

CURRENT APPOINTMENT

2022-present **Assistant Professor at the University of Washington, Seattle, WA.**
Paul G. Allen School of Computer Science & Engineering

EDUCATION

- 2016–2021 **Ph.D. at Stanford University, Computer Science Department, Palo Alto, CA.**
GPA 4.00 Doctor of philosophy: with a **distinction in teaching**.
◦ Co-advised by: Professor Fei-Fei Li and Professor Michael Bernstein
- 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**
◦ Co-advised by: Professor Fei-Fei Li and Professor Michael Bernstein
- 2009-2013 **B.Sc. at Cornell University, Electrical and Computer Engineering Department, Ithaca, NY.**
GPA 3.85 Bachelors of science in **electrical & computer engineering** with **Magna Cum Laude**
Bachelors of science in **computer science** with **Magna Cum Laude**

PAST APPOINTMENTS

- 2023-2024 **Visiting Faculty at Allen Institute of Artificial Intelligence, Seattle, WA.**
PRIOR computer vision team
- 2021-2022 **Research Scientist at Facebook, Menlo Park, CA.**
Facebook Artificial Intelligence Research (FAIR)
- 2021-2022 **Affiliate Assistant Professor at the University of Washington, Seattle, WA.**
Paul G. Allen School of Computer Science & Engineering

TEACHING EXPERIENCE

Instructor

- 2024-now **CSE 455: Computer Vision at University of Washington.**
2024 Spring: taught 168 students with a course staff of 6 teaching assistants
- 2023-now **CSE 493G1: Deep Learning at University of Washington.**
2024 Winter: taught 110 students with a course staff of 6 teaching assistants
2023 Spring: taught 90 students with a course staff of 6 teaching assistants
- 2023-2023 **CSE 599H: Artificial Intelligence vs. Intelligence Augmentation at University of Washington.**
2023 Winter: taught 20 PhD students with a course staff of 1 teaching assistant
- 2020-2021 **CS 231N: Convolutional Neural Networks for Visual Recognition at Stanford University.**
Co-instructed with Professor Fei-Fei Li and Danfei Xu
2021 Spring: taught 464 students with a course staff of 17 teaching assistants
2020 Spring: taught 578 students with a course staff of 19 teaching assistants
- 2017-2019 **CS 131: Computer Vision: Foundations and Applications at Stanford University.**
Co-instructed with Professor Juan Carlos Nieves
2019 Fall: taught 154 students with a course staff of 4 teaching assistants
2018 Fall: taught 94 students with a course staff of 3 teaching assistants
2017 Fall: taught 70 students with a course staff of 3 teaching assistants

Teaching Assistant

- 2015 **CS 131: Fundamentals of Computer Vision at Stanford University.**
2015 Fall: co-instructed by Professor Fei-Fei Li and Dr. Juan Carlos

2011-2013 CS 3110: Functional Programming at Cornell University.

2013 Spring: instructed by Professor Benjamin Ylvisaker

2012 Fall: instructed by Professor Ramin Zabih

2012 Spring: instructed by Professor Nate Foster

2011 Fall: instructed by Professor Ramin Zabih

HONORS AND AWARDS

Papers recognized

- 2024 CVPR Oral award for paper titled “Visual Program Distillation: Distilling Tools and Programmatic Reasoning into Vision-Language Models”
- 2024 ICLR Spotlight award for paper titled “Selective Visual Representations Improve Convergence and Generalization for Embodied AI”, awarded to top (5%) 114 out of 2997 submissions.
- 2023 CSCW Best paper honorable mention award for paper titled “Explanations can Reduce Overreliance on AI Systems during Decision-Making”
- 2023 NeurIPS Oral award for paper titled “Quilt-1M: One Million Image-Text Pairs for Histopathology”, awarded to top (0.6%) 77 papers out of 12000 submissions
- 2023 NeurIPS Oral award for paper titled “DataComp: In search of the next generation of multimodal datasets”, awarded to top (0.6%) 77 papers out of 12000 submissions
- 2023 NeurIPS ReALML workshop best paper finalist for paper titled “Agile Modeling: From Concept to Classifier in Minutes”, awarded to top 2 papers at the workshop
- 2023 CVPR Highlight award for paper titled “CREPE: Can Vision-Language Foundation Models Reason Compositionally?”, awarded to top (2.5%) 235 papers out of 9155 submissions
- 2021 ACL Outstanding Paper award for paper titled “Mind Your Outliers! Investigating the Negative Impact of Outliers on Active Learning through the Lens of Visual Question Answering”, awarded to top (0.2%) 7 papers out of 3350 submissions
- 2020 CSCW Best Paper Honorable Mention award for paper titled “Conceptual Metaphors Impact Perceptions of Human-AI Collaboration”, awarded to top (2.2%) 22 papers out of 1000+ submissions
- 2019 NeurIPS Oral for paper titled “HYPE: Human eYe Perceptual Evaluations of Generative Models”, awarded to top (0.5%) 36 papers out of 7000 submissions
- 2019 EMNLP-WNUT workshop Oral for paper titled “Determining Question-Answer Plausibility in Crowdsourced Datasets Using Multi-Task Learning”, awarded to top (10%) 6 papers out of 60 submissions
- 2017 UIST Best Paper Honorable Mention award for paper titled “Crowd Research: Open Scalable University Laboratories”
- 2017 CVPR Spotlight for paper titled “A Hierarchical Approach for Generating Descriptive Image Paragraphs”, awarded to top 5% of submissions
- 2016 CVPR Oral for paper titled “Visual Relationship Detection with Language Priors”, awarded to top 1% of submissions

