Shanmukha Ramakrishna Vedantam

Jersey City, NJ • Green Card Holder

LinkedIn • Google Scholar • +1-650-407-7845 • vrama@gatech.edu

Executive Summary

Ph.D., AI innovator, and scientist building humanlike AI, stronger deep learning foundations, and robust multimodal systems. With 20+ papers and 37K+ citations, my work has shaped the AI landscape over the past decade and continues to push the boundaries of intelligent machines.

Work Experience

AI Evangelist, Angel Investor, and Researcher Jan. 2024 - Present

Self-Employed, New York, NY

Researching cost-effective AI training to drive democratization, while bridging cutting-edge research with broader narratives on AI safety and accelerating last-mile deployment via. strategic investments

- Collaborating with leading research institutes, including Vector Institute, Princeton, and NUS
- Published a paper at ICLR 2024 on a scalable approach to dataset distillation
- Co-organized the Dataset Distillation Challenge at ECCV 2024 with colleagues from NUS and MIT
- Supporting innovation as an angel investor in Artificio AI and Sphere.app

Visiting Faculty Feb. 2023 - Dec. 2023

NYU Center for Data Science

Led research initiatives around multimodal AI, data efficient AI and distributionally robust AI

- Published papers at ICML 2023, CVPR 2023, and NeurIPS 2023
- Received the **Best Paper Award** at the ICLR 2023 Workshop on Multimodal Representation Learning (MRL)

Research Scientist Feb. 2019 - Feb. 2023

Meta Inc., Fundamental AI Research (FAIR), NYC

Formulated, owned and drove an ambitious research agenda on building provably robust AI that shares similar congitive underpinnings as humans

- Research led to six papers at top ML/AI venues like NeurIPS, ICML, and ICLR
- Mentored multiple interns and AI Residents while conceptualizing and executing on cross-functional collaborations with experts in NLP, CV, core machine learning, information theory and physics
- Released code, datasets and resources for the AI community on projects related to understanding OOD generalization, multimodal embeddings, and humanlike reasoning

Education

- **Georgia Tech, United States** 2013–2018
  Ph.D, Computer Science (transfer from Virginia Tech)
  Thesis: Interpretation, Grounding and Imagination for Machine Intelligence
  Committee: Devi Parikh (Advisor), Larry Zitnick, Kevin P. Murphy, Jacob Eisenstein, Dhruv Batra
- **Virginia Tech, 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. 37,406 citations with an H-index of 16 Nov. 2024
2. Best Paper Award, ICLR workshop on Multimodal Representation Learning (MRL) 2023

3. 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
4. ICLR travel award for attending the International Conference on Learning Representations 2018
5. Finalist for the Adobe Research Fellowship 2018
6. Finalist for the Adobe Research Fellowship 2016
7. 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

Conference Publications[1]

1. **Embarassingly Simple Dataset Distillation.**
   Yunzhen Feng, **Ramakrishna Vedantam**, Julia Kempe.
   International Conference on Learning Representations (ICLR), 2024
2. **Understanding the detrimental class-level effects of data augmentation.**
   Polina Kirichenko, Mark Ibrahim, Randall Balestriero, Diane Bouchacourt, **Ramakrishna Vedantam**, Hamed Firooz, Andrew Gordon Wilson.
   Neural Information Processing Systems (NeurIPS), 2023
3. **Hyperbolic Image-Text Representations.**
   Karan Desai, Maximilian Nickel, Tanmay Rajpurohit, Justin Johnson, **Ramakrishna Vedantam**.
   International Conference on Machine Learning (ICML), 2023
   Best Paper Award, Multimodal Representation Learning (MRL) Workshop colocated with ICLR, 2023
4. **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
5. **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
6. **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)***
7. **An Empirical Investigation of Domain Generalization in Empirical Risk Minimizers.**
   Ramakrishna Vedantam, David Lopez-Paz*, David J. Schwab*.
   Neural Information Processing Systems (NeurIPS), 2021
8. **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)***
9. **Learning Optimal Representations with the Decodable Information Bottleneck.**
   Yann Dubois, Douwe Kiela, David J. Schwab, **Ramakrishna Vedantam**.
   Neural Information Processing Systems (NeurIPS), 2020 ***(Spotlight)*** *[Top 4%]*

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

10. **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*

11. **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%, 102 citations]*

12. **Generative Models of Visually Grounded Imagination.**
Ramakrishna Vedantam, Ian Fischer, Jonathan Huang, Kevin P. Murphy.
International Conference on Learning Representations (ICLR), 2018 [Top 10%, 163 citations]

13. **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 [25,774 citations]
Also presented at *NIPS Workshop on Interpretable Machine Learning in Complex Systems, 2016*

