

# Ranjay Krishna

Stanford, CA - 94305  
✉ [ranjaykrishna@gmail.com](mailto:ranjaykrishna@gmail.com)  
🌐 [ranjaykrishna.com](http://ranjaykrishna.com)

## Education

- 2016–present **Ph.D. at Stanford University**, *Computer Science Department*, Palo Alto, CA.  
GPA 4.00 Doctor of philosophy: I enable machines that can expand their visual knowledge by interacting with and learning directly from people.
  - Co-advised by: Professor Fei-Fei Li and Professor Michael Bernstein
  - Focus areas: [computer vision](#), [machine learning](#), [human-computer interaction and natural language processing](#)
- 2014–2016 **M.Sc. at Stanford University**, *Computer Science Department*, Palo Alto, CA.  
GPA 3.98 Masters of science in computer science: with a [distinction in research](#)
  - Advised by: Professor Fei-Fei Li and Professor Michael Bernstein
  - Focus area: [artificial intelligence](#)
- 2009–2013 **B.Sc. at Cornell University**, *Electrical and Computer Engineering Department*, Ithaca, NY.  
GPA 3.85 Bachelors of science in [electrical engineering](#) with [Magna Cum Laude](#)
  - Focus areas: computer architecture and embedded systemsBachelors of science in [computer science](#) with [Magna Cum Laude](#)
  - Focus areas: compiler design and operating systems

## Teaching Experience

- 2020 **Instructor** Stanford CS 231N: Convolutional Neural Networks for Visual Recognition.
- 2019 **Instructor** Stanford CS 131: Computer Vision: Foundations and Applications.
- 2018 **Instructor** Stanford CS 131: Computer Vision: Foundations and Applications.
- 2017 **Instructor** Stanford CS 131: Computer Vision: Foundations and Applications.
- 2015 **Teaching Assistant** Stanford CS 131: Fundamentals of Computer Vision
- 2013 **Teaching Assistant** Cornell CS 3110: Functional Programming
- 2012 **Teaching Assistant** Cornell CS 3110: Functional Programming
- 2011 **Teaching Assistant** Cornell CS 3110: Functional Programming

## Honors and Awards

- 2019 NeurIPS Oral Paper for paper titled “HYPER: Human eYe Perceptual Evaluations of Generative Models”, awarded to top 0.5% 36 papers out of 7000 submissions
- 2019 Accell Innovation Scholar - chosen as 1 of 12 Stanford PhD Candidates in the School of Engineering to discover technology commercialization, opportunity evaluation and entrepreneurial leadership
- 2018 Brown Institute of Media Innovation Magic Grant - Awarded \$80,000 to study “Learning to Engage in Conversations to Train AI Systems”
- 2017 UIST Best Paper Honorable Mention award for paper titled “Crowd Research: Open Scalable University Laboratories”
- 2017 CVPR Spotlight award for paper titled “A Hierarchical Approach for Generating Descriptive Image Paragraphs”
- 2016 CVPR Oral award for paper titled “Visual Relationship Detection with Language Priors”, awarded to top 1% of submissions
- 2016 MIT PhD Fellowship - not accepted

- 2016 Christofer Stephenson Memorial Award for best Stanford CS Thesis for the thesis titled “Visual Genome: Crowdsourced Knowledge Representations”, awarded to 1 student at Stanford’s Computer Science Department
- 2015 Brown Institute for Media Innovation Grant - Awarded \$150,000 to research media innovations related to computer vision
- 2013 Winner of Cornell’s Electrical and Computer Engineering Department’s award for “Wireless Sign Language Translation” - Featured on Engadget, Sleashgear, Deaftech news, Element 14, The Economic Times, etc.
- 2010-2013 Cornell Dean’s list
- 2011 HKN: National Electrical and Computer Engineering Honor Society
- 2011 Tau Beta Pi, the Engineering Honor Society

