Chancharik Mitra

+1 (412) 708-4914 | cmitra@cs.cmu.edu | LinkedIn | Google Scholar | GitHub | Website | Pittsburgh, PA

Education

Carnegie Mellon University — Pittsburgh, PA, USA
M.S. in Machine Learning ***GPA: 3.95/4.00*** — *Aug. 2024 – Present*

University of California, Berkeley — Berkeley, CA, USA
B.S. in EECS - Research Honors; ***GPA: 3.87/4.00*** — *Aug. 2021 – May 2024*

Moon Area High School — Pittsburgh, PA, USA
High School Diploma; ***GPA weighted: 4.70/4.00 GPA unweighted: 4.00/4.00*** — *Aug. 2017 – Jun. 2021*

College Courses during High School: Harvard University, University of Pittsburgh, and Robert Morris University - all 4.00 GPA

Research Interests

My research aims to create AI systems that leverage pretrained knowledge to efficiently adapt to new tasks and environments. I pursue two core directions: (1) few-shot adaptation via mechanistic interpretability, targeting task-relevant representations without expensive fine-tuning, and (2) systematic generalization through cognitive and linguistic frameworks such as compositionality and pragmatics. I apply these methods to vision-language-action applications such as video reasoning, robotics, and education.

See research overview figure: https://chancharikmitra.github.io/media/research-overview.png

Research Experience

Carnegie Mellon University - Robotics Institute — Pittsburgh, PA
Research Assistant - Advised by Dr. Deva Ramanan — *May 2024 – Present*

- Conducting research on *few-shot mechanistic interpretability for model adaptation*, leveraging activation steering and feature extraction to enable **efficient adaptation of VLMs and VLAs** from limited examples across vision-language **classification**, multimodal **reward modeling**, **object detection**, and **robotic applications**
- Developing evaluation methods for spatiotemporal compositionality in video-language models, examining event properties such as object-attribute-relationship state transitions, event order, and speed
- Advancing compositional **video understanding and generation** through camera motion research: developed CameraBench and finetuned SOTA LMM for camera understanding; currently working on compositional cinematic video captioning and generation with fine-grained camera control

Berkeley AI Research (BAIR) Lab — Berkeley, CA
Researcher - Advised by Dr. Trevor Darrell and Dr. Dan Klein — *Apr. 2022 – August 2025*

- Advanced *few-shot mechanistic interpretability for model adaptation* through task vectors, activation steering, and feature extraction methods for Large Multimodal Models (LMMs), demonstrating superior performance over baselines including LoRA
- Developed methods for compositional vision-language understanding, leveraging scene graphs and chain-of-thought prompting to enhance multimodal reasoning capabilities
- Applied pragmatic reasoning principles to improve contextual language grounding in embodied settings for 3D object disambiguation
- Contributions **featured in IBM's Granite Vision models and AllenAI's MOLMO**

University of California, Berkeley – EECS Department — Berkeley, CA
Research Assistant - Advised by Dr. Gireeja Ranade and Dr. Narges Norouzi — *May 2023 – August 2025*

- Conducting research on multimodal question-answering and retrieval-augmented generation for educational AI applications
- *Edison*: Deployed LLM-based QA system in production for UC Berkeley's DATA 100 course (∼1000 students), with document retrieval pipeline and pedagogically-informed prompting strategies
- *Askademia*: Developed multi-institutional educational QA platform with advanced retrieval methods and evaluation frameworks for assessing pedagogical value

National Center for Supercomputing Applications (NCSA @ UIUC) Champaign, IL

Undergraduate Researcher - Advised by Dr. Aiman Soliman and Dr. Zeynep Madak-Erdogan *Apr. 2022 – Aug. 2024*

- Conducted research as an REU fellow and undergraduate researcher (post-program) on leveraging machine learning techniques to segment spatial transcriptomics data of ER+ breast cancer tissue samples
- First-authored paper published in the *IEEE MASS '22* conference and gave an oral presentation of the research at REUNS workshop (won **Best Paper** award) as well as at the Illinois Summer Research Symposium (ISRS)
- Developed an explainable, model-agnostic biomarker mining pipeline for spatial transcriptomics data

