Shanmukha Ramakrishna Vedantam

19 Dutch Street, Apt 43A, New York, NY, USA

+1-650-407-7845 • vrama@gatech.edu • https://vrama91.github.io/ • U.S. Green Card Holder

Professional Summary

Research Scientist with 4 years of experience working in topics related to multimodal learning, out of distribution (OOD) generalization, generative modeling, and representation learning. Previously, Research Scientist at FAIR, and Ph.D. from Georgia Tech (Google Ph.D. fellowship awardee).

Work Experience

Visiting Researcher Nov 2022 - Present

New York University (NYU) Center for Data Science (CDS)

- Collaborating with Prof. Andrew Gordon Wilson and Prof. Julia Kempe on core problems in AI, and Probabilistic Machine Learning
- Leading other joint research projects collaborations with UMich, Stanford, and FAIR /Meta AI
- Research topics: Robustness, Multimodal AI and Data Efficient Learning
- Submission to ICML, and ICLR Workshops, planned submissions to NeurIPS

Research Scientist Feb 2019 - Nov 2022

Meta Inc., Fundamental AI Research (FAIR) Team

- Led a research agenda around understanding out-of-distribution (OOD) generalization, multimodal learning, and structured representation learning resulting in publications at NeurIPS, ICML, ICLR and IJCAI
- Managed two AI Residents, and three interns, and collaborated frequently with experts from neuroscience, physics, deep learning and cognitive science backgrounds

Education

- **Georgia Institute of Technology, United States** 2017-2018
  Ph.D, Computer Science
  Advisor: Prof. Devi Parikh
  Thesis: Interpretation, Grounding and Imagination for Machine Intelligence
- **Virginia Polytechnic Institute and State University, United States** 2013-2016
  M.S., Computer Engineering
  Advisor: Prof. Devi Parikh
  Specialization: Computer Vision and Machine Learning
- **International Institute of Information Technology (IIIT), Hyderabad, India** 2009-2013
  Bachelor of Technology, Electronics and Communication Engg.
  Advisor: Prof. K. Madhava Krishna
  Specialization: Vision for Robotics

Honors and Achievements

Research Accomplishments and Awards

1. 19173 citations on Google scholar with an H-index of 15 Jan. 2023
2. Awarded the Google PhD Fellowship in Machine Perception, Speech Technology and Computer Vision
   One out of 5 awardees selected across North America, Europe and the Middle East 2018
3. Awarded the ICLR travel award for attending the International Conference on Learning Representations 2018
4. Awarded travel grant of USD 1000 for CVPR, 2017 under Google's Archimedes program 2017

Reviewing

1. Outstanding reviewer award at ICLR
   2021

2. Outstanding reviewer award at ICCV
2019

3. Outstanding reviewer award at CVPR
Awarded to 130 reviewers in the CVPR reviewer pool 2017

Undergraduate and Previous

1. Best Discussion Participant Award, Advanced Computer Vision Course, Virginia Tech Spring, 2014
2. Mentioned in Dean's List for excellence in academics at IIIT Hyderabad
Monsoon, 2009 & 2011, Spring 2012
3. Winner of Judges award and Peer award at Siemens CTT Intern Tech Challenge 2012
4. 3rd in global aerospace competition CANSAT organized by NASA, AAS and AIAA 2011
5. Top 20 rank in Regional Mathematics Olympiad Organized by National Board for Higher Mathematics (NBHM) from Gujarat State (qualified for Indian National Mathematics Olympiad)
2008
6. Finalist for the Bal Shree honor, conferred by the President of India for outstanding creativity in Science 2008
7. Awarded Chacha Nehru Scholarship for Artistic and Innovative Excellence from National Council of Educational Research and Training (NCERT) 2008
8. Attained All India Rank 134 in National Science Olympiad 2006
9. All India Rank 13 in Indian National Cartography Association (INCA) Map Quiz 2006

Journal Publications

1. **Adopting Abstract Images for Semantic Scene Understanding.**
C. Lawrence Zitnick, Ramakrishna Vedantam, and Devi Parikh.
Special Issue on the best papers at the 2013 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 2016

2. **Grad-CAM: Why did you say that? Visual Explanations from Deep Networks via Gradient-based Localization.**
Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, Dhruv Batra.
International Journal of Computer Vision (IJCV), 2020 [11515 citations]

Conference Publications[1]

1. **Hyperbolic Image-Text Representations.**
Karan Desai, Maximilian Nickel, Tanmay Rajpurohit, Justin Johnson, **Ramakrishna Vedantam**.
In submission to ICML, 2023