## Peer-Reviewed Academic Publications

- CSCW 2020 **Conceptual Metaphors Impact Perceptions of Human-AI Collaboration.**  
Pranav Khadpe, **Ranjay Krishna**, Li Fei-Fei, Jeffrey Hancock, Michael Bernstein  
*ACM Conference on Computer-Supported Cooperative Work and Social Computing, 2020*
- CVPR 2020 **Action Genome: Actions with Composable Spatio-temporal Scene Graphs.**  
Jingwei Ji, **Ranjay Krishna**, Li Fei-Fei, Juan Carlos Niebles  
*IEEE conference on Computer Vision and Pattern Recognition, 2020*
- HComp 2019 **AI-based Request Augmentation to Increase Crowdsourcing Participation.**  
Junwon Park, **Ranjay Krishna**, Pranav Khadpe, Li Fei-Fei, Michael S. Bernstein  
*AAAI Conference on Human Computation and Crowdsourcing, 2019*
- NeurIPS 2019 **HYPE: Human eYe Perceptual Evaluation for Generative Models.**  
Sharon Zhou\*, Mitchell Gordon\*, **Ranjay Krishna**, Austin Narcomey, Li Fei-Fei, Michael S. Bernstein  
*Advances in neural information processing systems, 2019*  
[Oral award \(top 0.53%\)](#)
- ICCV 2019 **Scene Graph Prediction with Limited Labels.**  
Vincent Chen, Paroma Varma, **Ranjay Krishna**, Michael S. Bernstein, Christopher Re, Li Fei-Fei  
*IEEE International Conference on Computer Vision, 2019*
- ICCV 2019 **Visual Relationships as Functions: Enabling Few-Shot Learning.**  
Apoorva Dornadula, Austin Narcomey, **Ranjay Krishna**, Michael S. Bernstein, Li Fei-Fei  
*IEEE International Conference on Computer Vision: Scene Graph Representation and Learning workshop, 2019*
- CVPR 2019 **Information Maximizing Visual Question Generation.**  
**Ranjay Krishna**, Michael Bernstein, Li Fei-Fei  
*IEEE conference on Computer Vision and Pattern Recognition, 2019*
- CVPR 2018 **Referring Relationships.**  
**Ranjay Krishna\***, Ines Chami\*, Michael Bernstein, Li Fei-Fei  
*IEEE conference on Computer Vision and Pattern Recognition, 2018*
- ICCV 2017 **Dense-Captioning Events in Videos.**  
**Ranjay Krishna**, Kenji Hata, Frederic Ren, Li Fei-Fei, Juan Carlos Niebles  
*IEEE International Conference on Computer Vision, 2018*
- UIST 2017 **Crowd Research: Open and Scalable University Laboratories.**  
Rajan Vaish, Snehal Kumar Gaikwad, Geza Kovacs, Andreas Veit, **Ranjay Krishna**, Imanol Arrieta Ibarra, Camelia Simoiu, Michael Wilber, Serge Belongie, Sharad C. Goel, James Davis, Michael Bernstein  
*ACM conference on User Interface Software and Technology, 2017*  
[Best paper honorable mention award](#)
- CVPR 2017 **A Hierarchical Approach for Generating Descriptive Image Paragraphs.**  
Jonathan Krause, Justin Johnson, **Ranjay Krishna**, Li Fei-Fei  
*IEEE conference on Computer Vision and Pattern Recognition, 2017*  
[Spotlight award \(top 5%\)](#)