Awards with grants

- 2024 Sony Faculty Innovation Award
- 2018 Brown Institute of Media Innovation Magic Award - Granted \$80,000 to study “Learning to Engage in Conversations to Train AI Systems”
- 2015 Brown Institute for Media Innovation Award - Granted \$150,000 to research media innovations related to computer vision

General

- 2021 Distinction in Teaching for designing and instructing 5 courses during my Ph.D. at Stanford
- 2019 Accell Innovation Scholar - chosen as 1 of 12 Stanford Ph.D. Candidates in the School of Engineering to discover technology commercialization, opportunity evaluation and entrepreneurial leadership
- 2016 MIT Ph.D. 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
- 2013 Senior Project Winner in Cornell's Electrical and Computer Engineering Department for "Wireless Sign Language Translation" - Featured on Engadget, Sleashgear, Deaftech news, Element 14, The Economic Times, etc.
- 2010-2013 Cornell Dean's list: Spring 2010, Fall 2010, Spring 2011, Fall 2011, Spring 2012, Fall 2012, Spring 2013
- 2011 HKN: National Electrical and Computer Engineering Honor Society
- 2011 Tau Beta Pi, the Engineering Honor Society

PEER-REVIEWED ACADEMIC PUBLICATIONS

- CVPR 2024 **Iterated Learning Improves Compositionality in Large Vision-Language Models.**
Chenhao Zheng, Jieyu Zhang, Aniruddha Kembhavi, Ranjay Krishna
IEEE conference on Computer Vision and Pattern Recognition, 2024
- CVPR 2024 **Modeling Collaborator: Enabling Subjective Vision Classification With Minimal Human Effort via LLM Tool-Use.**
Imad Eddine Toubal, Aditya Avinash, Neil Gordon Alldrin, Jan Dlabal, Wenlei Zhou, Enming Luo, Otilia Stretcu, Hao Xiong, Chun-Ta Lu, Howard Zhou, **Ranjay Krishna**, Ariel Fuxman, Tom Duerig
IEEE conference on Computer Vision and Pattern Recognition, 2024
- CVPR 2024 **Holodeck: Language Guided Generation of 3D Embodied AI Environments.**
Yue Yang, Fan-Yun Sun, Luca Weihs, Eli VanderBilt, Alvaro Herrasti, Winson Han, Jiajun Wu, Nick Haber, **Ranjay Krishna**, Lingjie Liu, Chris Callison-Burch, Mark Yatskar, Aniruddha Kembhavi, Christopher Clark
IEEE conference on Computer Vision and Pattern Recognition, 2024
- CVPR 2024 **Quilt-LLaVA: Visual Instruction Tuning by Extracting Localized Narratives from Open-Source Histopathology Videos.**
Mehmet Saygin Seyfioglu, Wisdom O. Ikezogwo, Fatemeh Ghezloo, **Ranjay Krishna**, Linda Shapiro
IEEE conference on Computer Vision and Pattern Recognition, 2024
- CVPR 2024 **Imitating Shortest Paths in Simulation Enables Effective Navigation and Manipulation in the Real World.**
Kiana Ehsani, Tanmay Gupta, Rose Hendrix, Jordi Salvador, Luca Weihs, Kuo-Hao Zeng, Kunal Pratap Singh, Yejin Kim, Winson Han, Alvaro Herrasti, **Ranjay Krishna**, Dustin Schwenk, Eli VanderBilt, Aniruddha Kembhavi
IEEE conference on Computer Vision and Pattern Recognition, 2024
- CVPR 2024 **Visual Program Distillation: Distilling Tools and Programmatic Reasoning into Vision-Language Models.**
Yushi Hu, Otilia Stretcu, Chun-Ta Lu, Krishnamurthy Viswanathan, Kenji Hata, Enming Luo, **Ranjay Krishna**, Ariel Fuxman
IEEE conference on Computer Vision and Pattern Recognition, 2024
- ICLR 2024 **Selective Visual Representations Improve Convergence and Generalization for Embodied AI.**
Ainaz Eftekhari, Kuo-Hao Zeng, Jiafei Duan, Ali Farhadi, Ani Kembhavi, **Ranjay Krishna**
International Conference on Machine Learning, 2024
Spotlight Paper award

