# Olga Russakovsky

Associate Professor, Computer Science Department Associate Director, Laboratory for Artificial Intelligence Princeton University

http://cs.princeton.edu/~olgarus olgarus@princeton.edu

## Education and training

- **Postdoctoral fellow**, Robotics Institute, Carnegie Mellon University, September 2015 June 2017 Advisors: Profs. Abhinav Gupta and Deva Ramanan
- **Ph.D. in computer science**, Stanford University, graduated September 2015 Thesis: "Scaling Up Object Detection," advisor: Prof. Fei-Fei Li
- M.S. in computer science, distinction in research, Stanford University, graduated June 2007 Thesis: "Algorithms for Training Conditional Log-Linear Models," advisor: Prof. Serafim Batzoglou
- B.S. in mathematics with distinction, Stanford University, graduated April 2007

### Awards

- Best Paper Award (track 3) at the ECCV 2024 Dataset Distillation Workshop. Awarded for "Vision-Language Dataset Distillation" (with Xindi Wu, Byron Zhang and Zhiwei Deng)
- Best Paper Honorable Mention, CHI 2023.
  Awarded for "'Help me help the AI': Understanding how explainability can support human-AI interaction." (with Sunnie S. Y. Kim, Elizabeth A. Watkins, Ruth Fong, Andrés Monroy-Hernández)
- PAMI Young Researcher Award, 2022
  Awarded annually for distinguished research contribution in computer vision to 1-2 researchers within seven years of their PhD
- **NSF CAREER Award**, 2022 Title: "Overcoming bias in computer vision: Building fairer systems and training diverse leaders"
- Princeton's School of Engineering & Applied Sciences Howard B. Wentz, Jr. Junior Faculty Award, 2021 Awarded for research contributions and for co-founding the AI4ALL nonprofit
- Princeton Phi Beta Kappa Award for Excellence in Undergraduate Teaching, 2021 Annual award to two faculty members selected by the students inducted into Phi Beta Kappa
- AnitaB.org's Emerging Leader Abie Award in Honor of Denice Denton, 2020 Awarded once every 2 years for high-quality research and significant positive impact on diversity
- **CRA-WP Anita Borg Early Career Award**, 2020 Awarded for significant contributions and outreach in CS and/or engineering
- Becominghuman.ai's 100 Brilliant Women in AI Ethics, 2019
- Princeton School of Engineering and Applied Sciences Distinction in Teaching Awards Awarded for courses taught in Spring 2018, Fall 2018, Fall 2022
- MIT Technology Review's 35 Innovators Under 35 Award, 2017
- PAMI Everingham Prize, 2016

Awarded for "a series of datasets and challenges since 2010 that have had such impact on the computer vision field. ImageNet built on the Caltech101/256 datasets, increasing the number of images by orders of magnitude and enabling the development of new algorithms." (shared with Alex Berg, Jia Deng, Fei-Fei Li and Wei Liu.)

- Outstanding Reviewer Awards, CVPR 2015 and CVPR 2016
- Foreign Policy's 100 Leading Global Thinkers, 2015 Awarded for co-founding Stanford AI Laboratory's outreach program (shared with Fei-Fei Li.)
- MIT EECS Rising Star Award, 2013 Awarded annually to "about 40 outstanding EECS graduate and postdoctoral women"
- National Science Foundation (NSF) Graduate Research Fellowship, 2007-2010
- Computing Research Association (CRA) Undergraduate research Award finalist, 2007

### **Publications**

(\* = equal contribution; links included only for pre-prints; citations at https://scholar.google.com/citations?hl=en&user=TB50wW8AAAAJ)

### Pre-prints

- 1. R. Wang, H. Huang, Y. Zhu, O. <u>Russakovsky</u>, Y. Wu. The Silent Prompt: Initial Noise as Implicit Guidance for Goal-Driven Image Generation. Dec 2024. https://arxiv.org/abs/2412.05101.
- 2. Y. Pei, R. Wang, Y. Yang, Y. Zhu, O. <u>Russakovsky</u>, Y. Wu. SOWing Information: Cultivating Contextual Coherence with MLLMs in Image Generation. Nov 2024. https://arxiv.org/abs/2411.19182.
- 3. S. Wang\*, Y. Zhu\*, R. Wang, A. Dharmasiri, O. <u>Russakovsky</u>, Y. Wu. DETER: Detecting Edited Regions for Deterring Generative Manipulations. Dec 2023. https://arxiv.org/abs/2312.10539.
- 4. Y. Zhu, Y. Wu, Z. Deng, O. <u>Russakovsky</u> and Y. Yan. Unseen image synthesis with diffusion models. Oct 2023. https://arxiv.org/abs/2310.09213.
- V. V. Ramaswamy, S. S. Y. Kim, R. Fong and O. <u>Russakovsky</u>. UFO: A unified method for controlling Understandability and Faithfulness Objectives in concept-based explanations for CNNs. Mar 2023. https://arxiv.org/abs/2303.15632.

Journal articles and monographs (peer-reviewed)

- A. Wang, A. Hertzmann and O. <u>Russakovsky</u>. Evaluating AI Fairness: Benchmark Suites, Not Leaderboards. *Patterns*, vol. 5, issue 11, 2024.
- X. Wu, B. Zhang, Z. Deng and O. <u>Russakovsky</u>. Vision-Language Dataset Distillation. *Transactions on Machine Learning Research (TMLR)*, 2024. Best Paper Award (track 3) at the ECCV 2024 Dataset Distillation Workshop.
- Z. Epstein, A. Hertzmann, and the Investigators of Human Creativity (M. Akten, H. Farid, J. Fjeld, M. R. Frank, M. Groh, L. Herman, N. Leach, R. Mahari, A. Pentland, O. <u>Russakovsky</u>, H. Schroeder, A. Smith). Art and the science of generative AI. *Science Perspectives*, vol. 380, pgs 1110-1111, 2023.
- 4. O. Yalcinkaya Simsek, O. <u>Russakovsky</u> & P. Duygulu. Learning actionness from action/background discrimination. *Signal, Image and Video Processing*, vol 17, pgs 1599–1606, 2023.
- A. Wang, A. Liu, R. Zhang, A. Kleiman, L. Kim, D. Zhao, I. Shirai, A. Narayanan and O. <u>Russakovsky</u>. REVISE: A Tool for Measuring and Mitigating Bias in Visual Datasets. *International Journal of Computer Vision (IJCV)*, 2022.