- CSCW 2017 **A Glimpse Far into the Future: Understanding Long-term Crowd Worker Accuracy.**  
Kenji Hata, **Ranjay Krishna**, Li Fei-Fei, Michael Bernstein  
*ACM Conference on Computer-Supported Cooperative Work and Social Computing, 2017*
- Stanford 2016 **Visual Genome: Crowdsourced Visual Knowledge Representations.**  
**Ranjay Krishna**  
[Christofer Stephenson Memorial Award for best Stanford CS Thesis](#)
- IJCV 2017 **Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations.**  
**Ranjay Krishna**, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li Jia-Li, David Ayman Shamma, Michael Bernstein, Li Fei-Fei  
*International Conference on Computer Vision, 2017*
- ECCV 2016 **Visual Relationship Detection with Language Priors.**  
Cewu Lu\*, **Ranjay Krishna\***, Michael Bernstein, Li Fei-Fei  
indicates equal contribution  
*European Conference on Computer Vision, 2016*  
[Oral award \(top 1%\)](#)
- CHI 2016 **Embracing Error to Enable Rapid Crowdsourcing.**  
**Ranjay Krishna**, Kenji Hata, Stephanie Chen, Joshua Kravitz, David Ayman Shamma, Li Fei-Fei, Michael Bernstein  
*ACM Conference on Human Computer Interaction, 2016*
- UIST 2015 **DAEMO: A Self-Governed Crowdsourcing Marketplace.**  
Stanford Crowdsourcing Collective  
*ACM Conference on User Interface Software and Technology, 2015*
- EMNLP 2015 **Generating Semantically Precise Scene Graphs from Textual Descriptions for Improved Image Retrieval.**  
Sebastian Schuster, **Ranjay Krishna**, Angel Chang, Li Fei-Fei and Christopher D. Manning  
*ACM conference on Empirical Methods in Natural Language Processing, Vision and Language Workshop, 2015*
- CVPR 2015 **Image Retrieval using Scene Graphs.**  
Justin Johnson, **Ranjay Krishna**, Michael Stark, Li-Jia Li, David Ayman Shamma, Michael Bernstein, Li Fei-Fei  
*IEEE conference on Computer Vision and Pattern Recognition, 2015*

---

## Manuscripts and Pre-prints

- ArXiv 2019 **Deep Bayesian Active Learning for Multiple Correct Outputs.**  
Khaled Jedoui, **Ranjay Krishna**, Michael S. Bernstein, Li Fei-Fei
- ArXiv 2019 **Open-Vocabulary Reinforcement Learning through Human Interaction.**  
**Ranjay Krishna**, Li Fei-Fei, Michael S. Bernstein
- ArXiv 2019 **Eevee: Transforming Image Content via High-Level User Goals.**  
Michelle Lam, **Ranjay Krishna**, Michael S. Bernstein
- ArXiv 2015 **SentenceRacer: A Game with a Purpose for Image Sentence Annotation.**  
Kenji Hata, Sherman Leung, **Ranjay Krishna**, Michael S. Bernstein, Li Fei-Fei

---

## Non-archival Publications

- UIST 2019 **Learning Social Strategies.**  
Junwon Park, **Ranjay Krishna**, Li Fei-Fei, Michael S. Bernstein
- CHI 2019 **Eevee: Transforming Images by Bridging High-level Goals and Low-level Edit Operations.**  
Michelle Lam, Gracie B. Young, Catherine Y. Xu, **Ranjay Krishna**, Michael Bernstein *ACM Conference on Human Computer Interaction, 2016*, Late-breaking work
- UIST 2018 **Engagement Learning: Generating AI Datasets by Engaging Online Participants.**  
**Ranjay Krishna\***, Donsuk Lee\*, Li Fei-Fei, Michael Bernstein *ACM User Interface Software and Technology Symposium, 2018*, Poster

---

## Research Experience

- 2014-present **Stanford Vision and Learning Group**, *Stanford, CA*.  
Advised by Professor Fei-Fei Li  
Extending visual representations through scene graphs
- 2014-present **Stanford Human Computer Interaction Group**, *Stanford, CA*.  
Advised by Professor Michael Bernstein  
Aligning A.I. goals with human interests
- 2017 **Google Machine Intelligence and Perception Group**, *Mountain View, CA*.  
Advised by Christian Szegedy  
Conducted research on Neural Program Synthesis
- 2017 **Stanford Graphics Group**, *Stanford, CA*.  
Advised by Professor Maneesh Agarwala  
Mapped visemes to phonemes
- 2016 **Facebook Artificial Intelligence Lab**, *New York City, NY*.  
Co-advised by Armand Joulin and Laurens Van der Maaten  
Studied visual relationships between objects in images
- 2014 **Cognition and Language Lab**, *Stanford, CA*.  
Advised by Professor Michael Frank  
Built a large dataset called Wordbank for learning childrens open vocabulary
- 2014-2016 **Yahoo Research**, *San Francisco, CA*.  
Advised by David Ayman Shamma  
Designed interfaces to speed up crowdsourcing by an order of magnitude
- 2010 **Integrated Circuits Design**, *Ithaca, NY*.  
Advised by Professor Alyosha Molnar  
Designed chips that are anti-symmetric angle sensitive through diffraction gratings