2. **Improving Selective Visual Question Answering by Learning from Your Peers.**
Corentin Dancette, Spencer Whitehead, Rishabh Maheshwary, **Ramakrishna Vedantam**, Stefan Scherer, Xinlei Chen, Matthieu Cord, Marcus Rohrbach.
IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2023

3. **Dont forget the nullspace! Nullspace occupancy as a mechanism for out of distribution failure.**
Daksh Idnani, Vivek Madan, Naman Goyal, David J. Schwab, **Ramakrishna Vedantam**.
International Conference on Learning Representations (ICLR), 2023

[1]Publications in top AI conferences (CVPR, ICML, NeurIPS, ICCV) typically have 20-25% acceptance rates

4. **COAT: Measuring Object Compositionality in Emergent Representations.**
   Sirui Xie, Ari S Morcos, Song-Chun Zhu, **Ramakrishna Vedantam**.
   International Conference on Machine Learning (ICML), 2022 ***(Short Oral)***

5. **An Empirical Investigation of Domain Generalization in Empirical Risk Minimizers.**
   Ramakrishna Vedantam, David Lopez-Paz*, David Schwab*.
   Neural Information Processing Systems (NeurIPS), 2021

6. **CURI: A Benchmark for Productive Concept Learning Under Uncertainty.**
   Ramakrishna Vedantam, Arthur Szlam, Maximilian Nickel, Ari Morcos, Brenden Lake.
   International Conference on Machine Learning (ICML), 2021 ***(Short Oral)***

7. **Learning Optimal Representations with the Decodable Information Bottleneck.**
   Yann Dubois, Douwe Keila, David J. Schwab, **Ramakrishna Vedantam**.
   Neural Information Processing Systems (NeurIPS), 2020 ***(Spotlight)*** [Top 4%]

8. **IR-VIC: Unsupervised Discovery of Sub-goals for Transfer in RL.**
   Nirbhay Modhe, Prithvijit Chattopadhyay, Mohit Sharma, Abhishek Das, Devi Parikh, Dhruv Batra, **Ramakrishna Vedantam**.
   International Joint Conference on Artificial Intelligence (IJCAI), 2020 [Top 12.6%]
   Also presented at *ICLR Workshop on Task Agnostic Reinforcement Learning, 2019*

9. **Probabilistic Neural-Symbolic Models for Interpretable Visual Question Answering.**
   Ramakrishna Vedantam, Karan Desai, Stefan Lee, Marcus Rohrbach, Dhruv Batra, Devi Parikh.
   International Conference on Machine Learning (ICML), 2019 ***(Long Oral)*** [Top 4.2%]

10. **Generative Models of Visually Grounded Imagination.**
    Ramakrishna Vedantam, Ian Fischer, Jonathan Huang, Kevin Murphy.
    International Conference on Learning Representations (ICLR), 2018 [Top 10%, 113 citations]

11. **Grad-CAM: Why did you say that? Visual Explanations from Deep Networks via Gradient-based Localization.**
    Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, **Ramakrishna Vedantam**, Devi Parikh, Dhruv Batra.
    International Conference on Computer Vision (ICCV), 2017
    Also presented at *NIPS Workshop on Interpretable Machine Learning in Complex Systems, 2016*

12. **Sound-Word2Vec: Learning Word Representations Grounded in Sounds.**
    Ashwin K. Vijayakumar, **Ramakrishna Vedantam**, Devi Parikh.
    Conference on Empirical Methods in Natural Language Processing (EMNLP), 2017

13. **Counting Everyday Objects in Everyday Scenes.**
    Prithvijit Chattopadhyay*, **Ramakrishna Vedantam***, Ramprasaath R. Selvaraju, Dhruv Batra, Devi Parikh.
    IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017 ***(Spotlight)*** [Top 8.2%, 128 citations]

14. **Context-aware Captions from Context-agnostic Supervision.**
    Ramakrishna Vedantam, Samy Bengio, Kevin Murphy, Devi Parikh, Gal Chechik.
    IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017 ***(Spotlight)*** [Top 8.2%, 134 citations]
    Also presented as an Oral at the *Bay Area Machine Learning Symposium (BayLearn), 2017.*

15. **Visual Word2Vec (`vis-w2v`): Learning Visually Grounded Word Embeddings using Abstract Scenes.**
    Satwik Kottur*, **Ramakrishna Vedantam***, José Moura, and Devi Parikh.
    IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016 [106 citations]

