

Yong Jae Lee

University of Wisconsin–Madison
Department of Computer Sciences
1210 W Dayton St
Madison, WI 53706 USA

Phone: (608) 262-1804
Email: yongjaelee@cs.wisc.edu
<https://cs.wisc.edu/~yongjaelee/>

APPOINTMENTS

Associate Professor (tenured) University of Wisconsin–Madison, Department of Computer Sciences	August 2021–Present
Cruise AI Visiting Faculty Cruise, San Francisco	July 2020–August 2021
Adjunct Associate Professor Associate Professor (tenured) Assistant Professor University of California, Davis, Department of Computer Science	July 2021–June 2022 July 2020–June 2021 July 2014–June 2020
Postdoctoral Fellow University of California, Berkeley, Department of EECS Carnegie Mellon University, Robotics Institute Supervised by Prof. Alexei A. Efros	Sept 2013–June 2014 Aug 2012–Aug 2013
Graduate Research Assistant University of Texas at Austin, Department of Computer Science Supervised by Prof. Kristen Grauman	June 2007–May 2012
Research Intern Microsoft Research, Interactive Visual Media Group Supervised by Dr. Larry Zitnick and Dr. Michael Cohen	June 2010–Sept 2010

EDUCATION

University of Texas at Austin, Austin, TX
Ph.D. in Electrical and Computer Engineering, May 2012
Thesis: “Visual Object Category Discovery in Images and Videos”
Supervised by Prof. Kristen Grauman

University of Texas at Austin, Austin, TX
M.S.E. in Electrical and Computer Engineering, August 2008
Thesis: “Foreground Focus: Finding Meaningful Features in Unlabeled Images”
Supervised by Prof. Kristen Grauman

University of Illinois at Urbana-Champaign, Champaign, IL
B.S. in Electrical Engineering (Honors), May 2006
Minor in Computer Science, Business

RESEARCH INTERESTS

My research interests are in computer vision and machine learning. In particular, I am interested in creating robust recognition systems that can learn to understand visual data with minimal human supervision.

SELECTED HONORS AND AWARDS

- Sony Focused Research Award, 2023
- SACM Student Choice Professor of the Year Award, 2022
Awarded to one professor in the CS department each year to recognize excellence in teaching, as voted on by the students.
- Adobe Data Science Research Award, 2022
- Best Paper Award, BMVC 2020
For the paper, “Delving Deeper into Anti-aliasing in ConvNets”, with Xueyan Zou, Fanyi Xiao, and Zhiding Yu.
- Sony Focused Research Award, 2020
- Most Innovative Award, COCO Object Detection Challenge, ICCV 2019
For the paper, “YOLACT: Real-time Instance Segmentation”, with Daniel Bolya, Chong Zhou, and Fanyi Xiao.
- UC Davis College of Engineering Outstanding Junior Faculty Award, 2019
Awarded to one Assistant Professor in the College of Engineering each year.
- Adobe Data Science Research Award, 2019
- AWS Machine Learning Research Award 2019
- National Science Foundation (NSF) CAREER Award, 2018
- AWS Machine Learning Research Award 2018
- Army Research Office (ARO) Young Investigator Program (YIP) Award, 2017
- UC Davis Hellman Foundation Fellowship, 2017
Awarded to Assistant Professors who exhibit potential for great distinction in their research. One of 15 awardees university-wide.
- Outstanding Reviewer Awards, ICCV 2013, CVPR 2015
- UT-Austin ECE dept. sole nominee for university-wide Outstanding Dissertation Award, 2012
- Doctoral Consortium Award, CVPR 2011

TEACHING EXPERIENCE

Faculty Instructor at UW–Madison:

- CS 639: Deep Learning for Computer Vision (Undergraduate course). Spring 2024. Enrollment: 104. Instructor rating: 6.03/7.0
- CS 839: Learning based Image Synthesis and Manipulation (Graduate course). Fall 2023. Enrollment: 33. Instructor rating: 6.18/7.0

- CS 639: Deep Learning for Computer Vision (Undergraduate course). Spring 2023. Enrollment: 102. Instructor rating: 5.70/7.0
- CS 839: Learning based Image Synthesis and Manipulation (Graduate course). Fall 2022. Enrollment: 42. Instructor rating: 6.38/7.0
- CS 839: Deep Learning for Visual Recognition (Graduate course). Spring 2022. Enrollment: 34. Instructor rating: 6.75/7.0

Faculty Instructor at UC Davis:

- ECS 174: Computer Vision (Undergraduate course). Spring 2020. Enrollment: 163. Instructor rating: 4.5/5.0
- ECS 269: Visual Recognition (Graduate course). Fall 2019. Enrollment: 54. Instructor rating: 4.5/5.0
- ECS 174: Computer Vision (Undergraduate course). Spring 2019. Enrollment: 171. Instructor rating: 4.4/5.0
- ECS 269: Visual Recognition (Graduate course). Fall 2018. Enrollment: 53. Instructor rating: 4.8/5.0
- ECS 174: Computer Vision (Undergraduate course). Spring 2018. Enrollment: 147. Instructor rating: 4.4/5.0
- ECS 289G: Visual Recognition (Graduate course). Winter 2018. Enrollment: 43. Instructor rating: 4.5/5.0
- ECS 174: Computer Vision (Undergraduate course). Spring 2017. Enrollment: 137. Instructor rating: 4.1/5.0
- ECS 289G: Visual Recognition (Graduate course). Fall 2016. Enrollment: 20. Instructor rating: 4.7/5.0
- ECS 289G: Visual Recognition (Graduate course). Fall 2015. Enrollment: 14. Instructor rating: 4.9/5.0
- ECS 189G: Intro to Computer Vision (Undergraduate course). Spring 2015. Enrollment: 141. Instructor rating: 4.2/5.0
- ECS 289H: Visual Recognition (Graduate course). Fall 2014. Enrollment: 15. Instructor rating: 4.9/5.0

Teaching Assistant at UT Austin:

- CS 376: Computer Vision (Undergraduate/Graduate course). Fall 2009, Spring 2011.