- ICLR 2024 **Davidsonian Scene Graph: Improving Reliability in Fine-grained Evaluation for Text-Image Generation.**
Jaemin Cho, Yushi Hu, Roopal Garg, Peter Anderson, **Ranjay Krishna**, Jason Baldridge, Mohit Bansal, Jordi Pont-Tuset, Su Wang
International Conference on Machine Learning, 2024
- VLDB 2024 **VOCALExplore: Pay-as-You-Go Video Data Exploration and Model Building.**
Maureem Daum, Enhao Zhang, Dong He, Brandon Hayes, **Ranjay Krishna**, Magdalena Balazinska
International Conference on Very Large Data Bases, 2024
- NeurIPS 2023 **SugarCrepes: Fixing Hackable Benchmarks for Vision-Language Compositionality.**
Cheng-Yu Hsieh, Jieyu Zhang, Zixian Ma, Aniruddha Kembhavi, **Ranjay Krishna**
Advances in neural information processing systems, 2023
- NeurIPS 2023 **OBJECT 3DIT: Language-guided 3D-aware Image Editing.**
Oscar Michel, Anand Bhattad, **Ranjay Krishna**, Tanmay Gupta, Aniruddha Kembhavi
Advances in neural information processing systems, 2023
- NeurIPS 2023 **Large Language Model as Attributed Training Data Generator: A Tale of Diversity and Bias.**
Yue Yu, Yuchen Zhuang, Jieyu Zhang, Yu Meng, Alexander Ratner, **Ranjay Krishna**, Jiaming Shen, Chao Zhang
Advances in neural information processing systems, 2023
- NeurIPS 2023 **Quilt-1M: One Million Image-Text Pairs for Histopathology .**
Wisdom Oluchi Ikezogwo, Mehmet Saygin Seyfioglu, Fatemeh Ghezloo, Dylan Stefan Chan Geva, Fatwir Sheikh Mohammed, Pavan Kumar Anand, **Ranjay Krishna**, Linda Shapiro
Advances in neural information processing systems, 2023
Oral Paper award
- NeurIPS 2023 **DataComp: In search of the next generation of multimodal datasets.**
Samir Yitzhak Gadre, Gabriel Ilharco, Alex Fang, Jonathan Hayase, Georgios Smyrnis, Thao Nguyen, Ryan Marten, Mitchell Wortsman, Dhruva Ghosh, Jieyu Zhang, Eyal Orgad, Rahim Entezari, Giannis Daras, Sarah Pratt, Vivek Ramanujan, Yonatan Bitton, Kalyani Marathe, Stephen Mussmann, Richard Vencu, Mehdi Cherti, **Ranjay Krishna**, Pang Wei Koh, Olga Saukh, Alexander Ratner, Shuran Song, Hannaneh Hajishirzi, Ali Farhadi, Romain Beaumont, Sewoong Oh, Alex Dimakis, Jenia Jitsev, Yair Carmon, Vaishal Shankar, Ludwig Schmidt
Advances in neural information processing systems, 2023
Oral Paper award
- NeurIPS 2023 **Cola: How to adapt vision-language models to Compose Objects Localized with Attributes?.**
Arijit Ray, Filip Radenovic, Abhimanyu Dubey, Bryan Plummer, **Ranjay Krishna**, Kate Saenko
Advances in neural information processing systems, 2023
- CoRL 2023 **AR2-D2: Training a Robot Without a Robot.**
Jiafei Duan, Yi Ru Wang, Mohit Shridhar, Dieter Fox, **Ranjay Krishna**
Conference on Robot Learning, 2023
- ICCV 2023 **TIFA: Text-to-Image Faithfulness Evaluation with Question Answering.**
Yushi Hu, Benlin Liu, Jungo Kasai, Yizhong Wang, Mari Ostendorf, **Ranjay Krishna**, Noah Smith
IEEE International Conference on Computer Vision, 2023
- ICCV 2023 **Agile Modeling: From Concept to Classifier in Minutes.**
Otilia Stretcu, Edward Vendrow, Kenji Hata, Krishnamurthy Viswanathan, Vittorio Ferrari, Sasan Tavakkol, Wenlei Zhou, Aditya Avinash, Emming Luo, Neil Alldrin, MohammadHossein Bateni, Gabriel Berger, Andrew Bunner, Chun-Ta Lu, Javier Rey, Giulia DeSalvo, **Ranjay Krishna**, Ariel Fuxman
IEEE International Conference on Computer Vision, 2023
- VLDB 2023 **EQUI-VOCAL: Synthesizing Queries for Compositional Video Events from Limited User Interactions.**
Enhao Zhang, Maureem Daum, Dong He, Brandon Hayes, **Ranjay Krishna**, Magdalena Balazinska
International Conference on Very Large Data Bases, 2023