16. **Learning Common Sense through Visual Abstraction.**
Ramakrishna Vedantam*, Xiao Lin*, Tanmay Batra, C. Lawrence Zitnick, and Devi Parikh.
IEEE International Conference on Computer Vision (ICCV), 2015
Also presented as an oral at *1st Workshop on Object Understanding for Interaction*, colocated with *ICCV, 2015*

17. **CIDEr: Consensus-based Image Description Evaluation.**
Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh.
IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015 [2983 citations]

* Equal Contribution

Workshops and arXiv Manuscripts

- **Hyperbolic Image-Text Representations.**
Karan Desai, Maximilian Nickel, Justin Johnson, **Ramakrishna Vedantam.**
ICLR Workshop on Multimodal Representation Learning (MRL), 2023

- **Understanding the class-specific effects of data augmentations.**
Polina Kirichenko, Randall Balestriero, Mark Ibrahim, **Ramakrishna Vedantam**, Hamed Firooz, Andrew Gordon Wilson.
ICLR Workshop on Trustworthy ML, 2023

- **Microsoft COCO Captions: Data Collection and Evaluation Server.**
Xinlei Chen, Hao Fang, Tsung-Yi Lin, **Ramakrishna Vedantam**, Saurabh Gupta, Piotr Dollar, C. Lawrence Zitnick.
arXiv:1504.00325 [1514 citations]

Professional Services

Reviewing:

- **Conference:** Reviewer for ICVGIP 2014 and 2018, ICCV 2015-2019, CVPR 2016-2020,2023, ECCV 2016-2020, ACCV 2016, ICVGIP 2016, BMVC 2017, NeurIPS 2017-2020,2022 ICLR 2018-2022, ICML 2018-2019, UAI 2018-2021
- **Journals:** Reviewer for International Journal of Computer Vision (IJCV), IEEE Transactions on Image Processing, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), Neural Computation, Computer Speech and Language, IEEE Open Journal of Signal Processing (2022)

Area Chair

- British Machine Vision Conference (BMVC), 2022

Past Research Internships

Research Internship Summer, 2018
Microsoft Research Cambridge, United Kingdom
- Generative Models for Concept Learning
- Supervisor: Nate Kushman, Matthew Johnson, and Sebastian Nowozin, Microsoft Research

Research Internship Summer, 2017
Facebook AI Research (FAIR), Menlo Park, CA
- Supervisor: Devi Parikh, Dhruv Batra and Marcus Rohrbach, Facebook AI Research/Georgia Tech

Research Internship Winter, 2017
Machine Perception Group, Google Research, Mountain View, CA
- Grounded latent variable generative models for images and semantics.
- Supervisor: Kevin Murphy, and Ian Fischer, Google Research

Research Internship Summer, 2016
Machine Perception Group, Google Research, Mountain View, CA
- Worked on a system to explain class discrimination conditioned on an image, using natural language

- Given an image, a target category and a distractor, explain why the image contains the target
- Supervisor: Gal Chechik, and Samy Bengio, Google Research and Google Brain

Research Internship Summer, 2014

Center for Visual Computing, Ecole Centrale de Paris/ INRIA - Saclay, France

- Worked on Loopy Part Models for Face Detection
- Supervisor: Prof. Iasonas Kokkinos, Ecole Centrale de Paris

Internship and AI Resident advising

- Karan Desai (Summer Intern), Facebook AI Research — Summer 2022
  Current position: Ph.D. student at U Michigan, Ann Arbor
- Daksh Idnani (AI Resident), Facebook AI Research — Fall 2021 - Spring 2022
  Current position: Stealth-mode Startup Founder
- Sirui Xie (Summer Intern), Facebook AI Research — Summer 2021
  Current position: Ph.D. student at University of California, Los Angeles
- Yann Dubois (AI Resident), Facebook AI Research — Fall 2019 - Spring 2020
  Next position: Ph.D. student at Stanford University
- Siddharth Ancha (Summer Intern), Facebook AI Research — Summer 2019
  Current position: Ph.D. student at Carnegie Mellon University
- Ananya Raval (MS Student), Georgia Institute of Technology — Fall 2017
  Current position: Software Development Engineer at Cisco
- Satwik Kottur (Intern), Virginia Tech (co-advised with Devi Parikh) — Summer 2015
  Current position: Research Scientist at Meta AI
- Prithvijit Chattopadhyay (Intern), Virginia Tech (co-advised with Devi Parikh) — Summer 2015
  Current position: Ph.D. student at Georgia Tech

Open Source Contributions

- Lead developer of the Domainbed Measures project opensourced by Facebook research
- Lead developer of the Productive Concept Learning project opensourced by Facebook research
- Lead developer of the Joint VAE project open sourced by Google
- Developer on the coco-caption project which implements commonly used image captioning metrics such as CIDEr, METEOR, BLEU, and ROUGE-L.
- Developer of the CIDEr project which implements the two versions of CIDEr (CIDEr and CIDEr-D) from our CVPR'15 paper on Consensus Based Image Description Evaluation.