- 6. S. S. Y. Kim, S. Zhang, N. Meister, O. <u>Russakovsky</u>. [Re] Don't Judge an Object by Its Context: Learning to Overcome Contextual Bias. *ReScience-C journal (through NeurIPS reproducibility challenge)*, 2021.
- 7. S. Yeung, O. <u>Russakovsky</u>, N. Jin, M. Andriluka, G. Mori and L. Fei-Fei. Every moment counts: dense detailed labeling of actions in complex videos. *International Journal of Computer Vision (IJCV)*, 2017.
- 8. A. Kovashka, O. <u>Russakovsky</u>, L. Fei-Fei and K. Grauman. Crowdsourcing in computer vision. *Foundations and Trends in Computer Graphics and Vision*, 10(3), 2016.
- O. <u>Russakovsky</u>\*, J. Deng\*, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. Berg and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3), 2015. Featured in MIT Tech Review. As of Dec 2024: 48,232 citations.

Conference articles (peer-reviewed and archival)

- X. Wu\*, D. Yu\*, Y. Huang\*, O. <u>Russakovsky</u> and S. Arora. ConceptMix: A Compositional Benchmark for Evaluating Text-to-Image Models. *Neural Information Processing Systems, Datasets&Benchmarks Track* (*NeurIPS*), 2024.
- 2. W. Yang, Y. Zhu, Z. Deng and <u>O. Russakovsky</u>. What is Dataset Distillation Learning? *International Conference on Machine Learning* (*ICML*), 2024.
- 3. A. Chen, I. Sucholutsky, O. <u>Russakovsky</u> and T. L. Griffiths. Analyzing the Roles of Language and Vision in Learning from Limited Data. *Annual Conference of the Cognitive Science Society* (*CogSci*), 2024.
- 4. W. Yang\*, B. Zhang\* and O. <u>Russakovsky</u>. ImageNet-OOD: Deciphering modern out-of-distribution detection algorithms. *International Conference on Learning Representations (ICLR)*, 2024.
- V. V. Ramaswamy, S. Y. Lin, D. Zhao, A. B. Adcock, L. van der Maaten, D. Ghadiyaram and O. <u>Russakovsky</u>. GeoDE: a Geographically Diverse Evaluation Dataset for Object Recognition. *Neural Information Processing Systems, Datasets&Benchmarks Track (NeurIPS)*, 2023.
- 6. Y. Zhu, Y. W, Z. Deng, O. <u>Russakovsky</u> and Y. Yan. Boundary Guided Learning-Free Semantic Control with Diffusion Models. *Neural Information Processing Systems (NeurIPS)*, 2023.
- 7. A. Wang and O. <u>Russakovsky</u>. Overcoming Bias in Pretrained Models by Manipulating the Finetuning Dataset. *International Conference on Computer Vision (ICCV)*, 2023. **Oral presentation.**
- 8. N. Meister\*, D. Zhao\*, A. Wang, V. V. Ramaswamy, R. Fong and O. <u>Russakovsky</u>. Gender artifacts in visual datasets. *International Conference on Computer Vision (ICCV)*, 2023.
- 9. N. Yoo and O. <u>Russakovsky</u>. Efficient, Self-Supervised Human Pose Estimation with Inductive Prior Tuning. *ICCV ROAD++ Workshop* (*ICCVW*), 2023.
- S. Rane, M. L. Nencheva, Z. Wang, C. Lew-Williams, O. <u>Russakovsky</u> and T. L. Griffiths. Predicting Word Learning in Children from the Performance of Computer Vision Systems. *Annual Conference of the Cognitive Science Society* (CogSci), 2023.
- S. S. Y. Kim, E. A. Watkins, O. <u>Russakovsky</u>, R. Fong and A. Monroy-Hernández. Humans, AI, and Context: Understanding End-Users' Trust in a Real-World Computer Vision Application. *Conference on Fairness, Accountability and Transparency (FAccT)*, 2023.
- V. V. Ramaswamy, S. S. Y. Kim, R. Fong and O. <u>Russakovsky</u>. Overlooked factors in concept-based explanations: Dataset choice, concept salience, and human capability. *Computer Vision and Pattern Recognition* (*CVPR*), 2023.

- S. S. Y. Kim, E. A. Watkins, O. <u>Russakovsky</u>, R. Fong and A. Monroy-Hernández. "Help me help the AI": Understanding how explainability can support human-AI interaction. *ACM CHI Conference on Human Factors in Computing Systems (CHI)*, 2023. Best Paper Honorable Mention.
- Z. Deng and O. <u>Russakovsky</u>. Remember the past: Distilling datasets into addressable memories for neural networks. *Neural Information Processing Systems (NeurIPS)*, 2022.
- 15. J. Chung, Y. Wu and O. <u>Russakovsky</u>. Enabling Detailed Action Recognition Evaluation Through Video Dataset Augmentation. *Neural Information Processing Systems, Datasets&Benchmarks Track (NeurIPS)*, 2022.
- 16. S. S. Y. Kim, N. Meister, V. V. Ramaswamy, R. Fong and O. <u>Russakovsky</u>. HIVE: Evaluating the Human Interpretability of Visual Explanations. *European Conference on Computer Vision (ECCV)*, 2022.
- 17. Z. Wang, Y. Wu, K. Narasimhan and O. <u>Russakovsky</u>. Multi-query video retrieval. *European Conference on Computer Vision (ECCV)*, 2022.
- M. Qu, Y. Wu, W. Liu, Q. Gong, X. Liang, O. <u>Russakovsky</u>, Y. Zhao and Y. Wei. SiRi: A Simple Selective Retraining Mechanism for Transformer-based Visual Grounding. *European Conference on Computer Vision* (ECCV), 2022.
- 19. K. Yang, J. Yao, L. Fei-Fei, J. Deng and O. <u>Russakovsky</u>. A Study of Face Obfuscation in ImageNet. *International Conference on Machine Learning (ICML)*, 2022. Featured in Wired.
- A. Wang, V. V. Ramaswamy and O. <u>Russakovsky</u>. Towards Intersectionality in Machine Learning: Including More Identities, Handling Underrepresentation, and Performing Evaluation. *Conference on Fairness, Accountability and Transparency (FAccT)*, 2022.
- C. Jimenez, O. <u>Russakovsky</u> and K. Narasimhan. CARETS: A Consistency And Robustness Evaluative Test Suite for VQA. *Association for Computational Linguistics (ACL)*, 2022. Oral presentation.
- D. Zhao, A. Wang and O. <u>Russakovsky</u>. Understanding and Evaluating Racial Biases in Image Captioning. International Conference on Computer Vision (ICCV), 2021.
- 23. A. Wang and O. <u>Russakovsky</u>. Directional Bias Amplification. *International Conference on Machine Learning* (*ICML*), 2021.
- V. V. Ramaswamy, S. S. Y. Kim and O. <u>Russakovsky</u>. Fair Attribute Classification through Latent Space De-Biasing. *Computer Vision and Pattern Recognition (CVPR)*, 2021.
- Z. Deng, K. Narasimhan, O. <u>Russakovsky</u>. Evolving Graphical Planner: Contextual Global Planning for Vision-and-Language Navigation. *Neural Information Processing Systems (NeurIPS)*, 2020.
- H. Law, Y. Teng, O. <u>Russakovsky</u> and J. Deng. CornerNet-Lite: Efficient Keypoint Based Object Detection. British Machine Vision Conference (BMVC), 2020.
- 27. A. Wang, A. Narayayan and O. <u>Russakovsky</u>. REVISE: A Tool for Measuring and Mitigating Bias in Visual Datasets. *European Conference on Computer Vision (ECCV)*, 2020. Spotlight presentation.
- Z. Wang, B. Feng, K. Narasimhan and O. <u>Russakovsky</u>. Towards Unique and Informative Captioning of Images. *European Conference on Computer Vision (ECCV)*, 2020.
- Z. Wang, K. Qinami, Y. Karakozis, K. Genova, P. Nair, K. Hata and O. <u>Russakovsky</u>. Towards Fairness in Visual Recognition: Effective Strategies for Bias Mitigation. *Computer Vision and Pattern Recognition* (CVPR), 2020.
- 30. F. Yu, Z. Deng, K. Narasimhan and O. <u>Russakovsky</u>. Take the Scenic Route: Improving Generalization in Vision-and-language Navigation. *CVPR Visual Learning with Limited Labels Workshop (CVPRW)*, 2020.