PUBLICATIONS

Books and Book Chapters

1. Fanyi Xiao and Yong Jae Lee. Localizing and Visualizing Relative Attributes. *Chapter in Visual Attributes*. R. Feris, C. Lampert, and D. Parikh, Editors. Springer, 2016.
2. Yong Jae Lee, Alexei A. Efros, and Martial Hebert. Discovering Mid-level Visual Connections in Space and Time. *Chapter in Visual Analysis and Geo-Localization of Large Scale Imagery*. A. Zamir, A. Hakeem, L. Van Gool, M. Shah, and R. Szeliski, Editors. Springer, 2016
3. Hyun Soo Park, Albert Ali Salah, Yong Jae Lee, Louis-Philippe Morency, Yaser Sheikh, Rita Cucchiara (Eds.). *Human Behavior Understanding*. Lecture Notes in Computer Science, Springer, 2014.

Journal Articles

1. Rafael E. P. Ferreira, Yong Jae Lee, and Joao Reboucas Dorea. Using Pseudo-Labeling to Improve Performance of Deep Neural Networks for Animal Identification. *Nature Scientific Reports*, 2023.
2. Xueyan Zou, Fanyi Xiao, Zhiding Yu, Yuheng Li, and Yong Jae Lee. Delving Deeper into Anti-aliasing in ConvNets. *International Journal of Computer Vision (IJCV)*, 2022. [[Invited article for best papers of BMVC 2020](#)]
3. Daniel Bolya*, Chong Zhou*, Fanyi Xiao, and Yong Jae Lee. YOLACT++: Better Real-time Instance Segmentation. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, August 2020. (*equal contribution)
4. Farzad Fereidouni, Austin Todd, Yuheng Li, Che-Wei Chang, Keith Luong, Avi Rosenberg, Yong Jae Lee, James W. Chan, Alexander Borowsky, Karen Matsukuma, Kuang-Yu Jen, and Richard Levenson. Dual-mode Emission and Transmission Microscopy for Virtual Histochemistry using Hematoxylin- and Eosin-Stained Tissue Sections. *Biomedical Optics Express*, Vol. 10, No. 12, November 2019.
5. M. Rezaur Rahman, Jinyoung Han, Yong Jae Lee, and Chen-Nee Chuah. Analyzing the Adoption and Cascading Process of OSN-Based Gifting Applications: An Empirical Study. *ACM Transactions on the Web (TWEB)*, Vol. 11, No. 2, April 2017.
6. Yong Jae Lee and Kristen Grauman. Predicting Important Objects for Egocentric Video Summarization. *International Journal of Computer Vision (IJCV)*, Vol. 114, No. 1, pp. 38-55, January 2015.
7. Jun-Yan Zhu, Yong Jae Lee, and Alexei A. Efros. AverageExplorer: Interactive Exploration and Alignment of Visual Data Collections. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, Vol. 33, No. 4, July 2014.
8. Yong Jae Lee and Kristen Grauman. Object-Graphs for Context-Aware Visual Category Discovery. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, Vol. 34, No. 2, pp. 346-358, February 2012.
9. Yong Jae Lee, C. Lawrence Zitnick, and Michael Cohen. ShadowDraw: Real-Time User Guidance for Freehand Drawing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, Vol. 30, No. 4, July 2011.
10. Yong Jae Lee and Kristen Grauman. Foreground Focus: Unsupervised Learning from Partially Matching Images. *International Journal of Computer Vision (IJCV)*, Vol. 85, No. 2, pp. 143-166, May 2009.

Peer Reviewed Conference Papers (acceptance rates typically ~ 2%-25%)

1. Thao Nguyen, Haotian Liu, Mu Cai, Yuheng Li, Utkarsh Ojha, and Yong Jae Lee. Yo'LLaVA: Your Personalized Language and Vision Assistant. In *Neural Information Processing Systems (NeurIPS)*, Vancouver, CA, December 2024.
2. Xueyan Zou, Linjie Li, Jianfeng Wang, Jianwei Yang, Mingyu Ding, Junyi Wei, Zhengyuan Yang, Feng Li, Hao Zhang, Shilong Liu, Arul Aravinthan, Yong Jae Lee*, and Lijuan Wang*. What Can Foundation Models' Embeddings Do? In *Neural Information Processing Systems (NeurIPS)*, Vancouver, CA, December 2024. (*equal advising)
3. Bocheng Zou*, Mu Cai*, Jianrui Zhang, and Yong Jae Lee. VGBench: Evaluating Large Language Models on Vector Graphics Understanding and Generation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Miami, FL, November 2024. (*equal contribution)