- ACL 2023 **Distilling Step-by-Step! Outperforming Larger Language Models with Less Training Data and Smaller Model Sizes** .
Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alex Jason Ratner, **Ranjay Krishna**, Chen-Yu Lee, and Tomas Pfister
Annual Meeting of the Association for Computational Linguistics (ACL) Findings, 2023
- CVPR 2023 **CREPE: Can Vision-Language Foundation Models Reason Compositionally?** .
Zixian Ma*, Jerry Hong*, Mustafa Omer Gul*, Mona Gandhi, Irena Gao, **Ranjay Krishna**
IEEE conference on Computer Vision and Pattern Recognition, 2023
Highlight Paper award (awarded to top 2.5% of submission)
- CSCW 2023 **Explanations Can Reduce Overreliance on AI Systems during Decision-Making.**
Helena Vasconcelos, Matthew Joerke, Tobias Gerstenberg, Michael Bernstein, **Ranjay Krishna**
ACM Conference on Computer-Supported Cooperative Work and Social Computing, 2023
- NeurIPS 2022 **Alignment as a Multi-agent Intrinsic Reward.**
Zixian Sunnie Ma, Rose Wang, Li Fei-Fei, Michael Bernstein, **Ranjay Krishna**
Advances in neural information processing systems, 2022
- PNAS 2022 **Learning to Interact and Interacting to Learn with Socially Situated Artificial Intelligence.**
Ranjay Krishna, Donsuk Lee, Li Fei-Fei, Michael S. Bernstein
Proceedings of the National Academy of Sciences of the United States of America, 2022
- CVPR 2022 **Measuring Compositional Consistency for Video Question Answering.**
Mona Gandhi*, Mustafa Omer Gul*, Eva Prakash, Madeleine Grunde-McLaughlin, **Ranjay Krishna**, Maneesh Agrawala
IEEE conference on Computer Vision and Pattern Recognition, 2022
- CIDR 2022 **Interactive Video Data Cleaning and Exploration.**
Maureen Daum*, Enhao Zhang*, Dong He, Magdalena Balazinska, Brandon Haynes, **Ranjay Krishna**, Apryle Craig, Aaron Wirsing
Conference on Innovative Data Systems Research, 2021
- ACL 2021 **Mind Your Outliers! Investigating the Negative Impact of Outliers on Active Learning through the Lens of Visual Question Answering.**
Siddharth Karamcheti, **Ranjay Krishna**, Fei-Fei, Chris Manning
Proceedings of the Annual Meeting of the Association for Computational Linguistics, 2021
Outstanding Paper award (awarded to top 0.2%)
- CVPR 2021 **AGQA: A Benchmark for Compositional Spatio-Temporal Reasoning.**
Madeleine Grunde-McLaughlin, **Ranjay Krishna**, Maneesh Agrawala
IEEE conference on Computer Vision and Pattern Recognition, 2021
- 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
Best Paper Honorable Mention award (awarded to top 2.2%)
- 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
- HCOMP2019 **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 paper (awarded to 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**, Manol 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 paper (awarded to 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, awarded to 1 student every year
- 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 paper (awarded to 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

- 2024 **Diffusion prior for few-shot dense prediction.**
Benlin Liu, Luming Tang, Yongming Rao, Wei-Chiu Ma, **Ranjay Krishna**
(in submission)
- 2024 **Coarse Correspondence unleashes the emergent 3D understanding capability in GPT-4V.**
Benlin Liu, Yiqin Wang, Yuhao Dong, Yongming Rao, Yansong Tang, Wei-Chiu Ma, **Ranjay Krishna**
(in submission)
- ArXiv 2024 **m&m's: A Benchmark to Evaluate Tool-Use for multi-step multi-modal Tasks.**
Zixian Ma, Weikai Huang, Jieyu Zhang, Tanmay Gupta, **Ranjay Krishna**
(in submission)
- 2024 **Videoshop: Localized Semantic Video Editing with Noise-Extrapolated Diffusion Inversion.**
Xiang Fan, Anand Bhattad, **Ranjay Krishna**
(in submission)
- 2024 **The Hard Positive Truth about Vision-Language Compositionality .**
Amita Kamath, Cheng-Yu Hsieh, Kai-Wei Chang, **Ranjay Krishna**
(in submission)
- 2024 **NaturalBench: Evaluating Vision-Language Models on Natural Adversarial Samples.**
Wenxuan Peng, Baiqi Li, Zhiqiu Lin, Jean de Dieu Nyandwi, Zixian Ma, Simran Khanuja, Deva Ramanan, **Ranjay Krishna**, Graham Neubig
(in submission)
- 2024 **Efficient Inference of Vision Instruction-Following Models with Elastic Cache.**
Zuyan Liu*, Benlin Liu*, Jiahui Wang, Yuhao Dong, Guangyi Chen, Jiwen Lu, **Ranjay Krishna**, Yongming Rao
(in submission)
- 2024 **BLINK: Multimodal Large Language Models Can See but Not Perceive .**
Xingyu Fu*, Yushi Hu*, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah Smith, Wei-Chiu Ma, **Ranjay Krishna**
(in submission)
- 2024 **ImageInWords: Unlocking Hyper-Detailed Image Descriptions .**
Roopal Garg, Andrea Burns, Burcu Karagol Ayan, Yonatan Bitton, Ceslee Montgomery, Yasumasa Onoe, Andrew Bunner, **Ranjay Krishna**, Jason Baldridge, Radu Soricut
(in submission)
- ArXiv 2024 **SPARO: Selective Attention Improves Recognition, Compositionality, and Robustness.**
Ankit Vani, Bac Nguyen, Samuel Lavoie-Marchildon, Aaron Courville, **Ranjay Krishna**
(in submission)
- ArXiv 2024 **The Colosseum: A Benchmark for Evaluating Generalization for Robotic Manipulation.**
Wilbert Pumacay*, Ishika Singh*, Jiafei Duan*, **Ranjay Krishna**, Jesse Thomason, Dieter Fox
(in submission)
- ArXiv 2024 **Training Language Model Agents without Modifying Language Models .**
Shaokun Zhang, Jieyu Zhang, Jiale Liu, Linxin Song, Chi Wang, **Ranjay Krishna**, Qingyun Wu
(in submission)
- ArXiv 2024 **Cultural and Linguistic Diversity Improves Visual Representations.**
Andre Ye, Sebastin Santy, Jena D. Hwang, Amy X. Zhang, **Ranjay Krishna**
(in submission)
- ArXiv 2023 **Designing LLM Chains by Adapting Techniques from Crowdsourcing Workflows.**
Madeleine Grunde-McLaughlin, Michelle S. Lam, **Ranjay Krishna**, Daniel S. Weld, Jeffrey Heer
(in submission)
- ArXiv 2023 **DreamSync: Aligning Text-to-Image Generation with Image Understanding Feedback.**
Jiao Sun, Deqing Fu, Yushi Hu, Su Wang, Royi Rassin, Da-Cheng Juan, Dana Alon, Charles Herrmann, Sjoerd van Steenkiste, **Ranjay Krishna**, Cyrus Rashtchian
(in submission)