- 31. K. Yang, K. Qinami, L. Fei-Fei, J. Deng and O. <u>Russakovsky</u>. Towards Fairer Datasets: Filtering and Balancing the Distribution of the People Subtree in the ImageNet Hierarchy. *Conference on Fairness, Accountability and Transparency (FAT\*)*, 2020. Featured in Wired.
- 32. J. Peterson\*, R. Battleday\*, T. Griffiths and O. <u>Russakovsky</u>. Human Uncertainty Makes Classification More Robust. *International Conference on Computer Vision (ICCV)*, 2019.
- 33. K. Yang, O. <u>Russakovsky</u> and J. Deng. SpatialSense: An Adversarially Crowdsourced Benchmark for Spatial Relation Recognition. *International Conference on Computer Vision (ICCV)*, 2019.
- J. Wang, O. <u>Russakovsky</u> and D. Ramanan. The more you look, the more you see: towards general object understanding through recursive refinement. *Winter Conference on Applications in Computer Vision (WACV)*, 2018.
- 35. G. Sigurdsson, O. <u>Russakovsky</u> and A. Gupta. What Actions are Needed for Understanding Human Actions in Videos? *International Conference on Computer Vision (ICCV)*, 2017.
- 36. S. Ganju, O. <u>Russakovsky</u> and A. Gupta. What's in a question: using visual questions as a form of supervision. *Computer Vision and Pattern Recognition (CVPR)*, 2017. **Spotlight presentation.**
- 37. A. Dave, O. <u>Russakovsky</u> and D. Ramanan. Predictive-corrective networks for action detection. *Computer Vision and Pattern Recognition (CVPR)*, 2017.
- 38. S. Yeung, V. Ramanathan, O. <u>Russakovsky</u>, L. Shen, G. Mori and L. Fei-Fei. Learning to learn from noisy web videos. *Computer Vision and Pattern Recognition (CVPR)*, 2017.
- 39. G. Sigurdsson, O. <u>Russakovsky</u>, I. Laptev, A. Farhadi and A. Gupta. Much ado about time: exhaustive annotation of temporal data. *Human Computation and Crowdsourcing Conference (HCOMP)*, 2016.
- 40. A. Bearman, O. <u>Russakovsky</u>, V. Ferrari and L. Fei-Fei. What's the point: semantic segmentation with point supervision. *European Conference on Computer Vision (ECCV)*, 2016.
- 41. S. Yeung, O. <u>Russakovsky</u>, G. Mori and L. Fei-Fei. End-to-end Learning of Action Detection from Frame Glimpses in Videos. *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- M. Vachovsky\*, G. Wu\*, S. Chaturapruek, O. <u>Russakovsky</u>, R. Sommer and L. Fei-Fei. Towards More Gender Diversity in CS through an Artificial Intelligence Summer Program for High School Girls. *ACM Special Interest Group on Computer Science Education* (*SIGCSE*), 2016.
- 43. O. <u>Russakovsky</u>, L.-J. Li and L. Fei-Fei. Best of both worlds: human-machine collaboration for object annotation. *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- 44. D. Modolo, A. Vezhnevets, O. <u>Russakovsky</u> and V. Ferrari. Joint calibration of Ensemble of Exemplar SVMs. *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- 45. J. Deng, O. <u>Russakovsky</u>, J. Krause, M. Bernstein, A. C. Berg and L. Fei-Fei. Scalable multi-label annotation. *ACM Conference on Human Factors in Computing Systems (CHI)*, 2014.
- 46. O. <u>Russakovsky</u>, J. Deng, Z. Huang, A. C. Berg, L. Fei-Fei. Detecting avocados to zucchinis: what have we done and where are we going? *International Conference on Computer Vision (ICCV)*, 2013.
- 47. O. <u>Russakovsky</u>, Y. Lin, K. Yu, L. Fei-Fei. Object-centric spatial pooling for image classification. *European Conference on Computer Vision (ECCV)*, 2012. Best poster award, Google PhD summit.
- 48. O. <u>Russakovsky</u> and L. Fei-Fei. Attribute learning in large-scale datasets. *Parts and Attributes Workshop of European Conference on Computer Vision (ECCVW)*, 2010.
- 49. E. Klingbeil, B. Carpenter, O. <u>Russakovsky</u> and A. Y. Ng. Autonomous operation of novel elevators for robot navigation. *International Conference on Robotics Automation (ICRA)*, 2010.
- 50. O. <u>Russakovsky</u> and A. Y. Ng. A Steiner tree approach to efficient object detection. *Computer Vision and Pattern Recognition (CVPR)*, 2010

51. S. S. Gross, O. <u>Russakovsky</u>, C. B. Do and S. Batzoglou. Training Conditional Random Fields for maximum labelwise accuracy. *Neural Information Processing Systems (NeurIPS)*, 2007.