4. Young Kyun Jang, Junmo Kang, Yong Jae Lee, and Donghyun Kim. MATE: Meet At The Embedding – Connecting Images with Long Texts. In *Findings of the Conference on Empirical Methods in Natural Language Processing (EMNLP Findings)*, Miami, FL, November 2024.
5. Yuheng Li, Haotian Liu, Mu Cai, Yijun Li, Eli Shechtman, Zhe Lin, Yong Jae Lee, and Krishna Kumar Singh. Removing Distributional Discrepancies in Captions Improves Image-Text Alignment. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Milan, Italy, October 2024.
6. Jianrui Zhang*, Mu Cai*, Tengyang Xie, and Yong Jae Lee. CounterCurate: Enhancing Physical and Semantic Visio-Linguistic Compositional Reasoning via Counterfactual Examples. In *Findings of the Association for Computational Linguistics (ACL Findings)*, Bangkok, Thailand, August 2024. (*equal contribution)
7. Mu Cai, Chenxu Luo, Yong Jae Lee, and Xiaodong Yang. Cross-Modal Self-Supervised Learning with Effective Contrastive Units for Point Clouds. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Abu Dhabi, UAE, October 2024.
8. Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved Baselines with Visual Instruction Tuning (LLaVA-1.5). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, June 2024. ([highlight](#), 2.8% acceptance rate)
9. Mu Cai, Haotian Liu, Siva Karthik Mustikovela, Gregory P. Meyer, Yuning Chai, Dennis Park, and Yong Jae Lee. ViP-LLaVA: Making Large Multimodal Models Understand Arbitrary Visual Prompts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, June 2024.
10. Thao Nguyen, Utkarsh Ojha, Yuheng Li, Haotian Liu, and Yong Jae Lee. Edit One for All: Interactive Batch Image Editing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, June 2024.
11. Moniek Smink, Haotian Liu, Doerte Doepfer, and Yong Jae Lee. Computer Vision on the Edge: Individual Cattle Identification in Real-time with ReadMyCow System. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, Waikoloa, HI, January 2024.
12. Yuexiang Zhai, Shengbang Tong, Xiao Li, Mu Cai, Qing Qu, Yong Jae Lee, and Yi Ma. Investigating the Catastrophic Forgetting in Multimodal Large Language Model Fine-Tuning. In *Conference on Parsimony and Learning (CPAL)*, Hong Kong, January 2024.
13. Kazuki Shin, Sankalp Yamsani, Roman Mineyev, Hongyu Chen, Nitish Gandhi, Yong Jae Lee, and Joohyung Kim. Exploring the Capabilities of a General-Purpose Robotic Arm in Chess Gameplay. In *Proceedings of the IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, Austin, TX, December 2023.
14. Haotian Liu*, Chunyuan Li*, Qingyang Wu, and Yong Jae Lee. Visual Instruction Tuning (LLaVA). In *Neural Information Processing Systems (NeurIPS)*, New Orleans, LA, December 2023. (*equal contribution) ([oral presentation](#), 0.5% acceptance rate)
15. Utkarsh Ojha*, Yuheng Li*, Anirudh Sundara Rajan*, Yingyu Liang, and Yong Jae Lee. Dissecting Knowledge Distillation: An Exploration of its Inner Workings and Applications. In *Neural Information Processing Systems (NeurIPS)*, New Orleans, LA, December 2023. (*equal contribution)
16. Thao Nguyen, Yuheng Li, Utkarsh Ojha, and Yong Jae Lee. Visual Instruction Inversion: Image Editing via Image Prompting. In *Neural Information Processing Systems (NeurIPS)*, New Orleans, LA, December 2023.

17. Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Wang, Lijuan Wang, Jianfeng Gao*, and Yong Jae Lee*. Segment Everything Everywhere All at Once. In *Neural Information Processing Systems (NeurIPS)*, New Orleans, LA, December 2023. (*equal advising)
18. Zeyi Huang, Andy Zhou, Zijian Ling, Mu Cai, Haohan Wang, and Yong Jae Lee. A Sentence Speaks a Thousand Images: Domain Generalization through Distilling CLIP with Language Guidance. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Paris, France, October 2023.
19. Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li*, and Yong Jae Lee*. GLIGEN: Open-Set Grounded Text-to-Image Generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Vancouver, CA, June 2023. (*equal advising)
20. Xueyan Zou*, Zi-Yi Dou*, Jianwei Yang*, Zhe Gan, Linjie Li, Chunyuan Li, Xiyang Dai, Jianfeng Wang, Lu Yuan, Nanyun Peng, Lijuan Wang, Harkirat Behl, Yong Jae Lee‡, and Jianfeng Gao‡. Generalized Decoding for Pixel, Image, and Language. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Vancouver, CA, June 2023. (*,‡ equal contribution)
21. Utkarsh Ojha*, Yuheng Li*, and Yong Jae Lee. Towards Universal Fake Image Detectors that Generalize Across Generative Models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Vancouver, CA, June 2023. (*equal contribution)
22. Haotian Liu, Kilho Son, Jianwei Yang, Ce Liu, Jianfeng Gao, Yong Jae Lee*, and Chunyuan Li*. REACT: Learning Customized Visual Models with Retrieval-Augmented Knowledge. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Vancouver, CA, June 2023. (*equal advising) ([highlight](#), 2.5% acceptance rate)
23. Zhuoran Yu, Yin Li, and Yong Jae Lee. InPL: Pseudo-labeling the Inliers First for Imbalanced Semi-supervised Learning. In *International Conference on Learning Representations (ICLR)*, May 2023.
24. Chunyuan Li*, Haotian Liu*, Liunian Harold Li, Pengchuan Zhang, Jyoti Aneja, Jianwei Yang, Ping Jin, Houdong Hu, Zicheng Liu, Yong Jae Lee, and Jianfeng Gao. ELEVATER: A Benchmark and Toolkit for Evaluating Language-Augmented Visual Models. In *Neural Information Processing Systems (NeurIPS)*, Datasets and Benchmarks Track, New Orleans, LA, December 2022. (*equal contribution)
25. Haotian Liu, Mu Cai, and Yong Jae Lee. Masked Discrimination for Self-Supervised Learning on Point Clouds. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Tel Aviv, Israel, October 2022.
26. Yuheng Li, Yijun Li, Jingwan Lu, Eli Shechtman, Yong Jae Lee, and Krishna Kumar Singh. Contrastive Learning for Diverse Disentangled Foreground Generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Tel Aviv, Israel, October 2022.
27. Haohan Wang, Zeyi Huang, Hanlin Zhang, Yong Jae Lee, and Eric Xing. Toward Learning Human-aligned Cross-domain Robust Models by Countering Misaligned Features. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, Eindhoven, Netherlands, August 2022.
28. Yang Xue, Yuheng Li, Krishna Kumar Singh, and Yong Jae Lee. GIRAFFE HD: A High-Resolution 3D-aware Generative Model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, LA, June 2022.