- ArXiv 2023 **Lasagna: Layered Score Distillation for Disentangled Object Relighting.**
Dina Bashkurova, Arijit Ray, Rupayan Mallick, Sarah Adel Bargal, Jianming Zhang, **Ranjay Krishna**, Kate Saenko
(in submission)
- ArXiv 2023 **Improving Interpersonal Communication by Simulating Audiences with Language Models.**
Ryan Liu, Howard Yen, Raja Marjeh, Thomas L. Griffiths, **Ranjay Krishna**
(in submission)
- ArXiv 2023 **EcoAssistant: Using LLM Assistant More Affordably and Accurately.**
Jieyu Zhang, **Ranjay Krishna**, Ahmed H. Awadallah, Chi Wang
(in submission)

BOOK CHAPTERS

- Daedalus 2022 **Searching for Computer Vision North Stars.**
Li Fei-Fei, **Ranjay Krishna**
Chapter of Daedalus Special issue on "AI & Society", 2022
- Springer 2021 **Visual Intelligence through Human Interaction.**
Ranjay Krishna, Mitchell Gordon, Li Fei-Fei, Michael Bernstein
Chapter of Artificial Intelligence for Human Computer Interaction: A Modern Approach, Springer 2021

NON-ARCHIVAL PUBLICATIONS

- CVPR 2024 **Retrieved Real Data Learns more Robust Representations than Synthetically Generated Data.**
Scott Geng, **Ranjay Krishna***, Pang Wei Koh*
IEEE conference on Computer Vision and Pattern Recognition - Workshop on Synthetic Data for Computer Vision
- CVPR 2023 **MIMIC: Masked Image Modeling with Image Correspondences.**
Kalyani Marathe*, Mahtab Bigverdi*, Nishat Khan, Tuhin Kundu, Aniruddha Kembhavi, Linda G. Shapiro, **Ranjay Krishna** *IEEE conference on Computer Vision and Pattern Recognition - Workshop for Learning 3D with Multi-View Supervision*
- WSDM 2024 **Scaling Up LLM Reviews for Google Ads Content Moderation.**
Wei Qiao, Tushar Dogra, Otilia Stretcu, Yu-Han Lyu, Tiantian Fang, Dongjin Kwon, Chun-Ta Lu, Enming Luo, Yuan Wang, Chih-Chun Chia, Ariel Fuxman, Fangzhou Wang, **Ranjay Krishna**, Mehmet Tek
ACM International Conference on Web Search and Data Mining, 2024
- ArXiv 2023 **Tool Documentation Enables Zero-Shot Tool-Usage with Large Language Models.**
Cheng-Yu Hsieh, Si-An Chen, Chun-Liang Li, Yasuhisa Fujii, Alexander Ratner, Chen-Yu Lee, **Ranjay Krishna**, Tomas Pfister
- ArXiv 2023 **LeetPrompt: Leveraging Collective Human Intelligence to Study Large Language Models.**
Sebastin Santy, Ayana Bharadwaj, Sahith Dambekodi, Alex Albert, Cathy Yuan, **Ranjay Krishna**
- ArXiv 2021 **On the Opportunities and Risks of Foundation Models.**
Center for Foundation Models at Stanford
- EMNLP 2020 **Determining Question-Answer Plausibility in Crowdsourced Datasets Using Multi-Task Learning.**
Rachel Gardner, Maya Varma, Clare Zhu, **Ranjay Krishna**. *The Fourth Workshop on Noisy User-generated Text at The 2020 Conference on Empirical Methods in Natural Language Processing*
Oral paper (awarded to top 10%)
- UIST 2019 **Learning Social Strategies.**
Junwon Park, **Ranjay Krishna**, Li Fei-Fei, Michael S. Bernstein. *ACM Symposium on User Interface Software and Technology, Late Breaking work.*