Work Experience

Carnegie Mellon University – Robotics Institute Pittsburgh, PA

Research Assistant - Advised by Dr. Deva Ramanan *May 2025 – August 2025*

- Conducting research on generalizable vision-language-action reasoning via few-shot mechanistic interpretability for model adaptation and compositional video understanding and generation
- For detailed research directions and related publications, please refer to the corresponding research section entry

Carnegie Mellon University – Machine Learning Department Pittsburgh, PA

10-718 Machine Learning in Practice: Teaching Assistant *Aug. 2025 – Present*

- Teaching assistant for graduate course covering ML specification, evaluation, robustness, interpretability, and other concepts essential for successful research and industry ML projects
- Responsibilities include hosting office hours, grading assignments, lecture support, and exam question development

PromptAI Berkeley, CA

Research Intern *Jun. 2024 – Sept. 2024*

- Developed an open-source video LMM model using InternVideo2 video encoder and LLaMA3 language model
- Created video-image in-context learning dataset using Panda-70M and OBELICS dataset to enable video in-context learning (ICL) capabilities
- Investigated the applications of video ICL and multimodal task vectors for open-source models, demonstrating that VILA-1.5 and LLaVA-OneVision can surpass GPT-4o on visual tasks using these methods

University of California, Berkeley – EECS Department Berkeley, CA

EECS 16B Designing Information Devices and Systems II: Lead Discussion and Content TA *Jun. 2022 – May 2024*

- Taught weekly discussion sections of over 50 students
- Developed discussion, midterm, and final exam problems for topics ranging from DC/AC circuits, differential equations, linear algebra, control, SVD/PCA, and classification
- For samples of my content creation, please refer to any discussion or exam on the EECS 16B course page
- Restructured course to emphasize circuit, Bode Plot, and SVD/PCA concepts

National Center for Supercomputing Applications (NCSA @ UIUC) Champaign, IL

REU FoDOMMaT Undergraduate Research Fellow *Jun. 2022 – Aug. 2023*

- Accepted to the selective REU-FoDOMMaT summer research program held at NCSA on-campus at the University of Illinois, Urbana-Champaign (UIUC). Invited to continue research post-program and to return as a Research Fellow the following summer.
- For more details, please refer to the corresponding research section entry.

Publications

More information (links, project pages, etc.) available on my website: https:// chancharikmitra.github.io/

Submitted (In Review):

[1] Activation Reward Models for Few-Shot Model Alignment. Tianning Chai*, Chancharik Mitra*, Brandon Huang, Gautam Rajendrakumar Gare, Zhiqiu Lin, Assaf Arbelle, Leonid Karlinsky, Rogerio Feris, Trevor Darrell, Deva Ramanan, and Roei Herzig. In: *Annual Meeting of the Association for Computational Linguistics (ACL)*. 2026. Submitted.

[2] Building a Precise Video Language with Human-AI Oversight. Chancharik Mitra*, Siyuan Cen*, Zhiqiu Lin*, Isaac Li, Yuhan Huang, Yu Tong Tiffany Ling, Hewei Wang, Irene Pi, Shihang Zhu, Yili Han, Yilun Du, and Deva Ramanan. In: *TBD*. 2026. Submitted.

[3] Mechanistic Finetuning of Vision-Language-Action Models via Few-Shot Demonstrations. Chancharik Mitra*, Yusen Luo*, Raj Saravanan*, Dantong Niu, Anirudh Pai, Jesse Thomason, Trevor Darrell, Abrar Anwar, Deva Ramanan, and Roei Herzig. In: *TBD*. 2026. Submitted.