29. Zeyi Huang*, Haohan Wang*, Dong Huang, Yong Jae Lee[†], and Eric Xing[†]. The Two Dimensions of Worst-case Training and the Integrated Effect for Out-of-domain Generalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, LA, June 2022. (*[†]equal contribution)
30. Maheen Rashid, Sofia Broome, Katrina Ask, Elin Hernlund, Pia Haubro Andersen, Hedvig Kjellstrom, and Yong Jae Lee. Equine Pain Behaviour Classification via Self-supervised Disentangled Pose Representation. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, January 2022.
31. Yuheng Li, Krishna Kumar Singh, Yang Xue, and Yong Jae Lee. PartGAN: Weakly-supervised Part Decomposition for Image Generation and Segmentation. In *Proceedings of the British Machine Vision Conference (BMVC)* 2021.
32. Yuheng Li, Yijun Li, Jingwan Lu, Eli Shechtman, Yong Jae Lee, and Krishna Kumar Singh. Generate Globally, Enhance Locally: High Resolution Conditional Image Generation with Class-specific Models. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, October 2021.
33. Haotian Liu*, Rafael A. Rivera-Soto*, Fanyi Xiao, and Yong Jae Lee. YolactEdge: Real-time Instance Segmentation on the Edge. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Xi'an, China, June 2021. (*equal contribution)
34. Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A. Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. Few-shot Image Generation via Cross-domain Correspondence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
35. Xueyan Zou, Linjie Yang, Ding Liu, and Yong Jae Lee. Progressive Temporal Feature Alignment Network for Video Inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
36. Utkarsh Ojha, Krishna Kumar Singh, and Yong Jae Lee. Generating Furry Cars: Disentangling Object Shape and Appearance across Multiple Domains. In *International Conference on Learning Representations (ICLR)*, May 2021.
37. Rajat Arora and Yong Jae Lee. SinGAN-GIF: Learning a Generative Video Model from a Single GIF. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, January 2021.
38. Utkarsh Ojha, Krishna Kumar Singh, Cho-Jui Hsieh, and Yong Jae Lee. Elastic-InfoGAN: Generative Modeling of Disentangled Representations in Class-Imbalanced Data. In *Neural Information Processing Systems (NeurIPS)*, Vancouver, Canada, December 2020.
39. Xueyan Zou, Fanyi Xiao, Zhiding Yu, and Yong Jae Lee. Delving Deeper into Anti-aliasing in ConvNets. In *Proceedings of the British Machine Vision Conference (BMVC)*, Manchester, United Kingdom, September 2020. ([oral presentation](#), 5.1% acceptance rate)
Best Paper Award (1 of 670 submissions)
40. Xiuye Gu, Weixin Luo, Michael Ryoo, and Yong Jae Lee. Password-conditioned Anonymization and Deanonimization with Face Identity Transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Glasgow, United Kingdom, August 2020.
41. Zainul Abi Din, Hari Venugopalan, Jaime Park, Andy Li, Weisu Yin, Haohui Mai, Yong Jae Lee, Steven Liu, Samuel T. King. Boxer: Preventing Fraud by Scanning Credit Cards. In *Proceedings of the USENIX Security Symposium (USENIX Security)*, August 2020.