Other

- 1. M. Coleman, O. <u>Russakovsky</u>, C. Allen-Blanchette and Y. Zhu. Discrete Diffusion Reward Guidance Methods for Offline Reinforcement Learning. https://openreview.net/forum?id=s4cSgzGudq. *Poster at the Sampling and Optimization in Discrete Space Workshop at ICML*, 2023.
- 2. S. Sudhakar, V. U. Prabhu, O. <u>Russakovsky</u> and J. Hoffman. ICON<sup>2</sup>: Reliably Benchmarking Predictive Inequity in Object Detection. https://arxiv.org/abs/2306.04482. *Poster at the Secure and Safe Autonomous Driving (SSAD) Workshop at CVPR*, 2023.
- 3. V. V. Ramaswamy, S. S. Y. Kim, N. Meister, R. Fong and O. <u>Russakovsky</u>. ELUDE: Generating interpretable explanations via a decomposition into labelled and unlabelled features. https://arxiv.org/abs/2206.07690. Short version at the *Explainable Artificial Intelligence for Computer Vision (XAI4CV) CVPR workshop*, 2022.
- 4. E. Zhao, D.-A. Huang, H. Liu, Z. Yu, A. Liu, O. <u>Russakovsky</u> and A. Anandkumar. Scaling Fair Learning to Hundreds of Intersectional Groups. https://openreview.net/forum?id=yjxVspo7gXt
- 5. A. Mani, W. Hinthorn, N. Yoo and O. <u>Russakovsky</u>. Point and Ask: Incorporating Pointing into Visual Question Answering. https://arxiv.org/abs/2011.13681. *Extended abstract at the Visual Question Answering Workshop at CVPR*, 2021.
- 6. J. Stroud, R. McCaffrey, R. Mihalcea, J. Deng and O. <u>Russakovsky</u>. Compositional Temporal Visual Grounding of Natural Language Event Descriptions. https://arxiv.org/abs/1912.02256.
- 7. M. Essaidi, O. <u>Russakovsky</u> and S. M. Weinberg. Fairness in Online Advertisement via Symmetric Auctions. Poster session at the *Conference on Web and Internet Economics (WINE)*, 2019.
- 8. E. Davydov and O. <u>Russakovsky</u>. Book chapter: "Introduction to Computer Science." *A Bioinformatics Guide for Molecular Biologists*. CSH Press, 2014.
- 9. O. <u>Russakovsky</u>, Y. Lin, K. Yu and L. Fei-Fei. Object-centric spatial pooling for image classification. *US Patent* 20130129199 A1.
- 10. S. Gould, O. <u>Russakovsky</u>, I. Goodfellow, P. Baumstarck, A. Y. Ng and D. Koller. The STAIR Vision Library. http://ai.stanford.edu/~sgould/svl, 2010.

## Selected Media

- Charter, in partnership with Time Magazine. How to make AI more diverse. December 10, 2023.
- NBC Today Show. How the summer program AI4ALL is helping reshape the future. August 1, 2023.
- ACM communications. *Trouble at the source*. Vol. 64, No. 12. December, 2021.
- VentureBeat. *ImageNet creators find blurring faces for privacy has a "minimal impact on accuracy."* March 16, 2021
- Wired. Researchers Blur Faces That Launched a Thousand Algorithms. March 15, 2021.
- IFLScience. Why Artificial Intelligence Is Biased Against Women. March 6, 2020.
- Wired. AI Is Biased. Here's How Scientists Are Trying to Fix It. December 19, 2019.
- New York Times. Dealing with Bias in Artificial Intelligence. November 19, 2019.

- Education Week. A Summer Camp With a Long Plan: Keeping Bias Out of Artificial Intelligence. Aug 28, 2019.
- The Princeton Packet. Summer program gives high schoolers hands-on experience with Artificial Intelligence. August 20, 2018.
- The Atlantic. The Future of AI Depends on High-School Girls. May 23, 2018.
- Princeton Alumni Weekly. Making Smart Machines Fair. June 6, 2018.
- Education Week. AI4All extends the power of Artificial Intelligence to high school girls. March 1, 2018.
- MIT Technology Review. The AI world will listen to these women in 2018. January 9, 2018.
- Wired. Meet the high schooler shaking up Artificial Intelligence. October 26, 2017.
- Australian Broadcasting Corporation. Science Friction extra: AI, eyes, girls and guys. October 7, 2017.
- Forbes. China's Rise In The Global AI Race Emerges As It Takes Over The Final ImageNet Competition. July 31, 2017.
- Quartz. The data that transformed AI research—and possibly the world. July 26, 2017
- Kathy Davis. Girl Power in the World of AI. June 2, 2017.
- EdTech. Stanford University's Artificial Intelligence Summer Camp Expands the World of Computer Science. September 2, 2016.
- Invited opinion piece at MIT Technology Review. AI's Research Rut. August 23, 2016.
- Motherboard. Can AI Help Gender Diversity Help AI? April 19, 2016.
- Foreign Policy. 100 Leading Global Thinkers: For cracking the STEM ceiling. December 1, 2015.
- Wired. This Girls' Summer Camp Could Help Change the World of AI. August 31, 2015.
- New Scientist. Computers are learning to see the world like we do. October 29, 2014.
- MIT Technology Review. *The Revolutionary Technique That Quietly Changed Machine Vision Forever*. September 9, 2014.
- CBC Radio. Teaching computers to see. September 5, 2014.
- New York Times. Computer Eyesight Gets a Lot More Accuracy. August 18, 2014.

### Invited talks

2024

- 1. ECCV Explainable Computer Vision workshop. *Human-centered approaches to explainable computer vision*. (with student Sunnie S. Y. Kim). Sept 29, 2024.
- 2. ECCV FAILED workshop. Fairness in Visual Recognition: The importance of the human perspective in computer vision. Sept 29, 2024.
- 3. **MIT Guest Lecture in 6.8300/6.8301.** *Fairness in Visual Recognition: The importance of the human perspective in computer vision.* Apr 26, 2024.

2023

4. Washington University Computer Science and Engineering. *Building Trustworthy (and Trusted) Computer Vision*. December 4, 2023.

- 5. ICCV Towards the Next Generation of Computer Vision Datasets workshop. *Reflecting on next generation datasets: The evolving and critical role of humans.* October 3, 2023.
- 6. ICCV "Quo Vadis, Computer Vision?" workshop. *Future AI research: Data* ∪ *Models* ∪ *Humans*. October 2, 2023.
- 7. ICML Data-centric Machine Learning Research workshop. *Data-centric Machine Learning: Tackling social bias in computer vision datasets.* (with Dr. Vikram V. Ramaswamy). July 29, 2023.
- 8. ICML Challenges in Deployable Generative AI workshop. *Art, Science and Challenges of Generative AI.* (with student Ye Zhu). July 28, 2023.
- 9. NSF Athena AI Institute seminar series. Trustworthy (and trusted) computer vision. January 17, 2023.

