

Amita Kamath

kamath.a.amita@gmail.com · <https://amitakamath.github.io> · (she/her)

Natural Language Processing · Vision-Language · Machine Learning

EDUCATION

University of California at Los Angeles Ph.D. Candidate in Computer Science <i>Awarded Graduate Dean's Scholar Award.</i> Advisors: Prof. Kai-Wei Chang, Prof. Ranjay Krishna (University of Washington)	2022 – Present GPA: 4.0
Stanford University M.S. in Computer Science (Specialization: Artificial Intelligence) <i>Awarded Distinction in Research.</i> Advisor: Prof. Percy Liang	2017 – 2020 GPA: 3.99
National Institute of Technology Karnataka (NITK), Surathkal, India B.Tech. in Computer Science and Engineering <i>Gold medalist (1st Rank).</i>	2013 – 2017 GPA: 9.54

RESEARCH EXPERIENCE

Facebook AI Research, Paris <i>PhD Summer Intern, advised by Hady Elsahar</i>	June 2024 – Present
★ Matryoshka Representation Learning for Structural Hierarchy <ul style="list-style-type: none">· Leveraging Matryoshka Representation Learning (Kusupati et al, 2024) to encode long-form text while preserving structurally hierarchical information and enabling compute flexibility. [work in progress]	
University of Washington <i>Visiting PhD student, advised by Prof. Ranjay Krishna</i>	January 2024 – Present
★ Hard Positives in Vision-Language Model Compositionality <ul style="list-style-type: none">· Revealed that existing studies of VL compositionality are lacking, as they only study label-altering perturbations to the model input (“hard negatives”), not label-retaining perturbations (“hard positives”).· Curated a benchmark of hard positives and showed that models perform poorly; further, methods to improve model performance on hard negatives exacerbate poor performance on hard positives.· Improved model performance on hard positives and hard negatives by generating both at scale and incorporating them into finetuning. [ECCV 2024]	
Allen Institute for Artificial Intelligence <i>PhD Summer Intern on Mosaic, advised by Jack Hessel, Khyathi Chandu, Jena Hwang</i>	June 2023 – September 2023
★ Reporting Bias in Vision-Language <ul style="list-style-type: none">· Studying how people caption images in web-scale image-text corpora, leveraging theories of cognitive science and pragmatics, and identifying how what people do and do not mention in captions impacts what models do and do not learn when trained on these captions, regardless of scale. [work in progress]	
University of California at Los Angeles <i>Graduate Research Assistant, advised by Prof. Kai-Wei Chang</i> <i>Work done in collaboration with Jack Hessel, Allen Institute for AI</i>	September 2022 – Present
★ Text encoders bottleneck compositionality in contrastive vision-language models <ul style="list-style-type: none">· Created a benchmark to quantify the amount of information lost by text encoders of vision-language models such as CLIP, that compress the text input into a single vector.· Created a benchmark to test vision-language models on various types of reasoning, and showed that the amount of information lost by the text encoder predicts overall model performance. [EMNLP 2023]	

★ **Investigating vision-language models’ struggle to understand spatial relations.**

- Constructed three benchmarks to evaluate spatial reasoning capabilities of vision-language models, testing model understanding of situations that are not supported by training priors, e.g., a mug *under* a table.
- Showed that popular vision-language pretraining corpora contain little reliable data for learning spatial relationships, and that basic modeling interventions like up-weighting and finetuning are not sufficient to address the challenges posed by our benchmarks. [EMNLP 2023]

Allen Institute for Artificial Intelligence

October 2020 – October 2022

Predoctoral Young Investigator, advised by Prof. Aniruddha Kembhavi

Work done in collaboration with Prof. Derek Hoiem, UIUC

★ **Webly Supervised Concept Expansion for General Purpose Vision Models**

- Curated 40K queries, obtained 1M image-query pairs from search engine data (for only \$150!) and converted them to a VQA-style dataset containing over 10K visual concepts (including tail concepts).
- Finetuned three general purpose vision-language models on this data, and studied how the models transferred the knowledge of new concepts to tasks such as captioning, localization, and action- and attribute-focused datasets. [ECCV 2022]

★ **Towards General Purpose Vision Systems**

- Co-created an end-to-end, task-agnostic vision-language model and an evaluation that tests model generality in terms of: (1) flexibility of architecture, (2) ability to transfer knowledge between tasks, and (3) ability to learn new tasks quickly.
- Conducted detailed analyses of the model performance on the new evaluation benchmark, identifying multiple kinds of knowledge transfer between tasks. [CVPR 2022, oral]

Stanford University

September 2018 – June 2020

Graduate Research Assistant, advised by Prof. Percy Liang

★ **Selective Question Answering under Domain Shift**

- Proposed a novel setting encapsulating the practical necessity of knowing when to abstain on test data that differs from the training data.
- Showed that QA models are overconfident on out-of-domain (OOD) examples relative to in-domain examples, leading to poor performance in mixed settings.
- Explored multiple selective prediction methods, showing that OOD data from even a non-target distribution can improve accuracy significantly when used to train a calibrator [ACL 2020].