42. Yuheng Li, Krishna Kumar Singh, Utkarsh Ojha, and Yong Jae Lee. MixNMatch: Multifactor Disentanglement and Encoding for Conditional Image Generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, June 2020.
43. Krishna Kumar Singh, Dhruv Mahajan, Kristen Grauman, Yong Jae Lee, Matt Feiszli, and Deepti Ghadiyaram. Don't Judge an Object by Its Context: Learning to Overcome Contextual Bias. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, June 2020. ([oral presentation](#), 5.7% acceptance rate)
44. Zhongzheng Ren, Zhiding Yu, Xiaodong Yang, Ming-Yu Liu, Yong Jae Lee, Alexander Schwing, and Jan Kautz. Instance-aware, Context-focused, and Memory-efficient Weakly-supervised Object Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, June 2020.
45. Maheen Rashid, Hedvig Kjellstrom, and Yong Jae Lee. Action Graphs: Weakly-supervised Action Localization with Graph Convolution Networks. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, Aspen, CO, March 2020.
46. Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. YOLACT: Real-time Instance Segmentation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Seoul, Korea, November 2019. ([oral presentation](#), 4.3% acceptance rate)
Most Innovative Award, COCO Object Detection Challenge, ICCV 2019
47. Fanyi Xiao, Haotian Liu, and Yong Jae Lee. Identity from here, Pose from there: Self-supervised Disentanglement and Generation of Objects using Unlabeled Videos. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Seoul, Korea, November 2019.
48. Krishna Kumar Singh*, Utkarsh Ojha*, and Yong Jae Lee. FineGAN: Unsupervised Hierarchical Disentanglement for Fine-Grained Object Generation and Discovery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, June 2019. (*equal contribution) ([oral presentation](#), 5.5% acceptance rate)
49. Krishna Kumar Singh and Yong Jae Lee. You reap what you sow: Using Videos to Generate High Precision Object Proposals for Weakly-supervised Object Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, June 2019.
50. Xiuye Gu, Yijie Wang, Chongruo Wu, Yong Jae Lee, and Panqu Wang. HPLFlowNet: Hierarchical Permutohedral Lattice FlowNet for Scene Flow Estimation on Large-scale Point Clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, June 2019.
51. Krishna Kumar Singh, Santosh Divvala, Ali Farhadi, and Yong Jae Lee. DOCK: Detecting Objects by transferring Common-sense Knowledge. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Munich, Germany, September 2018.
52. Fanyi Xiao and Yong Jae Lee. Video Object Detection with an Aligned Spatial-Temporal Memory. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Munich, Germany, September 2018.
53. Zhongzheng Ren, Yong Jae Lee, and Michael Ryoo. Learning to Anonymize Faces for Privacy Preserving Action Detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Munich, Germany, September 2018.
54. Mingyang Zhou, Runxiang Cheng, Yong Jae Lee, and Zhou Yu. A Visual Attention Grounding Neural Model for Multimodal Machine Translation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium, November 2018.

55. Zhongzheng Ren and Yong Jae Lee. Cross-Domain Self-supervised Multi-task Feature Learning using Synthetic Imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, UT, June 2018.
56. Wenjian Hu, Krishna Kumar Singh*, Fanyi Xiao*, Jinyoung Han, Chen-Nee Chuah, and Yong Jae Lee. Who Will Share My Image? Predicting the Content Diffusion Path in Online Social Networks. In *ACM International Conference on Web Search and Data Mining (WSDM)*, Los Angeles, CA, February 2018. (*equal contribution)
57. Maheen Rashid, Sofia Broome, Pia Andersen, Karina Glerup, and Yong Jae Lee. What should I annotate? An automatic tool for finding video segments for EquiFACS annotation. In *International Conference on Methods and Techniques in Behavioral Research (Measuring Behavior)*, Manchester, UK, June 2018.
58. Pia Andersen, Karina Glerup, Jennifer Wathan, Britt Coles, Hedvig Kjellstrom, Sofia Broome, Yong Jae Lee, Maheen Rashid, Claudia Sonder, Erika Rosenberger, and Deborah Forster. Can a Machine Learn to See Horse Pain? An Interdisciplinary Approach Towards Automated Decoding of Facial Expressions of Pain in the Horse. In *International Conference on Methods and Techniques in Behavioral Research (Measuring Behavior)*, Manchester, UK, June 2018.
59. Krishna Kumar Singh and Yong Jae Lee. Hide-and-Seek: Forcing a Network to be Meticulous for Weakly-supervised Object and Action Localization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Venice, Italy, October 2017.
60. Fanyi Xiao, Leonid Sigal, and Yong Jae Lee. Weakly-supervised Visual Grounding of Phrases with Linguistic Structures. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, June 2017.
61. Maheen Rashid, Xiuye Gu, and Yong Jae Lee. Interspecies Knowledge Transfer for Facial Keypoint Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, June 2017.
62. Chenyou Fan, Jangwon Lee, Mingze Xu, Krishna Kumar Singh, Yong Jae Lee, David Crandall and Michael Ryoo. Identifying First-Person Camera Wearers in Third-Person Videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, June 2017.
63. Zhongzheng Ren, Adriana Noronha, Annie Vogel Ciernia, and Yong Jae Lee. Who Moved My Cheese? Automatic Annotation of Rodent Behaviors with Convolutional Neural Networks. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, Santa Rosa, CA, March 2017.
64. Krishna Kumar Singh and Yong Jae Lee. End-to-End Localization and Ranking for Relative Attributes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Amsterdam, Netherlands, October 2016.
65. Fanyi Xiao and Yong Jae Lee. Track and Segment: An Iterative Unsupervised Approach for Video Object Proposals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, June 2016. ([spotlight presentation](#), 9.7% acceptance rate)
66. Krishna Kumar Singh, Fanyi Xiao, and Yong Jae Lee. Track and Transfer: Watching Videos to Simulate Strong Human Supervision for Weakly-Supervised Object Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, June 2016.
67. Fanyi Xiao and Yong Jae Lee. Discovering the Spatial Extent of Relative Attributes. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Santiago, Chile, December 2015. ([oral presentation](#), 3.3% acceptance rate)