Published:

[4] CameraBench: Towards Understanding Camera Motions in Any Video. Zhiqiu Lin*, Siyuan Cen*, Daniel Jiang, Jay Karhade, Hewei Wang, Chancharik Mitra, Tiffany Ling, Yuhan Huang, Sifan Liu, Mingyu Chen, et al. In: *Conference on Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*. 2025. **Spotlight**.

[5] Enhancing Few-Shot Vision-Language Classification with Sparse Attention Vectors. Chancharik Mitra, Brandon Huang, Tianning Chai, Zhiqiu Lin, Assaf Arbelle, Rogerio Feris, Leonid Karlinsky, Trevor Darrell, Deva Ramanan, and Roei Herzig. In: *International Conference on Computer Vision (ICCV)*. 2025.

[6] Analyzing Pedagogical Quality and Efficiency of LLM Responses with TA Feedback to Live Student Questions. Chancharik Mitra*, Mihran Miroyan*, Rishi Jain*, Gireeja Ranade, and Narges Norouzi. In: *ACM Technical Symposium on Computer Science Education (SIGCSE)*. 2025.

[7] EduMod-LLM: A Modular Approach for Designing Flexible and Transparent Educational Assistants. Meenakshi Mittal, Rishi Khare, Mihran Miroyan, Chancharik Mitra, and Narges Norouzi. In: *AAAI Conference on Artificial Intelligence*. 2025.

[8] Raising the Bar: Automating Consistent and Equitable Student Support with LLMs. Meenakshi Mittal*, Azalea Bailey*, Victoria Phelps*, Rose Niousha, Mihran Miroyan, Chancharik Mitra, Rishi Jain, Gireeja Ranade, and Narges Norouzi. In: *ACM Technical Symposium on Computer Science Education (SIGCSE) Poster*. 2025.

[9] Less is More Tokens: Efficient Math Reasoning via Difficulty-Aware Chain-of-Thought Distillation. Abdul Waheed*, Chancharik Mitra*, Laurie Z. Wang*, Deva Ramanan, and Bhiksha Raj. In: *NeurIPS Workshop on Multimodal Algorithmic Reasoning*. 2025.

[10] Multimodal Task Vectors Enable Many-Shot Multimodal In-Context Learning. Brandon Huang*, Chancharik Mitra*, Assaf Arbelle, Leonid Karlinsky, Trevor Darrell, and Roei Herzig. In: *Conference on Neural Information Processing Systems (NeurIPS)*. 2024.

[11] Compositional Chain-of-Thought Prompting for Large Multimodal Models. Chancharik Mitra, Brandon Huang, Trevor Darrell, and Roei Herzig. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024.

[12] Which One? Leveraging Context Between Objects and Multiple Views for Language Grounding. Chancharik Mitra*, Abrar Anwar*, Rodolfo Corona, Dan Klein, Trevor Darrell, and Jesse Thomason. In: *North American Chapter of the Association for Computational Linguistics (NAACL)*. 2024.

[13] Elevating Learning Experiences: Leveraging Large Language Models as Student-Facing Assistants in Discussion Forums. Chancharik Mitra*, Mihran Miroyan*, Rishi Jain*, Vedant Kumud, Gireeja Ranade, and Narges Norouzi. In: *ACM Technical Symposium on Computer Science Education (SIGCSE) Poster*. 2024.

[14] RetLLM-E: Retrieval-Prompt Strategy for Question-Answering on Student Discussion Forums. Chancharik Mitra*, Mihran Miroyan*, Rishi Jain*, Vedant Kumud, Gireeja Ranade, and Narges Norouzi. In: *AAAI Conference on Artificial Intelligence*. 2024.

[15] Geospatial Analysis To Quantify Spatial Heterogeneity Of Tumor Microenvironment. Jin Young Yoo, Chancharik Mitra, Yang Yue, Aiman Soliman, and Zeynep Madak-Erdogan. *Journal of the Endocrine Society*. 2023.