Talks

CURI: A Benchmark For Productive Concept Learning Under Uncertainty

- AAAI Symposium on Conceptual Abstraction and Analogy in Natural and Artificial Intelligence. (November, 2020)

Concept Abstraction and Generalization for Machine Learning

- ConCats Seminar, New York University (October, 2020)
- Open Data Science Conference (ODSC) (March, 2022)

Learning Optimal Representations with the Decodable Information Bottleneck

- IIT Kanpur (May, 2021)

Connecting Vision and Language via. Interpretation, Grounding, and Imagination

- Courant Institute, New York University (May, 2019)
- University of Oxford (July, 2018)
- Google DeepMind, London (May, 2018)
- Microsoft Research, Cambridge (May, 2018)
- Facebook AI Research, Menlo Park (April, 2018)
- University of California, Berkeley (April, 2018)
- Google, Mountain View (May, 2018)

- Allen Institute for AI Research (AI2), Seattle (May, 2018)
- Toyota Technological Institute (TTI), Chicago (April, 2018)
- Indian Institute of Science (IISc), Bangalore (December, 2017)
- International Institute of Information Technology (IIIT), Hyderabad (December, 2017)

Panel Discussion

- Panelist on the "All things Attention" Workshop 2022 (co-located with NeurIPS, 2022).

Other Projects

- **Loopy Part Models for Face Detection** *INRIA-Saclay*
  Advisor: Iasonas Kokkinos and Dhruv Batra Summer, 2014
  Augmented the Deformable Parts Model (DPM) based face detector and landmark estimator with loopy part models. Utilized dual decomposition and an augmented lagrangian technique called ADMM (Alternating Direction Method of Multipliers) to solve the resulting inference problem efficiently, often achieving zero primal dual gap. Applied the model to get results comparable to the state of the art for detection and landmark localization
- **Understanding and Predicting Importance** *Virginia Tech*
  Advisor: Devi Parikh Spring, 2014
  Formulated importance prediction in abstract images as a structured prediction problem, where importance is defined as the likelihood of an object in an image being mentioned in a description. Incorporated task related insights into feature and model (structure) selection. Predicted importance of objects at upto 86 % accuracy on the Abstract-50S dataset
- **CanSat 2011** *IIIT Hyderabad*
  Advisor: K.S. Rajan Summer, 2010
  Designed, fabricated and launched into the lower space an autonomous micro-satellite carrying a large raw hen's egg intact - from *launch* to *landing*. Ground station set up to monitor the mini-satellite. Involved in CanSat testing, circuitry and sensor integration teams

Coursework

- **Graduate Coursework**: Computer Vision Systems, Advanced Computer Vision, Introduction to Machine Learning, Probabilistic Graphical Models, Independent Study - Deep Learning, Numerical Analysis and Software, Data Analytics-2, Convex Optimization, Deep Learning for Perception, Bayesian Statistics, Mathematical Foundations of Machine Learning, Computability and Algorithms
- **Selected Undergraduate Coursework**: Mobile Robotics, Artificial Neural Networks, Speech Signal Processing, Medical Image Processing, Engineering Systems, Data Structures, Operating Systems and Algorithms

Skills

- **Programming Languages**: Python, C++, Lua, Matlab
- **Libraries**: PyTorch, Tensorflow, Caffe, NLTK (Natural Language ToolKit)
- **Human Computation**: Amazon Mechanical Turk

Extra Curricular

- Volunteer for Women in Machine Learning (WiML) mentoring program, 2022
- Volunteered in organizing Mid-Atlantic Computer Vision (MACV) workshop at Virginia Tech
- Regular participation in Computer Vision and Machine Learning Reading Group at Virginia Tech
- Hosted all the Talks at Felicity - 2011, annual college fest of IIIT Hyderabad
- Coordinator and Founder- Entrepreneurship Cell at IIIT Hyderabad
- Class Representative for ECE Undergraduate batch
- Member, Students Parliament (Monsoon 2012 and Spring 2013)
- Campus Ambassador for Teach for India at IIIT (2011 to 2012)
- Trained in Carnatic Classical music for 7 years

References

- Devi Parikh (FAIR/ Georgia Tech)
- Leon Bottou (FAIR)
- Kevin Murphy (Google)