68. Tinghui Zhou, Yong Jae Lee, Stella Yu, and Alexei A. Efros. FlowWeb: Joint Image Set Alignment by Weaving Consistent, Pixel-wise Correspondences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, June 2015. (oral presentation, 3.3% acceptance rate)
69. Hyun Oh Song, Yong Jae Lee, Stefanie Jegelka, and Trevor Darrell. Weakly-supervised Discovery of Visual Pattern Configurations. In *Neural Information Processing Systems (NIPS)*, Montreal, Canada, December 2014.
70. Yong Jae Lee, Alexei A. Efros, and Martial Hebert. Style-aware Mid-level Representation for Discovering Visual Connections in Space and Time. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Sydney, Australia, December 2013. (oral presentation, 2.5% acceptance rate)
71. Yong Jae Lee, Joydeep Ghosh, and Kristen Grauman. Discovering Important People and Objects for Egocentric Video Summarization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Providence, RI, June 2012.
72. Yong Jae Lee, Jaechul Kim, and Kristen Grauman. Key-Segments for Video Object Segmentation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Barcelona, Spain, November 2011.
73. Yong Jae Lee and Kristen Grauman. Face Discovery with Social Context. In *Proceedings of the British Machine Vision Conference (BMVC)*, Dundee, Scotland, August 2011.
74. Yong Jae Lee and Kristen Grauman. Learning the Easy Things First: Self-Paced Visual Category Discovery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Colorado Springs, CO, June 2011.
75. Yong Jae Lee and Kristen Grauman. Object-Graphs for Context-Aware Category Discovery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, San Francisco, CA, June 2010. (oral presentation, 4.5% acceptance rate)
76. Yong Jae Lee and Kristen Grauman. Collect-Cut: Segmentation with Top-Down Cues Discovered in Multi-Object Images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, San Francisco, CA, June 2010.
77. Yong Jae Lee and Kristen Grauman. Shape Discovery from Unlabeled Image Collections. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Miami, FL, June 2009.
78. Yong Jae Lee and Kristen Grauman. Foreground Focus: Finding Meaningful Features in Unlabeled Images. In *Proceedings of the British Machine Vision Conference (BMVC)*, Leeds, United Kingdom, September 2008. (oral presentation, 12.5% acceptance rate)
79. Changhai Xu, Yong Jae Lee, and Benjamin Kuipers. Ray-based Color Image Segmentation. In *Proceedings of the Canadian Conference on Computer and Robot Vision (CRV)*, Windsor, Ontario, May 2008.

INVITED TALKS

- Colloquium, Washington University in St. Louis, October 2024.
“Generalist Vision-Language Models.”
- Invited Talk, ML seminar series, KLA, October 2024.
“Next Steps in Generalist Multimodal Models.”

- Invited Talk, Midwest Computer Vision Workshop, September 2024.
“Generalist Vision-Language Models.”
- Invited Talk, London Machine Learning Meetup, September 2024.
“Next Steps in Generalist Multimodal Models.”
- Distinguished AI Speaker Seminar, Salesforce AI, September 2024.
“Next Steps in Generalist Multimodal Models.”
- Keynote Talk, Korean Conference on Computer Vision, Busan, South Korea, August 2024.
“Learning to Understand Our Multimodal World.”
- Distinguished Speaker, National Yang Ming Chiao Tung University H&J Global Chair Colloquium, Hsinchu, Taiwan, August 2024.
“Learning to Understand Our Multimodal World.”
- Distinguished Speaker, National Yang Ming Chiao Tung University H&J Global Chair Colloquium, Hsinchu, Taiwan, August 2024.
“Generating Images and Detecting Them.”
- Invited Talk, Samsung, Seoul, South Korea, July 2024.
“Learning to Understand Our Multimodal World.”
- Invited Talk, Nota AI, Seoul, South Korea, July 2024.
“Next Steps in Generalist Multimodal Models.”
- Invited Talk, Korea University, Seoul, South Korea, July 2024.
“Next Steps in Generalist Multimodal Models.”
- Invited Talk, Seoul National University, Seoul, South Korea, July 2024.
“Next Steps in Generalist Multimodal Models.”
- Invited Talk, Meta Reality Labs, Redmond, WA, June 2024.
“Learning to Understand Our Multimodal World with Minimal Supervision.”
- Invited Talk, CVPR Workshop on Computer Vision in the Wild, Seattle, WA, June 2024.
“Building Steerable Generalist Multimodal Models.”
- Invited Talk, CVPR Workshop on What is Next in Multimodal Foundation Models?, Seattle, WA, June 2024.
“Learning to Understand Our Multimodal World.”
- Invited Talk, CVPR Workshop on Prompting in Vision, Seattle, WA, June 2024.
“Visual Prompting in Large Multimodal Models.”
- Invited Talk, CVPR Tutorial on Generalist Agent AI, Seattle, WA, June 2024.
“Generalist Multimodal Models.”
- Keynote Talk, Embedded Vision Summit, Santa Clara, CA, May 2024.
“Learning to Understand Our Multimodal World with Minimal Supervision.”
- AI Speaker Series, Upwork, May 2024.
“Learning to Understand Our Multimodal World with Minimal Supervision.”
- Computer Science Colloquium, University of Pittsburgh, April 2024.
“Learning to Understand Our Multimodal World with Minimal Supervision.”
- Microsoft Research, April 2024.
“Next Steps in Large Multimodal Models.”