[16] Spatial Analysis of Tumor Heterogeneity Using Machine Learning Techniques. Chancharik Mitra, Jin Young Yoo, Zeynep Madak-Erdogan, and Aiman Soliman. In: *IEEE International Conference on Mobile Ad Hoc and Smart Systems (MASS)*. 2022. **Best Paper Award** at REUNS Workshop.

Awards & Achievements

NeurIPS Datasets and Benchmarks Spotlight for *CameraBench* | 2025

National Science Foundation Graduate Research Fellowship | 2025

UC Berkeley "Outstanding Graduate Student Instructor Award" | 2024

Berkeley EECS Research Honors Program Recipient | 2023

Google CS Research Mentorship Program Recipient | 2023

2x UC Berkeley College of Engineering Dean's List Award | top 10% semester GPA - 2022, 2023

Best Paper Award | Awarded "Best Paper" at the *8th National Workshop for REU Research in Networking and Systems* colocated with the *19th IEEE International Conference on Mobile Ad-Hoc and Smart Systems* - 2022

USACO Gold | Achieved a perfect score on USA Computing Olympiad Bronze and Silver division competitions, earning Gold-level USACO standing - 2020/2021

2x American Invitational Mathematics Exam (AIME) Qualifier | Earned a top 5% score nationally on the American Mathematics Competition 12 (AMC 12) exam to qualify for the AIME exam, going on to score in the top 20% of AIME qualifiers - 2020/2021

Projects

All About Autoencoders | *GitHub*

- Comprehensive theoretical and practical survey of autoencoder architectures for Berkeley's CS 182 Deep Learning
- Covers vanilla, variational, and vector-quantized VAE implementations, theoretical exercises, and illustrative visualizations including latent space exploration, denoising, and ablation studies
- Received full marks plus extra credit for exceeding project requirements
- Resource incorporated into a comprehensive machine learning course at UIUC

Volunteering & Student Organizations

Computer Science Mentors Berkeley, CA

Senior Teaching Mentor *Jun. 2022 – May 2024*

- Volunteer Senior Mentor for CS 70 (Discrete Mathematics & Probability) – STEM Excellence through Equity and Diversity (SEED) Scholars section
- Responsibilities include teaching a section of students, leading a group of junior mentors, and contributing to CS 70's content task force: recording video walkthroughs of problems as well as creating exam review material

IEEE – Berkeley Student Branch Berkeley, CA

Technical Operations Officer *Jun. 2022 – August 2023*

- Responsible for maintaining IEEE's server which houses the organization's website and technical project resources

STEM Coding Lab Pittsburgh, PA

Volunteer Coding Instructor and Curriculum Developer *Aug. 2018 – June 2022*

- STEM Coding Lab is a nonprofit organization with the mission of teaching children from underserved communities in Pittsburgh how to code
- Teach programming classes (e.g. Python, Java, HTML/CSS/Javascript) of 20+ students (grades K-12) multiple times a week
- Developed an 8-week Python course that was used to teach students at the Pittsburgh Public Schools during the COVID-19 pandemic

Skills

Programming: Python, Pytorch, NumPy, Scikit Learn, Java, C, Assembly, Tensorflow, SQL

Technologies and Software: Git, Linux, Windows, MacOS, Arduino, Tableau, Android Studio

Languages: English (Native), Bengali (Native), Hindi (Advanced), German (Intermediate)

Relevant Coursework

CMU: Computer Vision (16-720) - Grade A+, Advanced Natural Language Processing (11-711) - Grade A+, Machine Learning in Practice - Grade A+

Berkeley: CS 288 Graduate-level Natural Language Processing - Grade A, PLANTBI 200B Graduate-level Genomics and Computational Biology - Grade A, CS 182 Deep Learning (completed graduate course requirements for extra credit) - Grade A, CS 188 Artificial Intelligence - Grade A+, CS 189 Machine Learning - Grade A, MCELLBI 132 Biology of Human Cancer - Grade A, EECS 16B Designing Information Devices and Systems I (Linear Algebra, Differential Equations, and Circuits) - Grade A+