- ArXiv 2019 **Deep Bayesian Active Learning for Multiple Correct Outputs.**
Khaled Jedoui, **Ranjay Krishna**, Michael S. Bernstein, Li Fei-Fei
- 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
- CVPR 2018 **The ActivityNet Large-Scale Activity Recognition Challenge 2018 Summary.**
Bernard Ghanem, Juan Carlos Niebles, Cees Snoek, Fabian Caba Heilbron, Humam Alwassel, Victor Escorcia, **Ranjay Krishna**, Shyamal Buch, Cuong Duc Dao *IEEE conference on Computer Vision and Pattern Recognition - The ActivityNet Large-scale Activity Recognition Challenge Workshop, 2018*
- CVPR 2017 **ActivityNet Challenge 2017 Summary.**
Bernard Ghanem, Juan Carlos Niebles, Cees Snoek, Fabian Caba Heilbron, Humam Alwassel, Victor Escorcia, **Ranjay Krishna**, Shyamal Buch, Cuong Duc Dao *IEEE conference on Computer Vision and Pattern Recognition - The ActivityNet Large-scale Activity Recognition Challenge Workshop, 2017*
- ArXiv 2015 **SentenceRacer: A Game with a Purpose for Image Sentence Annotation.**
Kenji Hata, Sherman Leung, **Ranjay Krishna**, Michael S. Bernstein, Li Fei-Fei

RESEARCH WORK EXPERIENCE

- 2021-2022 **Meta Fundamental Artificial Intelligence Research (FAIR), Menlo Park, CA.**
Developing multi-modal correctable models
- 2017 **Google Machine Intelligence and Perception Group, Mountain View, CA.**
Advised by Dr. Christian Szegedy
Conducted research on Neural Program Synthesis
- 2017 **Stanford Graphics Group, Stanford, CA.**
Advised by Professor Maneesh Agarwala
Spatio-temporal scene graphs
- 2016 **Facebook Artificial Intelligence Lab, New York City, NY.**
Co-advised by Dr. Armand Joulin and Dr. 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 Dr. 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

ENGINEERING WORK 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 an 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

- 2024 **Evaluating Foundation models.**
 ◦ CVPR 2024 workshop on Evaluation for Generative Foundation Models [link]
- 2023-2024 **Vision-Language Compositionality.**
 ◦ CVPR 2024 workshop on Computer Vision with Humans in the Loop [link]
 ◦ DUB Seminar 2024 at the University of Washington's [link]
 ◦ ICCV 2023 workshop on On Closing The Loop Between Vision And Language [link]
 ◦ CVPR 2023 workshop on New Frontiers in Vision and Language Reasoning [link]
- 2023 **Embodied Intelligence.**
 ◦ AAAI 2023 Inaugural Summer Symposium on Embodied Intelligence [link]
- 2023 **Visual Intelligence from Human Interactions.**
 ◦ Microsoft Research, Seattle, USA
- 2022 **Self-supervision for 3D Dense Visual Representations.**
 ◦ Amazon Science Hub, Seattle, USA
- 2022 **Video Organization and Interactive Analytics.**
 ◦ Cisco, Remote talk
- 2021-2023 **Visual Intelligence through Human Cognition.**
 ◦ ICCV 2023 workshop on Scene Graphs and Graph Representation Learning (SG2RL) [link]
 ◦ CS 520 seminar talk, Stanford, USA
- 2021 **Visual Intelligence through Human Interaction.**
 ◦ ICML 2021 workshop on Human-in-the-loop learning, Virtual
 ◦ SKC Science & Technology Webinar Series, Virtual

- 2021 **Visual Intelligence through Human Learning.**
- Yale University, New Haven, USA
 - University of Toronto, Toronto, Canada
 - Samsung, Toronto, Canada
 - Robotics Institute, Carnegie Mellon University, Pittsburgh, USA
 - Stanford University, Stanford, USA
 - University of Southern California, Los Angeles, USA
 - University of Washington, Seattle, USA
 - Cornell University, Ithaca, USA
- 2020 **Learning to Interact and Interacting to Learn.**
- Vision group, University of Austin, Texas, USA
 - Graphics group, Stanford University, Stanford, USA
 - Princeton University, Princeton, USA
 - Snap 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

- 2022-present Principal Investigator of my own research group at the University of Washington
- 2022-present Co-director of RAIVN research group at the University of Washington
- 2016-2021 Led my own research group at Stanford at the intersection of Computer Vision and Human-Computer Interaction
- 2016-2021 Lead vision + language research sub-group at Stanford Vision and Learning Group
- 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

International service

- 2024 Local co-chair for IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2024
- 2021-2023 Guest Editor for Special Issue on Graph Learning for Computer Vision at Transactions at Pattern Analysis and Machine Intelligence (TPAMI)