★ **Evaluation and Improvement of Question Answering using Local Perturbations**

- Generated diverse local perturbations of QAs, using only a small number of crowdworker-labeled examples to classify these as meaning-altering or meaning-preserving.
- Showed that existing QA models had poor local decision boundaries, and that augmenting the training data with perturbations improved performance.

University of British Columbia

May 2016 – December 2016

Mitacs Summer Research Fellow, advised by Prof. Karthik Pattabiraman

★ **Dynamic Invariant Detection for Cyber-Physical System Security**

- Mined dynamic system properties of two cyber-physical systems, a smart electric meter and a smart insulin pump, using Association Rule Mining and related techniques.
- Co-created a dynamic Intrusion Detection System based on these invariants, reducing false positives by 31% and false negatives by 93% [FSE 2017].

Indian Institute of Science

May 2015 – December 2015

IAS Summer Research Fellow, advised by J. Lakshmi and Prof. S. K. Nandy

★ **Secure Virtualization on a Network-on-Chip**

- Completed a detailed security analysis of a novel Network-on-Chip (NoC) architecture for secure virtualization under a wide variety of virtual workloads.
- Suggested protocols to be followed by architecture components to prevent various NoC attacks.
- Upgraded the NoC design so it could be implemented on an FPGA.

WORK EXPERIENCE

Amazon AI

June 2020 – September 2020

Applied Science Intern, advised by Saab Mansour

- Created a test-bed for Intent Classification and Slot Labeling under two realistic forms of distribution shift: (1) from synthetic data to real data, and (2) between different data collection methods.
- Explored methods to improve performance of state-of-the-art models on this test-bed, including modified pre-training, modified training losses, and self-training.

Nutanix

June 2018 – September 2018

Software Engineering Intern, advised by Bala Neerumalla

- Built an orchestration service for Cloud Security. Proposed and initiated work on an Intrusion Detection System based on Association Rule Mining of system logs.

PUBLICATIONS

* denotes equal contribution

- [Amita Kamath](#), Cheng-Yu Hsieh, Kai-Wei Chang, Ranjay Krishna. **The Hard Positive Truth about Vision-Language Compositionality**, *European Conference on Computer Vision (ECCV)*, 2024.
- [Amita Kamath](#), Jack Hessel, Kai-Wei Chang. **Text encoders are performance bottlenecks in contrastive vision-language models**, *Empirical Methods in Natural Language Processing (EMNLP)*, 2023. URL: <https://aclanthology.org/2023.emnlp-main.301/>
- [Amita Kamath](#), Jack Hessel, Kai-Wei Chang. **What’s “up” with vision-language models? Investigating their struggle to understand spatial relations.**, *Empirical Methods in Natural Language Processing (EMNLP)*, 2023. URL: <https://aclanthology.org/2023.emnlp-main.568/>
- [Amita Kamath](#)*, Christopher Clark*, Tanmay Gupta*, Eric Kolve, Derek Hoiem, Aniruddha Kembhavi. **Webly Supervised Concept Expansion for General Purpose Vision Models**, *European Conference on Computer Vision (ECCV)*, 2022. URL: <https://arxiv.org/pdf/2202.02317.pdf>
- Tanmay Gupta, [Amita Kamath](#), Aniruddha Kembhavi, Derek Hoiem. **Towards General Purpose Vision Systems: An End-to-End, Task-Agnostic Vision-Language Architecture**, *Conference on Computer Vision and Pattern Recognition (CVPR) (oral)*, 2022. URL: <https://arxiv.org/pdf/2104.00743.pdf>
- [Amita Kamath](#), Robin Jia, Percy Liang. **Selective Question Answering under Domain Shift**, *Association of Computational Linguistics (ACL)*, 2020. URL: <https://aclanthology.org/2020.acl-main.503.pdf>
- [Amita Kamath](#). **Selective Prediction under Domain Shift for Question Answering**, Masters Thesis advised by Prof. Percy Liang and Prof. Christopher Manning (Research Distinction), 2020. URL: <https://amitakamath.github.io/AmitaKamathMSThesis.pdf>
- [Undergraduate] Maryam Aliabadi, [Amita Kamath](#), Julien Gascon-Samson, Karthik Pattabiraman. **ARTINALI: Dynamic Invariant Detection for Cyber-Physical System Security**, *ACM SIGSOFT Symposium on Foundations of Software Engineering (FSE)*, 2017. URL: <http://blogs.ubc.ca/karthik/files/2017/07/ARTINALI-FSE17.pdf>
- [Undergraduate] [Amita Kamath](#), Chirag Jamadagni, Kevin Mathew, Abhijith Anilkumar, Mohit Tahiliani. **GCPiN: Group Caching for Privacy in Named Data Networking**, *IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*, 2017.
- [Undergraduate] [Amita Kamath](#), Chirag Jamadagni, K Chandrasekaran. **VirtTorrent: BitTorrent for Inter-VM File Distribution**, *International Conference on IoT and Cloud Computing (ICC)*, 2016.
- [Undergraduate] Dhruv Chand, Sunil Nayak, K Bhat, Shivani Parikh, Yuvraj Singh, [Amita Kamath](#). **A mobile application for Women’s Safety: WoSApp**, *IEEE Region 10 Conference (TENCON)*, 2015.