2022

- 10. ECCV Out-of-distribution Generalization in Computer Vision workshop. *Trustworthy (and trusted) computer vision: The transparency story.* October 24, 2022.
- 11. ECCV Real-world Surveillance workshop. *Trustworthy (and trusted) computer vision: The fairness part of the story.* October 24, 2022.
- 12. Keynote at the ECCV Responsible Computer Vision workshop. *Trustworthy (and trusted) computer vision: The fairness and transparency story.* October 23, 2022.
- 13. ECCV Computer Perception workshop. *Trustworthy (and trusted) computer vision: The benchmarking story.* October 23, 2022.
- 14. ECCV Adversarial Robustness in the Real World workshop. *Trustworthy (and trusted) computer vision: The benchmarking story.* October 23, 2022.
- 15. ECCV Robust Vision Challenge workshop. *Trustworthy (and trusted) computer vision: The evaluation story.* October 23, 2022.
- **16.** Google Media Understanding Research Speaker Series. Fairness in Visual Recognition: Redesigning the Datasets, Improving the Models and Diversifying the AI Leadership. September 7, 2022.
- 17. Keynote panelist at the Fairness, Accountability and Transparency (FAccT) Conference. *Implementing Intersectionality in Algorithmic Fairness.* June 21, 2022.
- 18. **CVPR AI for content creation workshop.** *Fairness in the context of AI-created content.* (with students Vikram V. Ramaswamy and Dora Zhao). June 19, 2022.
- 19. Stanford guest lecture in CS231N. Fairness in Visual Recognition: Redesigning the Datasets, Improving the Models and Diversifying the AI Leadership. May 31, 2022.
- 20. University of Amsterdam Deep Vision Lectures. Fairness in Visual Recognition: Redesigning the Datasets, Improving the Models and Diversifying the AI Leadership. January 27, 2022.

2021

- 21. Plenary speaker at the Indian Conference on Computer Vision, Graphics, and Image Processing. Fairness in Visual Recognition: Redesigning the Datasets, Improving the Models and Diversifying the AI Leadership. December 19, 2021.
- 22. NeurIPS Data-centric AI workshop. Past and future of data-centric AI. December 14, 2021.
- 23. NeurIPS ImageNet workshop. *Fairness and privacy aspects of ImageNet*. (with student Kaiyu Yang). December 13, 2021.

- 24. NeurIPS Datasets and Benchmarks track. *Broadening our thinking: responsible data collection*. December 9, 2021.
- 25. Duke Guest Lecture in ECE 661. Fairness in Visual Recognition. November 4, 2021.
- 26. Symposium on AI challenges at UCLA. *A time for reckoning and reflection in computer vision (where are we with ethics?)* October 28, 2021.
- 27. Vector Institute Computer Vision symposium. Fairness in Visual Recognition: Redesigning the Datasets, Improving the Models and Diversifying the AI Leadership. October 27, 2021.
- 28. Amazon. Fairness in Visual Recognition: Redesigning the Datasets, Improving the Models and Diversifying the AI Leadership. September 28, 2021.
- 29. Michigan Institute for Data Science (MIDAS). Fairness in Visual Recognition: Redesigning the Datasets, Improving the Models and Diversifying the AI Leadership. September 13, 2021.
- 30. ICML workshop on Socially Responsible ML. *Revealing, quantifying, analyzing and mitigating bias in visual recognition.* July 27, 2021.
- 31. **CVPR Learning with Limited and Imperfect Data Workshop.** *Mitigating bias and privacy concerns in visual data.* (with students Angelina Wang, Vikram V. Ramswamy and Kaiyu Yang). June 20, 2021.
- 32. CVPR VQA workshop. *Models, metrics, tasks and fairness in vision and language.* (with postdoc Zhiwei Deng and students Zeyu Wang, Arjun Mani and Dora Zhao). June 19, 2021.
- 33. Keynote at the CVPR Women in Computer Vision workshop. *Perception, interaction and fairness: key components of visual recognition.* (with students Nobline Yoo, Angelina Wang and Sunnie S. Y. Kim). June 19, 2021.
- 34. Keynote at the CVPR LatinX in AI workshop. *Towards fairness and inclusion in computer vision*. (with student Vikram V. Ramaswamy). June 19, 2021.
- 35. **MIT Guest Lecture in 6.819/6.869.** Fairness in Visual Recognition: Redesigning the Datasets, Improving the Models and Diversifying the AI Leadership. May 5, 2021.
- 36. Columbia University Computer Vision Seminar. Fairness in visual recognition. April 6, 2021.
- 37. UChicago Center for Data and Computing (CDAC) Distinguished Speaker Series. Fairness in visual recognition. Jan 25, 2021.

2018-2020

- 38. ECCV Fair Face Recognition and Analysis Workshop. Fairness in visual recognition. Aug 28, 2020.
- 39. MIT Vision Seminar. Fairness in visual recognition. Aug 18, 2020.
- 40. Stanford Vision and Learning lab. Fairness in visual recognition. July 6, 2020.
- 41. CVPR's Seventh Workshop on Fine-Grained Visual Categorization. *Revealing and mitigating biases in visual datasets*. June 19, 2020.
- 42. CMU VASC seminar. Fairness in visual recognition. April 20, 2020.
- 43. TTI Chicago. Fairness in visual recognition. Oct 14, 2019.
- 44. **CVPR workshop on Bias Estimation in Face Analytics.** *Strategies for mitigating social bias in visual recognition.* June 17, 2019.
- 45. ICML workshop on Identifying and Understanding Deep Learning Phenomena. *Strategies for mitigating social bias in deep learning systems.* June 15, 2019.
- 46. University of Pennsylvania GRASP seminar. Computer vision meets fairness. April 19, 2019.
- 47. CVPR workshop on Vision with Biased or Scarce Data. Fairness in computer vision. June 22, 2018.

- 48. CVPR workshop on DeepVision. Fairness in computer vision. June 18, 2018.
- 49. Keynote at O'Reilly AI conference. AI will change the world. Who will change AI? May 2, 2018
- 50. O'Reilly AI conference. Five reasons why fairness is important and relevant in computer vision. May 2, 2018.
- 51. Cornell Tech. The Human Side of Computer Vision. Feb 23, 2018.
- 52. Applied ML days conference at EPFL. The Human Side of Computer Vision. Jan 29, 2018.