Workshop organization

- 2024 Co-organized "Synthetic Data for Computer Vision" workshop at IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2024
- 2024 Co-organized "Compositionality in Computer Vision" workshop at IEEE European Conference on Computer Vision (ECCV) 2024
- 2023 Co-organized "Artificial Intelligence and Human-Computer Interaction" workshop at IEEE International Conference on Machine Learning (ICML) 2023
- 2023 Co-organized "Compositionality in Computer Vision" workshop at IEEE International Conference on Computer Vision (ICCV) 2023
- 2022 Co-organized "Artificial Intelligence and Human-Computer Interaction" workshop at IEEE International Conference on Machine Learning (ICML) 2023
- 2022 Co-organized "Compositionality in Computer Vision" workshop at IEEE International Conference on Computer Vision (ECCV) 2022
- 2021 Co-organized "Compositionality in Computer Vision" workshop at IEEE International Conference on Computer Vision (ICCV) 2021
- 2020 Co-organized "International Challenge on Compositional and Multimodal Perception" workshop at IEEE European Conference on Computer Vision (ECCV) 2020
- 2019-2020 Co-organizer and Guest Editor for an IEEE TPAMI special issue on "Graphs in Computer Vision"
- 2020 Co-organized for "Compositionality in Computer Vision" workshop at IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2020 at Seattle, USA
- 2020 Co-organized "The ActivityNet Large Scale Activity Recognition Challenge" workshop at IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2020, at Seattle, USA
- 2019 Lead organizer for "Scene Graph Representation and Learning" workshop at IEEE International Conference on Computer Vision 2019 at Seoul (ICCV), Korea
- 2018 Co-organized "The ActivityNet Large Scale Activity Recognition Challenge" workshop at IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2018, at Salt Lake City, USA

- 2017 Co-organized “The ActivityNet Large Scale Activity Recognition Challenge” workshop at IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017, at Honolulu, USA
- 2017 Program committee for “Groupsight: Workshop on Human Computation for Image and Video Analysis” workshop at The AAAI Conference on Human Computation and Crowdsourcing (HCOMP) 2017

Academic reviewer

- 2021-present Area chair from UIST, NeurIPS, CVPR
- 2016-present Reviewed papers from CVPR, CHI, UIST, CSCW, IJCV, ECCV, ICCV, TPAMI, NeurIPS, ICML

University service

- 2020-2021 Co-organized Stanford Human Computer Interaction group’s weekly talk series
- 2018-2021 Organized Stanford Vision and Learning group’s weekly talk series
- 2019-2020 Reviewed faculty applications as part of the Faculty Search Committee for Stanford’s Computer Science Department
- 2019-2020 Organized Stanford Human Computer Interaction group’s weekly talk series
- 2019-2020 Organized Stanford Bernstein group’s weekly talk series
- 2018-2019 Reviewed Ph.D. applications for Stanford’s Computer Science department as part of the Applications Committee
- 2015-2016 Organized Stanford Artificial Intelligence Entrepreneurship Club
- 2015 Organized logistics for Stanford’s AI outreach summer program

PH.D. ADVISING

- 2024-present Chenhao Zheng
- 2024-present Amita Kamath co-advised Kai-Wei Chang
- 2023-present Xiang Fan
- 2023-present Scott Geng co-advised with Pang Wei Koh (recipient of the [NSF Graduate Fellowship](#))
- 2023-present Zixian Ma
- 2023-present Sebastin Santy
- 2023-present Benlin Liu
- 2023-present Jieyu Zhang (recipient of the [Apple Graduate Fellowship](#))
- 2022-present Jiafei Duan co-advised with Dieter Fox (recipient of the [Singapore’s National Science Scholarship](#))
- 2022-present Ainaz Eftekhari co-advised with Ali Farhadi
- 2022-present Cheng-Yu Hsieh

MENTORSHIP

Current Postdocs

- 2023-present Wei-Chiu Ma (University of Washington) with Ali Farhadi

Current Ph.D. students

- 2023-present Yushi Hu (Princeton University) with Noah Smith
- 2023-present Ryan Liu (Princeton University) with Tom Griffiths
- 2023-present Kaitlin Flores (University of Washington) with Linda Shapiro
- 2023-present Wisdom Ikezogwo (University of Washington) with Linda Shapiro
- 2023-present M. Saygin Seyfioglu (University of Washington) with Linda Shapiro

2022-present Mahtab Bigverdi (University of Washington) with Linda Shapiro
2022-present Kalyani Marathe (University of Washington) with Linda Shapiro
2022-present Madeleine Grunde-McLaughlin (University of Washington) with Jeff Heer and Dan Weld
2022-present Ankit Vani (University of Montreal) with Aaron Courville
2022-present Arijit Ray (Boston University) with Kate Saenko and Bryan Plummer
2022-present Dong He (University of Washington) with Magdalena Balazinska
2022-present Enhao Zhang (University of Washington) with Magdalena Balazinska