---

## Industry Experience

- 2014 **Maps Enterprise Team, Google Inc.**, *Mountain View, CA*.  
Data scientist intern  
Improved performance of VectorDB, the backend for Google Maps for Businesses
- 2013-2014 **Cloud and Kernel Teams, MongoDB Inc.**, *New York City, NY*.  
Software solutions architect  
Revamped the company website and added HR tools for sales and recruiting  
Implemented the C++ driver for MongoDB and designed and tested the redaction framework  
Integrated two-factor authentication for backups, automation and monitoring
- 2013 **Adwords Team, Google Inc.**, *Mountain View, CA*.  
Software engineering intern  
Created a searchable, analytically hierarchical model of the experiments conducted on adwords to optimize the quality and revenue and track these experiments  
Designed and developed an evaluation tool to study NGO's that combat societal issues in the United Kingdom and India through the Google Impact Challenge
- 2013 **Open Source Ruby on Rails Contributor**, *Ithaca, NY*.  
Cached AST's produced by SQL queries made to AREL  
Added simpler build and destroy functionality to Rake, a software management tool for both mysql and postgresql databases  
Abstracted away the type checking conducted by the Schema Dumper to individual databases to avoid unnecessary
- 2012 **Office Exchange Team, Microsoft Inc.**, *Redmond, WA*.  
Automated pre-build server validation for Microsoft Exchange from a manual process of 25 days to 6 minutes, saving 72% of fixed cost

---

## Non-Engineering Work Experience

- 2014-2016 **Stanford Venture Capital Group, Stanford, CA.**  
Consulted for General Catalyst and Red Point Ventures on potential investments in specific verticals
- 2013 **Influenza Tracking, New York City, NY.**  
Ran a freelance project to use e-prescriptions from patients to track the spread of influenza visually
- 2009-2012 **Cornell Desktop Support, Ithaca, NY.**  
Ghosted and processed images for all the staff and faculty computers and micros in Cornell
- 2010 **Department of Economic Affairs, Ministry of Finance, New Delhi, India.**  
Advised by Kaushik Basu, Chief Economist of the World Bank  
Investigated and built a model on inflation targeting applications and implementations

---

## Invited talks

- 2020 **Learning to Interact and Interacting to Learn.**  
◦ Snapchat research, Los Angeles, USA
- 2020 **Conceptual Metaphors Impact Perceptions of Human-AI Collaboration .**  
◦ Stanford Vision and Learning Lab, Stanford, CA
- 2020 **Scene Graphs as a Symbolic Visual Representation.**  
◦ Keynote talk at CVPR workshop on Diagram Image Retrieval and Analysis, Seattle, USA [link]
- 2020 **Compositionality in Computer Vision.**  
◦ CVPR workshop on compositionality, Seattle, USA [link]
- 2019 **Scene Graph Representation and Learning.**  
◦ ICCV workshop on Scene Graphs, Seoul, Korea [link]
- 2019 **What's new in Computer Vision?.**  
◦ SystemX Alliance Spring Workshop, Stanford, CA [link]
- 2019 **Learning to Engage in Conversations for AI Systems.**  
◦ Oval Seminar, Stanford University, CA [link]  
◦ Thomson Reuters, New York City, NY [link]
- 2018 **Artistic Computer Vision.**  
◦ Stanford Vision and Learning Lab, Stanford, CA  
◦ Stanford HCI reading group, Stanford, CA
- 2018 **Trust and Transparency in Artificial Intelligence.**  
◦ MediaX at Stanford, Stanford, CA [link]
- 2018 **The Building Blocks of Computer Vision.**  
◦ Stanford HCI workshop, Stanford, CA
- 2017 **Beyond Perception.**  
◦ Indian Institute of Technology, New Delhi, India  
◦ Indian Institute of Technology, Guwahati, India
- 2017 **Understanding Semantics.**  
◦ Sanskriti High School, New Delhi, India
- 2017 **Dense-Captioning Events in Video.**  
◦ IEEE Conference on Computer Vision and Pattern Recognition ActivityNet Challenge Workshop
- 2016 **Visual Relationship Detection with Language Priors.**  
◦ European conference on Computer Vision, Amsterdam, Netherlands  
◦ Stanford Vision Group, Stanford, CA  
◦ Stanford HCI Group, Stanford, CA
- 2016 **Embracing Error to Enable Rapid Crowdsourcing.**  
◦ ACM Conference on Human Computer Interaction, San Jose, CA  
◦ Stanford HCI Group, Stanford, CA