- HP Labs, March 2024.
“Large Multimodal Vision-Language Models for Image Understanding.”
- Adobe, San Jose, CA, December 2023.
“Large Multimodal Vision-Language Models for Image Understanding.”
- Carnegie Mellon University, Robotics Institute VASC Seminar, October 2023.
“Large Multimodal Vision-Language Models for Image Generation and Understanding.”
- MLLM Talk, University of Wisconsin–Madison, October 2023.
“Large Multimodal Vision-Language Models for Image Generation and Understanding.”
- Yonsei University, LG Energy Solution, Krafton, MakinaRocks, August 2023.
“Large Multimodal Vision-Language Models for Image Generation and Understanding.”
- MLLM Talk, University of Wisconsin–Madison, February 2023.
“On Detecting Fake (AI-generated) Images.”
- Keynote Talk, Hyundai & Seoul National University AI Consortium, April 2022.
“Towards 3D Visual Understanding with Minimal Human Supervision.”
- SILO Talk, University of Wisconsin–Madison, March 2022.
“Towards 3D Visual Understanding with Minimal Human Supervision.”
- Keynote Talk, KAIST AI Graduate School, October 2021.
“Controllable Visual Content Generation without Expensive Supervision.”
- Keynote Talk, OpenMMLab Tutorial, CVPR 2021.
“Real-time Instance Segmentation with the YOLACT Family.”
- AI Summer School, Seoul National University, August 2020.
“Learning Disentangled Visual Representations with Minimal Human Supervision.”
- School of Interactive Computing Seminar, Georgia Tech, April 2020.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Computer Sciences Seminar, University of Wisconsin–Madison, April 2020.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Computer Science Colloquium, Purdue University, March 2020.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Computer Science Seminar, Brown University, March 2020.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Computer Vision Seminar, UC Berkeley, February 2020.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Cruise, San Francisco, CA, February 2020.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Google Research, Mountain View, CA, October 2019.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Computer Science Seminar, University of British Columbia, September 2019.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Computer Science Seminar, University of California, Irvine, September 2019.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Snap Research, Santa Monica, CA, September 2019.
“Learning to Understand Visual Data with Minimal Human Supervision.”

- ECE Seminar, University of Southern California, September 2019.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Seminar, LG Research, Seoul, Korea, July 2019.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Computer Science Colloquium, University of Maryland, College Park, April 2019.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- Seminar, NVIDIA Research, Santa Clara, CA, March 2019.
“Learning to Understand Visual Data with Minimal Human Supervision.”
- CVPR Area Chair Workshop, San Diego, February 2019.
“FineGAN: Unsupervised Hierarchical Disentanglement for Fine-Grained Object Generation and Discovery.”
- Center for Mind and Brain Seminar, University of California, Davis, February 2019.
“Learning to Localize, Generate, and Anonymize Objects with Indirect Supervision.”
- Carnegie Mellon University, Robotics Institute VASC Seminar, October 2018.
“Learning to Localize and Anonymize Objects with Indirect Supervision.”
- Data Science Institute Seminar, Lawrence Livermore National Laboratory, September 2018.
“Learning to Localize and Anonymize Objects with Indirect Supervision.”
- Computer Vision Seminar, KAIST University, Seoul, Korea, August 2018.
“Learning to Localize and Anonymize Objects with Indirect Supervision.”
- Korean Conference on Computer Vision, Seoul, Korea, July 2018.
“Cross-Domain Self-supervised Multi-task Feature Learning using Synthetic Imagery.”
- Computer Vision Seminar, UC Berkeley, July 2018.
“Learning to Localize and Anonymize Objects with Indirect Supervision.”
- Center for Vision, Cognition, Learning, and Autonomy Seminar, UCLA, February 2018.
“Learning to See with Weak Human Supervision.”
- Computational Neuroscience Seminar, University of California, Davis, November 2017.
“Computer Vision for Animal Behavior Recognition.”
- Seminar, Intel Labs, Santa Clara, CA, October 2017.
“Learning to See with Minimal Human Supervision.”
- Seoul National University, UNIST, Samsung, LG, Naver, Kakao, Korea, August 2017.
“Learning Where and How to Look for Weakly-Supervised Visual Recognition.”
- Keynote Talk, VALSE Seminar, China, April 2017.
“Learning Where and How to Look for Weakly-Supervised Visual Recognition.”
- Keynote Talk, 1st International Workshop on Human Activity Analysis with Highly Diverse Cameras, Santa Rosa, CA, March 2017.
“Hide-and-Seek: Forcing a Network to be Meticulous for Weakly-Supervised Object and Action Localization.”
- Computer Science Colloquium, Sonoma State University, September 2016.
“Attending to the Right Image Regions for Weakly-Supervised Visual Recognition.”
- Seminar, Lawrence Livermore National Laboratory, September 2016.
“Discovering the Relevant Image Regions for Weakly-Supervised Visual Recognition.”

- Machine Learning Seminar, University of California, Davis, December 2015.
“Watching Videos to Simulate Strong Human Supervision for Visual Recognition.”
- Korean Conference on Computer Vision, Seoul, Korea, August 2015.
“Interactive Exploration and Alignment of Visual Data Collections.”
- Computational Neuroscience Seminar, University of California, Davis, May 2015.
“Visual Data Mining in Image Collections.”
- Computer Science Departmental Colloquium, University of California, Davis, April 2014.
“Visual Data Mining in Image Collections.”
- Computer Science Departmental Seminar, University of Massachusetts, Amherst, March 2014.
“Visual Data Mining in Image Collections.”
- Creative Technologies Lab, Adobe Research, San Francisco, CA, February 2014.
“Visual Data Mining in Image Collections.”
- IEEE International Conf. on Computer Vision, Sydney, Australia, December 2013.
“Style-aware Mid-level Representation for Discovering Visual Connections in Space and Time.”
- IEEE CVPR Workshop on Egocentric Vision, Providence, RI, June 2012.
“Discovering Important People and Objects for Egocentric Video Summarization.”
- Carnegie Mellon University, Robotics Institute VASC Seminar, Pittsburgh, December 2011.
“Visual Category Discovery in Images and Videos.”
- Microsoft Research, MSRNE Vision Seminar, Boston, MA, December 2011.
“Visual Category Discovery in Images and Videos.”
- ACM SIGGRAPH conference, Vancouver, Canada, August 2011.
“ShadowDraw: Real-Time User Guidance for Freehand Drawing.”
- Microsoft Research, Vision Technology (VisTech) Seminar, Redmond, WA, June 2010.
“Collect-Cut: Segmentation with Top-Down Cues Discovered in Multi-Object Images.”
- IEEE Conf. on Computer Vision and Pattern Recognition, San Francisco, CA, June 2010.
“Object-Graphs for Context-Aware Category Discovery.”
- British Machine Vision Conference, Leeds, U.K., September 2008.
“Foreground Focus: Finding Meaningful Features in Unlabeled Images.”