TEACHING EXPERIENCE

- Head Teaching Assistant for Stanford CS221 (Artificial Intelligence), Summer 2019, with instructor Robin Jia. Managed a team of 6 TAs and 100+ students, taught sections, held office hours, and mentored several students.
- Teaching Assistant for Stanford CS224n (NLP with Deep Learning), Winter 2020, with Prof. Christopher Manning. Worked with a team of 23 TAs for 450+ students, taught sections, held office hours, and mentored 10+ student course project teams.
- Teaching Assistant for Stanford CS224n (NLP with Deep Learning), Winter 2019, with Prof. Christopher Manning. Worked with a team of 20 TAs for 400+ students to develop new assignments and re-implement existing assignment code in PyTorch. Held office hours and mentored 10+ student course project teams.
- Teaching Assistant for Stanford CS221 (Artificial Intelligence), Autumn 2018, with Prof. Percy Liang and Kelvin Guu. Worked with a team of 15 TAs for 400+ students to refine course assignments. Held office hours and mentored 10+ student course project teams.

PROFESSIONAL AND DEPARTMENT SERVICE

- Reviewer at ACL, EMNLP, CVPR, and NeurIPS Workshop on Distribution Shifts (DistShift).
- Student Member of the Stanford MSCS Admissions Committee, 2019. Reviewed student applications and discussed with faculty on the committee to select the incoming MS students.
- Organizer of the weekly Stanford NLP Reading Group, Autumn 2019, attended by graduate students and faculty of the NLP Group.

AWARDS

- Graduate Dean's Scholar Award, awarded by UCLA 2022-24.
- Computer Science Department First-Year Fellowship, awarded by UCLA 2022-23.
- Distinction in Research, awarded by the Stanford CS Department, 2020. For significant research conducted during my MS degree and the approval of a thesis based on the same.
- Top 3 Projects Award, Stanford CS224n (NLP with Deep Learning) 2018. My teammate and I studied adversarial attacks on Question Answering systems and devised methods to improve robustness against the same. Awarded 2nd Prize out of 145+ projects.
- Stanford-CISPA Research Fellowship, 2018. To perform work on Reinforcement Learning for web security.
- Grace Hopper Celebration Award Recipient, awarded by the Stanford CS Department, 2018.
- Mitacs Globalink Summer Research Fellowship, 2016. This fellowship is awarded by the Government of Canada to the top 5% of 10,000+ student applicants from developing countries based on academic performance and research potential, to perform research in top Canadian universities.
- Indian Academies of Science (IAS) Summer Research Fellowship, 2015. This fellowship is awarded to the top 5% of 25,000+ student applicants based on outstanding academic performance, to perform research in top Indian universities.
- Highest SGPA and CGPA awards, 2015-17. Awarded by the Department of Computer Science, NITK to the student with highest semester/cumulative GPA.
- Selected as one of 10 student volunteers from across the country for COMSNETS 2017, the top computer networks conference in India.
- INSPIRE, SPDC, DASA scholarships, 2010-2017. Awarded by the Government of India for outstanding academic performance.

RELEVANT COURSEWORK

Undergraduate	Data Structures and Algorithms, Design and Analysis of Algorithms, Continuous and Discrete Math, Discrete Mathematical Structures, Artificial Intelligence, Distributed Systems, Advanced Data Structures, Computer Networks
Graduate	Artificial Intelligence, Machine Learning, NLP with Deep Learning, CNNs for Visual Recognition, Deep Generative Models, Cryptography, Natural Language Understanding, Information Retrieval and Web Search, Mining Massive Datasets

SKILLS

Programming	Python, PyTorch, Tensorflow, C, C++ (proficient) Javascript, Java, MATLAB, HTML (basic)
Languages	English, Konkani (native) Hindi, Kannada, French (basic)
Extra-curricular	Social Dancing (performed in Stanford's The Nutcracker, 2019) Art (Creative Coordinator of Artists' Forum, NITK, 2017) Handicrafts (Guinness World Record [group] for largest crocheted blanket, 2016) MaRRS Spelling Bee (Second place in national and international levels, 2012) Creative Writing (published two short stories through Scholastic)