2016 **Visual Genome - Crowdsourced Visual Knowledge Representations.**

- Stanford Natural Language Processing Group, Stanford, CA
- Stanford Vision Group, Stanford, CA

---

## Leadership

- 2016-present Leading vision + language research sub-group at Stanford Vision and Learning Group
- 2016-present Leading Visual Genome research sub-group at the intersection of Computer Vision and Human-Computer Interaction
- 2014-2015 Elected Graduate Student Representative at Stanford Computer Science Department
- 2012-2013 Elected Second Vice President of Kappa Alpha Literary Society chapter at Cornell University
- 2011-2012 Elected Social Chair of Theta Tau, a professional engineering fraternity
- 2012-2013 Elected Officer at Association of Computer Science Undergraduates, Cornell Chapter of ACM
- 2010-2013 Elected Membership Coordinator for Aiesec - the worlds largest student organization

---

## Professional activities

### Workshop organization

- 2020 Co-organizer for “International Challenge on Compositional and Multimodal Perception” workshop at IEEE European Conference on Computer Vision 2020
- 2019-2020 Co-organizer and Guest Editor for an IEEE TPAMI special issue on “Graphs in Computer Vision”
- 2020 Co-organizer for “Compositionality in Computer Vision” workshop at IEEE Conference on Computer Vision and Pattern Recognition 2020 at Seattle, USA
- 2020 Co-organized “The ActivityNet Large Scale Activity Recognition Challenge” workshop at IEEE Conference on Computer Vision an Pattern Recognition 2020, at Seattle, USA
- 2019 Lead organizer for “Scene Graph Representation and Learning” workshop at IEEE International Conference on Computer Vision 2019 at Seoul, Korea
- 2018 Co-organized “The ActivityNet Large Scale Activity Recognition Challenge” workshop at IEEE Conference on Computer Vision an Pattern Recognition 2018, at Salt Lake City, USA
- 2017 Co-organized “The ActivityNet Large Scale Activity Recognition Challenge” workshop at IEEE Conference on Computer Vision an Pattern Recognition 2017, at Honolulu, USA
- 2017 Program committee for “Groupsight: Workshop on Human Computation for Image and Video Analysis” workshop at HCOMP 2017

### Academic reviewer

- 2016-present Reviewed papers from CVPR, CHI, UIST, CSCW, IJCV, ECCV, ICCV, TPAMI

### University service

- 2019-present Reviewed faculty applications as part of the Faculty Search Committee for Stanford’s Computer Science Department
- 2019-present Organize Stanford Bernstein group’s weekly talk series
- 2018-present Organize Stanford Vision and Learning group’s weekly talk series
- 2018-2019 Reviewed and helped select the next batch of Stanford’s Computer Science department’s PhD students as part of the Applications Committee
- 2015-2016 Organized Stanford Artificial Intelligence Entrepreneurship Club
- 2015 Organized logistics for Stanford’s AI outreach summer program