During postdoc (2015-2017)

- 53. CVPR workshop on Visual Understanding by Learning from Web Data. *Towards Web-scale Video Understanding*. July 26, 2017.
- 54. IBM Watson. The Human Side of Computer Vision. Sep 6, 2016.
- 55. University of Edinburgh. The Human Side of Computer Vision. Aug 5, 2016.
- 56. University of Oxford. The Human Side of Computer Vision. Aug 4, 2016.
- 57. Amazon Lab126. The Human Side of Computer Vision. July 12, 2016.
- 58. Princeton University CS colloquium. The Human Side of Computer Vision. April 14, 2016.
- 59. University of Michigan. The Human Side of Computer Vision. April 4, 2016.
- 60. University of Southern California. The Human Side of Computer Vision. March 23, 2016.
- 61. Facebook AI Research. The Human Side of Computer Vision. March 18, 2016.
- 62. TTI Chicago. The Human Side of Computer Vision. Feb 24, 2016.
- 63. University of Illinois Urbana-Champaign. The Human Side of Computer Vision. Feb 22, 2016.
- 64. Michigan State University. Scaling Up Object Detection. Jan 22, 2016.
- 65. Cornell University. Scaling Up Object Detection. Dec 11, 2015.
- 66. NeurIPS workshop Women in Machine Learning. What's the point: semantic segmentation with point supervision. Dec 7, 2015
- 67. University of Texas Austin. Scaling Up Object Detection. Dec 4, 2015.
- 68. University of Pittsburgh. Scaling Up Object Detection. Dec 2, 2015.
- 69. University of Washington. Scaling Up Object Detection. Nov 23, 2015
- 70. National Robotics Engineering Center. Scaling Up Object Detection. Nov 20, 2015.
- 71. University of Southern California. Scaling Up Object Detection. Nov 10, 2015.
- 72. Disney Research. Scaling Up Object Detection. Oct 14, 2015.
- 73. UC San Diego. Scaling Up Object Detection. Oct 2, 2015.
- 74. Caltech. Scaling Up Object Detection. Oct 1, 2015.

During PhD (2010-2015)

- 75. Dropbox. Scaling Up Object Detection. July 29, 2015.
- 76. Xerox PARC. Scaling Up Object Detection. July 28, 2015.
- 77. Simon Fraser University. Scaling Up Object Detection. July 24, 2015.
- 78. University of British Columbia. Scaling Up Object Detection. July 23, 2015.
- 79. NVIDIA. Scaling Up Object Detection. July 21, 2015.
- 80. Google. Scaling Up Object Detection. June 24, 2015.
- 81. CVPR workshop ChaLearn Looking at People. Best of Both Worlds: Human-Machine Collaboration for

Object Annotation. June 12, 2015.

- 82. Carnegie Mellon University VASC seminar. Designing and Overcoming Challenges in Large-Scale Object Detection. Mar 20, 2015.
- 83. NVIDIA GPU Technology Conference. *ImageNet Large Scale Visual Recognition Challenge*. (with Alexander Berg). Mar 19, 2015.
- 84. UC Irvine. Designing and Overcoming Challenges in Large-Scale Object Detection. Jan 15, 2015.
- 85. Baidu. ImageNet Large Scale Visual Recognition Challenge. Jan 7, 2015.
- 86. NeurIPS workshop on Challenges in Machine Learning. *ImageNet Large Scale Visual Recognition Challenge*. Dec 12, 2014.
- 87. Yahoo! Research Labs. ImageNet Large Scale Visual Recognition Challenge. Dec 8, 2014.
- 88. UC Berkeley. ImageNet Large Scale Visual Recognition Challenge. Nov 18, 2014.
- 89. Photo App Meetup. ImageNet Large Scale Visual Recognition Challenge. Sep 25, 2014.
- 90. Apple. ImageNet Large Scale Visual Recognition Challenge. Sep 18, 2014.
- 91. Australian National University. Analysis of Large Scale Visual Recognition. Nov 29, 2013.
- 92. Bay Area Vision Meeting. Analysis of Large Scale Visual Recognition. (with Fei-Fei Li). Oct 4, 2013.
- 93. UC Berkeley. Analysis of Large Scale Visual Recognition. June 5, 2013.
- AAAI Symposium on Weakly Supervised Learning from Multimedia. Object-Centric Spatial Pooling for Image Classification. Mar 25, 2013.
- 95. ECCV workshop on Parts and Attributes. Attribute Learning in Large-Scale Datasets. Sep 10, 2010.

Outreach talks

- 96. **Stanford University Math Camp.** July 2015, Aug 2016, July 2017, June 2020, July 2020, July 2021, July 2022, two in July 2023, July 2024
- 97. AI Helps Ukraine Charity Conference. December 8, 2022.
- 98. Stanford AI4ALL. June 2018, June 2019, July 2020, July 2021.
- 99. Boston University AI4ALL. Aug 2018.
- 100. Simon Fraser University AI4ALL. July 2018.
- 101.CMU AI4ALL. July 2018.
- 102. GirlCode summer camp. Aug 2015.

### **Employment history**

- Associate professor, Computer Science Department, Princeton University, 2023-now
  - Associate director, Princeton Laboratory for Artificial Intelligence, 2024-now
  - Affiliated faculty: Princeton Center for Statistics and Machine Learning (since 2017), Princeton Center for Information Technology Policy (since 2018), Princeton Dialogues on AI and Ethics (since 2019), Princeton Program in Cognitive Science (since 2021)
- Assistant professor, Computer Science Department, Princeton University, 2017-2023
- Postdoctoral fellow, Robotics Institute, Carnegie Mellon University, 2015-2017
- Research assistant with Prof. Fei-Fei Li, Stanford Vision lab, 2010-2015

- Research intern in the Media Analytics team, NEC Labs America, summers 2011-2013
- Research assistant with Prof. Andrew Ng, Stanford University, 2007-2010
- Residential counselor and teaching assistant, Stanford University Math Camp, summers 2007-2008
- Instructor, Educational Program for Gifted Youth middle school math course, summer 2007
- Undergraduate research assistant with Prof. Serafim Batzoglou, Stanford University, 2005-2007