Current masters students

2023-present Mayank Kumar (University of Washington)
2022-present Mona Gandhi (University of Pennsylvania)

Current undergraduate students

2024-present Peter Sushko (University of Washington)
2023-present Abhinav Bandari (University of Washington)
2023-present Ben Caffè (University of Washington)
2023-present Dylan Bunarto (University of Washington)
2023-present Jun Wang (University of Washington)
2023-present Chenhao Zheng (University of Michigan)
2023-present Ayana Bharadwaj (University of Washington)
2023-present Andre Ye (University of Washington)
2020-present Helena Vasconcelos (Stanford University)

Past Ph.D. students (placement after mentorship)

2020-2021 Madeleine Grunde-McLaughlin (continued CS Ph.D. at University of Washington)
2020-2021 Anelise Newman (continued CS Ph.D. Stanford University)
2020-2021 Siddharth Karamcheti (continued CS Ph.D. at Stanford University)
2019-2020 Jingwei Ji (continued CS Ph.D. at Stanford University)
2018-2019 Mitchell Gordon (continued CS Ph.D. at Stanford University)

Past masters students (placement after mentorship)

2022-2023 Oscar Michel (placement: CS Ph.D. at New York University with Saining Xie)
2022-2023 Kunal Pratap Singh (placement: CS Ph.D. at EPFL with Amir Zamir)
2019-2020 Shubhang Desai (placement: Applied ML Scientist at Microsoft)
2018-2020 Pranav Khadpe (placement: CS Ph.D. at Carnegie Mellon University co-advised by Chinmay Kulkarni and Geoff Kaufman)
2017-2019 Junwon Park (placement: program manager at Microsoft)
2017-2019 Apoorva Dornadula (placement: co-founder of Viralspace.ai startup)
2017-2019 Vincent Chen (placement: co-founder of Snorkel.ai startup)
2016-2018 Donsuk lee (placement: CS Ph.D. at University of South California advised by Yan Liu)
2016-2018 Ines Chami (placement: CS Ph.D. at Stanford University advised by Christopher Re)
2017-2017 Mohana Moorthy (placement: autonomy engineer at Uber)
2015-2017 Kenji Hata (placement: CS Ph.D. at Princeton University advised by Olga Russakovsky)
2016-2016 Vincent Sitzmann (placement: EE Ph.D. at Stanford University advised by Gordon Wetzstein)
2015-2016 Oliver Groth (placement: CS Ph.D. at Oxford University advised by Andrea Vedaldi)

- 2016-2017 Frederic Ren (placement: software engineer at Visa)
2016-2016 Yutian Li (placement: software engineer at Conscripton)

Past undergraduate students (placement after mentorship)

- 2020-2023 Zixian Ma (placement: CS Ph.D. at University of Washington with me)
2018-2022 Omer Gul (placement: CS Ph.D. at Cornell University with Yoav Artzi)
2021-2022 Mona Gandhi (placement: CS M.Sc. at University of Pennsylvania)
2021-2022 Jerry Hong (placement: Product Designer at Airbnb)
2020-2021 Madeleine Grunde-McLaughlin (placement: CS Ph.D. at University of Washington)
2020-2021 Kimberly Te (placement: Business Analyst at McKinsey & Company)
2017-2021 Austin Narcomey (placement: CS Ph.D. at Yale University)
2018-2021 Khaled Jedoui (placement: CS Ph.D. at Stanford University)
2017-2019 Michelle Lam (placement: CS Ph.D. at Stanford University)
2016-2019 Sho Arora (placement: machine learning engineer at Marcari)
2017-2018 Jihyeon Janel Lee (placement: CS M.Sc. at Stanford University)
2018-2018 Daniel Cai (placement: software engineer at Zoox)
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 M.Sc. at Stanford University)

PRESS

- 2024 **Holodeck: Creating 3D simulated spaces with diverse 3D object assets.**
 - Interesting Engineering - "A real Holodeck? Scientists recreate Star Trek technology with ChatGPT"
 - TechTimes - "Engineers Use ChatGPT, Video Game Assets to Recreate Star Trek's Holodeck"
 - TechXplore - "Engineers recreate Star Trek's Holodeck using ChatGPT and video game assets"
- 2024 **Revolutionizing Content Moderation in Digital Advertising: A Scalable LLM Approach.**
 - Matketpost - "Scaling Up LLM Reviews for Google Ads Content Moderation"
- 2022 **Tech that analyzes videos wins top prize at University of Washington computer science showcase.**
 - Geekwire - "VOCAL: Video Organization and Interactive AnaLytics"
- 2022 **Is this a deer I see? Socially aware AI adapts by asking questions of humans.**
 - Techxplore - "Socially situated AI"
- 2022 **When AI asks dumb questions, it gets smart fast.**
 - Science - "Socially situated AI"
- 2021 **Consumers Like Chatbots to Be Smartbut Not Too Smart.**
 - Wall Street Journal - "Scene Graphs"
- 2020 **Three ways Computer Vision is transforming marketing.**
 - Forbes - "Ranjay Krishna"
- 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" [video]

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"