Doing Memory for Mobile Manipulation**
Anuva Banwasi, William Muckelroy III, Priya Sundaresan, Linfeng Zhao, Jeannette Bohg, Cherie Ho
In submission.
Keywords: Robot memory, interaction-grounded memory & planning, Robot Learning, Mobile Manipulation.
- MessyNav: Zero-shot Navigation via Manipulation in Messy Environments**
Cherie Ho, Kenneth Llontop, Dylan Zhou, **Anuva Banwasi**, Priya Sundaresan, Carlota Parés-Morlan, Jeannette Bohg
In submission.
Keywords: Robot Learning: Found Models, Planning, Autonomous Navigation, Perception
- Self Evaluation Using Zero-shot Learning**
Anuva Banwasi, Xinghua Sun, Rohith Ravindranath, March Vazquez
2023 5th International Conference on Robotics and Computer Vision (ICRCV), Nanjing, China
doi: 10.1109/ICRCV59470.2023.10329149.
Keywords: Large language models; visual language models; language grounding for vision and video understanding; robotics
Date published: 30 November 2023, Link: <https://doi.org/10.1109/ICRCV59470.2023.10329149>
- Promoting Electric Vehicle Growth through Infrastructure and Policy: A Forecasting Analysis**
Anuva Banwasi, Adele M Sinai, Brennan Xavier McManus.
10th International Conference on Time Series and Forecasting, Gran Canaria, Spain, Engineering Proceedings. 2024.
Date published: 18 July 2024, Link: <https://doi.org/10.3390/engproc2024068060>
- A Comparative Analysis of Guardrail Frameworks for LLMs and Enhancement with Ensemble Techniques**
Anuva Banwasi, Samuel Friedman, Michael Khanzadeh
64th Annual Meeting of the Association for Computational Linguistics (ACL 2026)
Submitting to ACL 2026, Link: <https://github.com/llmguardrails/LLMGuardrailsPaper>
- Hands Off, Splat On: Dexterous Grasp Evaluator with 3D Gaussian Encodings**
Anuva Banwasi, Yi Du, Ashwin Mahendran
Link: <https://github.com/Dude346/CS231A>