### Service

Research community service

- Program chair, European Conference on Computer Vision (ECCV) 2024
- Steering committee member, Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track 2021-now
- Associate editor, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2020-2023
- Workshop chair, International Conference on Computer Vision (ICCV) 2021, Computer Vision and Pattern Recognition (CVPR) 2023
- Doctoral consortium chair, CVPR 2019
- Publicity and press chair, CVPR 2016
- Area chair, WACV 2016, NeurIPS 2019, CVPR 2018-23,25, ICCV 2021, ECCV 2022, NeurIPS D&B track 2022
- Workshop co-organizer:
  - "Workshop on responsible data" at CVPR 2024
  - "Future of Computer Vision Datasets" at CVPR 2021
  - "Responsible Computer Vision" at CVPR 2021
  - "International Challenge on Compositional and Multimodal Perception" at ECCV 2020
  - "Compositionality in Computer Vision" at CVPR 2020
  - "ImageNet Large Scale Visual Recognition Challenge" at ICCV 2013, ECCV 2014, ICCV 2015, ECCV 2016, CVPR 2017
  - "BigVision: International Workshop on Large Scale Visual Recognition and Retrieval" at CVPR 2015-16
  - "WiCV: Women in Computer Vision workshop" at CVPR 2015 (first year of the now-annual event)
- Tutorial co-organizer
  - "Algorithmic Fairness: Why it's hard, and why it's interesting" at CVPR 2022
  - "Human-centered AI for Computer Vision" at CVPR 2022
  - "Fairness in Computer Vision: Datasets, Algorithms, and Implications" at FAccT 2022
  - "ImageNet Large Scale Visual Recognition Challenge" at CVPR 2015
- **Reviewer** for journals: Proceedings of the National Academy of Science (PNAS), Transactions on Pattern Analysis and Machine Intelligence (TPAMI), Computer Vision and Image Understanding (CVIU), Pattern Recognition (PR)

• **Reviewer** for conferences: Computer Vision and Pattern Recognition (CVPR), European Conference on Computer Vision (ECCV), International Conference on Computer Vision (ICCV), Neural Information Processing Systems (NeurIPS), International Conference on Machine Learning (ICML), International Conference on Learning Representations (ICLR), Conference on Human Computation and Crowdsourcing (HCOMP), Conference on Computer Graphics and Interactive Techniques (SIGGRAPH), British Machine Vision Conference (BMVC). **Outstanding reviewer awards** at CVPR 2015 and 2016.

Other professional activities

- Participant, FBI Scientific Working Group on Artificial Intelligence, 2019-now
- Board member, Common Visual Data Foundation (CVDF), 2016-now
- Panelist, National Science Foundation (NSF) 2017, 2020, 2022
- Commentator, Columbia Law Review Symposium on Common law for the age of AI, 2019.

University service (except outreach)

- Member, ORFE Department Senior Hiring Search Committee, 2023-2024
- Senior faculty mentor to an Assistant Professor in Computer Science, 2023-now
- Member, Princeton Research Data Service Advisory Committee, 2022-2023
- Member, Executive Committee of the Robotics and Intelligent Systems Certificate Program, 2018-now
- Member, School of Engineering Metropolis Initiative Steering Committee, 2018-2023
- Co-founder and faculty mentor of the Princeton CS monthly Research Inclusion Social Event, 2017-2023
- Undergraduate advising, COS AB students, 2019-2023
- Member, Humanities Computing Curriculum Committee, 2020-2022
- Research area lead, COS PhD admissions and recruiting, 2018, 2019, 2021, 2022, 2025
- Member, Princeton CS Infrastructure Advisory Board, 2018-2019
- Member, School of Engineering Innovation Grant Proposal Review Committee, 2018
- Student member of the Stanford CS Department Faculty Search committee, 2015
- **Co-founder** of the Stanford Women in AI group with quarterly events, 2014-2015
- Founder of the weekly Stanford Computer Vision reading group, 2008-2014

#### Outreach K-12

• Co-founder and Board Member, AI4ALL nonprofit, 2016-now.

AI4ALL is a nonprofit working to increase diversity and inclusion in Artificial Intelligence. We create pipelines for underrepresented talent through education and mentorship programs. I became Chair of the Board in February of 2024. More at http://ai-4-all.org.

- Co-founder and co-director, Princeton AI4ALL outreach summer camp, 2018-now The camp teaches AI technology and policy to rising 11th graders from underrepresented groups, in a partnership between the Princeton Computer Science Department, the Princeton Center for Information Technology Policy, and the AI4ALL nonprofit. More at http://ai4all.princeton.edu.
- Co-founder and co-director, Stanford AI4ALL outreach summer camp, 2015-2017

The camp teaches AI to high school girls in a three-week technically rigorous curriculum. The camp was featured in Wired, we published a research study on its impact in SIGCSE 2016, and its success inspired the creation of the national AI4ALL nonprofit. More at http://ai4all.stanford.edu.

## Teaching

- COS 324: "Introduction to Machine Learning" (Spring 2022)
- COS 429: "Computer Vision" (Fall 2017, Fall 2019, Fall 2020, Fall 2021, Spring 2023, Fall 2024)
- COS 529: "Advanced Computer Vision" (Spring 2019)
- COS Graduate seminars: "Advanced Topics in CS: Visual Recognition" (Spring 2018, Fall 2022), "Advanced Topics in CS: Computer Vision Research Skills" (Fall 2018)
- COS Undergraduate Independent Work seminars: "AI Education" (Spring 2019), "Fairness in Visual Recognition" (Fall 2020, Spring 2021), "Computer Vision for Social Good" (Spring 2021, Spring 2025)
- Head teaching assistant; CS228: Probabilistic graphical models; Stanford University (Winter 2010)
- Head teaching assistant; CS221: Artificial Intelligence; Stanford University (Fall 2009)

## Advising and mentoring

Current students and postdocs

- Allison Chen, PhD student, Computer Science (recipient of the NSF Graduate Fellowship)
- Jihoon Chung, PhD student, Computer Science (recipient of the CS Outstanding Teaching Award)
- Amaya Dharmasiri, PhD student, Computer Science
- Max Gonzalez Saez-Diez, MS student, Computer Science
- Sunnie S. Y. Kim, PhD student, Computer Science (recipient of the NSF Graduate Fellowship and a Siebel Scholar)
- Xinran Liang, PhD student, Computer Science
- Xindi Wu, PhD student, Computer Science
- William Yang, PhD student, Computer Science
- Nobline Yoo, MS student, Computer Science
- Tyler Zhu, PhD student, Computer Science
- Ye Zhu, Postdoctoral fellow, Computer Science

#### Alumni: postdocs

- Dr. Yu Wu, Postdoctoral Scholar, 2021-2022. He became a Professor at Wuhan University.
- Dr. Zhiwei Deng, Postdoctoral Scholar, 2019-2022. He joined Google Research.

#### Alumni: PhD students

• **Dr. Angelina Wang,** PhD thesis *Operationalizing responsible machine learning: from equality towards equity,* Computer Science department, 2024. She was a **Siebel Scholar** and a recipient of the **NSF Graduate** 

**Fellowship**. She started her postdoc position at Stanford and will be an **Assistant Professor** at the Cornell Tech Information School starting fall 2025.