14. **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

15. **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%, 187 citations]*

16. **Context-aware Captions from Context-agnostic Supervision.**
Ramakrishna Vedantam, Samy Bengio, Kevin P. Murphy, Devi Parikh, Gal Chechik.
IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017 ***(Spotlight)*** *[Top 8.2%, 169 citations]*
Also presented as an Oral at the *Bay Area Machine Learning Symposium (BayLearn), 2017.*

17. **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 [118 citations]

18. **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 [109 citations]
Also presented as an oral at *1^{st} Workshop on Object Understanding for Interaction*, colocated with *ICCV, 2015*

19. **CIDEr: Consensus-based Image Description Evaluation.**
Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh.
IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015 [5,202 citations]

* Equal Contribution

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

Workshops and arXiv Manuscripts

- **Embarassingly Simple Dataset Distillation.**
Yunzhen Feng, **Ramakrishna Vedantam**, Julia Kempe.
Workshop on Advancing Neural Network Training (WANT), NeurIPS, 2023
- **Hyperbolic Image-Text Representations.**
Karan Desai, Maximilian Nickel, Justin Johnson, **Ramakrishna Vedantam.**
ICLR Workshop on Multimodal Representation Learning (MRL), 2023 ***(Spotlight)***[Best Paper Award]
- **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 [2,754 citations]

Professional Services

Reviewing:

- **Conference:** Reviewer for ICCV 2015-2019,25 CVPR 2016-2020,2023-24,25, ECCV 2016-2020, ACCV 2016, BMVC 2017, NeurIPS 2017-2020,2022, ICLR 2018-2022, ICML 2018-2019, UAI 2018-2021,25
- **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

Chair

- Technical Chair, Workshop on Applications of Deep Learning (WALDA) 3.0, 2023

Workshop Organization

- The First Dataset Distillation Challenge, co-located with ECCV 2024

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 P. 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

- Yunzhen Feng (Co-advised with Julia Kempe), New York University (NYU) Summer-Fall 2023
  Current Position: Ph.D. student at NYU
- Polina Kirichenko (NYU-FAIR collaboration), Facebook AI Research Fall 2022
  Current Position: Research Scientist at Meta AI
- Karan Desai (Summer Intern), Facebook AI Research Summer 2022
  Current position: Founding Member of Technical Staff at WorldLabs Inc.
- Daksh Idnani (AI Resident), Facebook AI Research Fall 2021 - Spring 2022
  Current position: Founder and CTO, Sphere.app
- Sirui Xie (Summer Intern), Facebook AI Research Summer 2021
  Current position: Research Scientist at Google DeepMind
- Yann Dubois (AI Resident), Facebook AI Research Fall 2019 - Spring 2020
  Current position: Ph.D. student at Stanford University
- Siddharth Ancha (Summer Intern), Facebook AI Research Summer 2019
  Current position: Postdoc at Massachusetts Institute of Technology (MIT)
- Prithvijit Chattopadhyay (Summer Intern), Virginia Tech Summer 2015
  Current position: Research Scientist at NVIDIA Research
- Satwik Kottur (Winter Intern), Virginia Tech Fall 2014
  Current position: Research Scientist at Meta AI

Thesis Co-Supervision

- Ananya Raval (Co-advised with Devi Parikh) Georgia Institute of Technology (Gatech)
  Generation of Linux Commands Using Natural Language Descriptions (M.S. Thesis) May 2018

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

Tutorial on Multimodal Machine Learning
- NYU Center for Data Science (Andrew Gordon Wilson Lab) March 2023

Measuring and Structuring Generalizable Deep Representations
- NYU Center for Data Science Feb. 2023
- Department of Computer Science, Princeton University Feb. 2023
- Indian Institute of Science (IISC) May 2023
- Indian Institute of Technology (IIT), Delhi April 2023

- Indian Institute of Technology (IIT), Madras April 2023
- International Institute of Information Technology (IIIT), Hyderabad April 2023
- Indian Institute of Technology (IIT), Hyderabad May 2023
- Department of Computer Science, University of Michigan Sep. 2023
- Stevens Institute of Technology Feb. 2024

Concept Abstraction and Generalization for Machine Learning

- ConCats Seminar, New York University Oct. 2020
- Open Data Science Conference (ODSC), Boston March 2022
- AAAI Symposium on Conceptual Abstraction and Analogy Nov. 2020

Learning Optimal Representations with the Decodable Information Bottleneck

- Indian Institute of Technology (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 Dec. 2017
- International Institute of Information Technology (IIIT), Hyderabad Dec. 2017

Panel Discussion

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

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