---

## Research mentoring and advising

### Current PhD students

- 2020-present Siddharth Karamcheti (Stanford University)
- 2018-present Khaled Jedoui (Stanford University)

### Current Masters students

- 2019-present Shubhang Desai (Stanford University)
- 2017-present Austin Narcomey (Stanford University)

### Current Undergraduate students

- 2020-present Madeleine Grunde-McLaughlin (Stanford University)
- 2020-present Helena Vasconcelos (Stanford University)
- 2020-present Zixian Ma (Stanford University)
- 2018-present Omer Gul (Stanford University)
- 2018-present Shubhang Desai (Stanford University)

### Past PhD students (placement after mentorship)

- 2018-2019 Mitchell Gordon (continued CS PhD at Stanford University)
- 2018-present Pranav Khadpe (continued CS PhD at Carnegie Mellon University)
- 2019-present Jingwei Ji (continued CS PhD at Stanford University)

### Past Masters students (placement after mentorship)

- 2017-2019 Junwon Park (placement: Microsoft as program manager)
- 2017-2019 Apoorva Dornadula (placement: co-founded startup viralspace.ai)
- 2017-2019 Vincent Chen (placement: co-founded startup)
- 2016-2018 Donsuk lee (placement: CS PhD at University of South California advised by Yan Liu)
- 2016-2018 Ines Chami (placement: CS PhD at Stanford University advised by Christopher Re)
- 2017-2017 Mohana Moorthy (placement: Uber as autonomy engineer)
- 2015-2017 Kenji Hata (placement: CS PhD at Princeton University advised by Olga Russakovsky)
- 2016-2016 Vincent Sitzmann (placement: EE PhD at Stanford University advised by Gordon Wetzstein)
- 2015-2016 Oliver Groth (placement: CS PhD at Oxford University advised by Andrea Vedaldi)
- 2016-2017 Frederic Ren (placement: Visa as software engineer)
- 2016-2016 Yutian Li (placement: Conscripton in South Korea as software engineer)

### Past Undergraduate students (placement after mentorship)

- 2017-2019 Michelle Lam (placement: CS PhD at Stanford University)
- 2016-2019 Sho Arora (placement: Marcari as machine learning engineer)
- 2017-2018 Jihyeon Janel Lee (placement: CS masters at Stanford University)
- 2018-2018 Daniel Cai (placement: Zoox as software engineer)
- 2018-2018 Buck Bukaty (continued undergraduate CS at Stanford University)
- 2016-2016 Gavin Mai (continued undergraduate CS at Stanford University)
- 2015-2015 Joshua Kravitz (continued undergraduate CS at Stanford University)
- 2015-2015 Stephanie Chen (continued undergraduate CS at Stanford University)
- 2015-2015 Sherman Leung (placement: CS masters at Stanford University)



---

## Press

- 2020 **Three ways Computer Vision is transforming marketing.**
  - Forbes - "Scene Graphs"
- 2018 **Engaging in Conversations to train AI systems.**
  - Stanford University - "Ranjay Krishna"
- 2018 **NOVA Wonders: Can we build a brain - Documentary.**
  - PBS - "Ranjay Krishna"
- 2016 **Brown Institute of Media Innovation.**
  - Columbia University - "Ranjay Krishna"
- 2016 **Visual Relationship Detection with Language Priors.**
  - RspVision - "Ranjay Krishna"
- 2013 **Sign Language Translation.**
  - YouTube - "Sign Language Translator - The Sound of Signing"
  - Enggaget - "Sign language translator turns gestures into spoken letters, makes for a better world"
  - Slashgear - "Sign Language Translator glove interprets gestures "
  - Economic Times of India - "power glove that translates sign language into spoken words"
  - The Tech Journal - "Sign language translator turns gestures into spoken letters"
  - Amrully - "Sign language translator turns gestures into spoken letters"
  - ZDNet - "Sign language translator turns gestures into spoken letters"