- **Dr. Vikram V. Ramaswamy**, PhD thesis *Tackling bias within Computer Vision Models*, Computer Science department, 2023. He joined the Princeton CS department as a **Teaching Faculty**.
- **Dr. Zeyu Wang,** PhD thesis *Tackling imperfections in data for building real-world computer vision systems*, Electrical Engineering department, 2022. He was awarded the **SEAS award for excellence**. He started his postdoc position at Stanford.

Alumni: Master's students

- **Dorothy Zhao**, MS research on fairness in computer vision, 2022. She joined the Sony AI Ethics team and then started her **CS PhD at Stanford** in 2023.
- Danqi Liao, MS thesis *Selective Feature Aggregation for Single Frame Supervised Temporal Action Localization*, 2022. She started her CS PhD at Yale in 2022.
- Felix Yu, early PhD work on video activity recognition and vision-and-language navigation, 2021. He joined Reverie Labs.

Alumni: Princeton undergraduate senior theses

- Matthew Coleman. (co-advised with Prof. Chrisine Allen-Blanchette.) Latent Diffusion Policies, 2023.
- Rohan Jinturkar. *Why Did The Model Fail? Empowering Non-Experts to Explore Computer Vision Models through Interactive Dialogue*, 2023. He won the Outstanding CS Senior Thesis prize and the Sigma Xi Book Award for Outstanding Undergraduate Research.
- Sing Yu (Phoebe) Lin. Cost Efficiency of Generating a Geodiverse Dataset for Object Classification, 2023.
- William Olson. Faces at Face Value: An Analysis of Face Recognition Technology Policy and Performance, 2023.
- Jayson Wu. Dress for Success: An Incremental Approach Towards Fashion Compatibility Featurization, 2023.
- Nobline Yoo. Efficient, template-based, self-supervised 2D human pose estimation, 2023.
- Zishuo (Byron) Zhang. *Learning in the Wild: Challenges With Out-of-Distribution Data*, 2023. He won the Outstanding CS Senior Thesis prize and the Sigma Xi Book Award for Outstanding Undergraduate Research.
- Nicole Meister. *Towards Interpretable and Fair Computer Vision*, 2022. She won the School of Engineering Calvin Dodd MacCracken Senior Thesis prize, the Sigma Xi Book Award for Outstanding Undergraduate Research, and the NSF Graduate Fellowship. She started her ELE PhD at Stanford in 2022.
- **Ryan Zhang.** *Redundancy Reduction in Vision Transformers,* 2022. He won the Outstanding CS Senior Thesis prize and the Sigma Xi Book Award for Outstanding Undergraduate Research.
- **Rishwanth Raghu**, *Analyzing and Improving Object Navigation in Indoor Environments*, 2022. He started his CS MS at Princeton in 2022.
- Henry Wang. Reconstructing Sound from Visual Piano Performances: An Overview, 2022.
- Arjun Mani. *Point and Ask: Incorporating Pointing Into Visual Question Answering*, 2021. He won the Outstanding CS Senior Thesis prize, Sigma Xi Book Award for Outstanding Undergraduate Research, NSF Graduate Fellowship, and was a finalist for Hertz Fellowship. He started CS PhD at Columbia in 2021.
- Sharon Zhang. *Contextual Bias and Interpretability in Image Classification*, 2021. She won the Middleton Miller'29 Prize for Independent Work. She started her CS PhD at Stanford in 2021.
- Dorothy Zhao. Understanding and Evaluating Racial Biases in Image Captioning, 2021. She won the

Outstanding CS Senior Thesis prize and the Sigma Xi Book Award for Outstanding Undergraduate Research. She started her CS MS at Princeton in 2021.

- Jessica Ho, Effects of Dataset Bias on Conditional Generative Adversarial Networks for Urban Scene Understanding, 2020.
- Gregory McCord, Evaluating Compositionality of Vision and Language Models, 2020.
- Emmanuel Teferi, An Exploratory Analysis of ML Models in CADx Design, 2020.
- Phillip Yoon, Improving Sound Separation and Localization Using Audio-Visual Scene Analysis, 2020.
- Andrew Zeng, Using Computer Vision to Model Fashion Outfit Compatibility, 2020.
- Berthy Feng, *Moving from Recognition to Reasoning in Image Captioning*, 2019. She won the Sigma Xi Book Award for Outstanding Undergraduate Research and started her CME PhD at Caltech in 2019.
- **Ioannis Karakozis**, *Fundamental Techniques for Bias Mitigation in Deep Visual Recognition*, 2019. He won the Sigma Xi Book Award for Outstanding Undergraduate Research.
- Ryan McCaffrey, Toward Zero-Shot Action Recognition for Video Moment Localization, 2019.
- **Rohan Doshi**, *Zero-Shot Semantic Segmentation*, 2018. He won the Sigma Xi Book Award for Outstanding Undergraduate Research.
- William Hinthorn, *Inferring Intent from Pointing with Computer Vision*, 2018. He won the Outstanding Computer Science Senior Thesis Prize.
- Prem Nair, An Exploration of Multi-class Multi-domain Image Classification, 2018.

Mentoring non-Princeton students

- Koome Murungi, Undergraduate student at Swarthmore (AI4ALL alum), summer 2024
- Lindsey Ehrlich, High school student, summer 2023.
- Kelly Lao, High school student, summer 2023.
- Ozge Yalcinkaya, PhD student in Computer Engineering at Hacettepe University, 2020-2021. Our joint work won the Alper Atalay Best Student Paper Award at SIU, 2022.
- Angelina Hasina Rajoelimbololona, African Masters in Machine Intelligence Rwanda, 2020-2021
- Noa Souccar, High school student (AI4ALL alumna), summer 2020
- Achal Dave, PhD student advised by Deva Ramanan at CMU, 2015-2017 (while I was a postdoc)
- Gunnar Sigurdsson, PhD student advised by Abhinav Gupta at CMU, 2016-2017 (while I was a postdoc)
- Jingyan Wang, PhD student advised by Deva Ramanan at CMU, 2015-2017 (while I was a postdoc)
- Serena Yeung, PhD student advised by Fei-Fei Li at Stanford, 2015-2017 (while I was a PhD student/postdoc)
- Siddha Ganju, Master's student at CMU, 2016-2017 (while I was a postdoc)
- Amy Bearman, Undergraduate student at Stanford, 2015-2016 (while I was a PhD student)
- Sean Ma, Master's student at Stanford, 2013-2014 (while I was a PhD student)