INVITED PANELS

- Embedded Vision Summit, Santa Clara, CA, May 2024.
“Multimodal LLMs at the Edge: Are We There Yet?”
- CVPR Workshop on Prompting in Vision, Seattle, WA, June 2024.

PATENTS

- C. Lawrence Zitnick, Yong Jae Lee, and Michael Cohen. Realtime User Guidance for Freehand Drawing. US Patent US8827710B2. Granted Sep 2014.

LIVE DEMOS

- Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. YOLACT: Real-time Instance Segmentation. ICCV 2019.
- Zhongzheng Ren, Yong Jae Lee, Hyun Jong Yang, and Michael Ryoo. Activity-Preserving Face Anonymization for Privacy Protection. ECCV 2018.

PROFESSIONAL SERVICE ACTIVITY

- **Program Chair:** WACV 2023
- **Area Chair:** NeurIPS 2024, ECCV 2024, ICLR 2024, CVPR 2024, NeurIPS 2023, ICCV 2023, CVPR 2023, AAAI 2023, ACCV 2023, CVPR 2021, ACCV 2020, ECCV 2020, ICCV 2019, CVPR 2019, ICVGIP 2016
- **Associate Editor:** IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2020-Present
- **Associate Editor:** Journal of Artificial Intelligence Research (JAIR), 2019-2022
- **Associate Editor:** Journal of Visual Communication and Image Representation (JVCI), 2016-2021
- **Best Paper Award Committee:** AAAI 2023
- **Senior Program Committee:** IJCAI 2016, AAAI 2020, IJCAI 2020
- **Conference Program Committee:**
 - CVPR 2011-2018, 2020, 2022 ([Outstanding Reviewer Award CVPR 2015](#))
 - ICCV 2009, 2011, 2013, 2015, 2017 ([Outstanding Reviewer Award ICCV 2013](#))
 - ECCV 2012, 2014, 2018, 2022
 - NeurIPS 2015, 2018, 2019, 2020
 - ICML 2018, 2019
 - ACCV 2014
 - WACV 2021
 - BMVC 2021
- **Journal Reviewer:** International Journal of Computer Vision (IJCV), IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), ACM SIGGRAPH, ACM SIGGRAPH ASIA, ACM Multimedia, Pacific Graphics, Computer Graphics Forum, IEEE Transactions on Human-Machine Systems
- **Workshop Program Committee:** International Workshop on Attention/Intention Understanding (AIU), ACCV 2018; 1st International Workshop on Egocentric Perception, Interaction and Computing (EPIC), ECCV 2016; 4th Workshop on Web-scale Vision and Social Media (VSM), ECCV 2016; 4th International Workshop on Large Scale Visual Recognition and Retrieval (BigVision), CVPR 2016; 1st Workshop on Storytelling with Images and Videos, ECCV 2014; Workshop on Learning From Unlabeled Videos (LUV) CVPR 2020; NewInML NeurIPS 2021; Transformers for Vision (T4V) CVPR 2022
- **Panelist:** NSF 2024, 2022, 2018, 2017, 2016; DOE 2016
- **Chair / Co-Chair:** Demo & Exhibition, CVPR 2019
- **Organizer:** 4th Workshop on Egocentric (First-Person) Vision, CVPR 2016

- **Organizer:** 3rd Workshop on Egocentric (First-Person) Vision, CVPR 2014
- **Organizer:** 5th Workshop on Human Behavior Understanding, ECCV 2014
- **Faculty Advisor:** AggieMentors Outreach Program, UC Davis, 2018-2020
- **Mentor:** AggieMentors Science Fair Mentor, UC Davis, 2017-2019
- **Mentor:** Doctoral Consortium Mentor, ICCV 2021, CVPR 2024
- **Judge:** AggieMentors Science Fair Judge, UC Davis, 2016; RoboPlay C-STEM Video Judge, UC Davis, 2016-2020
- **Organizer:** UT-Austin Vision and Learning Reading Group, 2008-2012
- **UC Davis Committees:**
 - Research and Library Committee (2014-2020)
 - Graduate Admissions Committee (2014-2016, 2017-2018)
 - Faculty Recruitment Committee (NLP 2016-2017, Comp. Neuroscience 2017-2018, ML 2019-2020)
 - Faculty Representative Committee (2018-2020)
 - Data Science Committee (2019-2020)
- **UW–Madison Committees:**
 - Department Vision Committee (2023-Present)
 - Professional Programs Committee (2021-Present)
 - Office of the Vice Chancellor for Research and Graduate Education Research Committee (2022-Present)

ADVISING ACTIVITY

- Fanyi Xiao (PhD student), 2015-2020
 - Thesis: “Multimodal Learning with Minimal Human Supervision from Videos and Natural Language”
 - [UC Davis Best Graduate Researcher in Computer Science Award 2018](#)
 - Next position: Research Scientist at Amazon Research
- Krishna Kumar Singh (PhD student), 2015-2020
 - Thesis: “Understanding and Generating Visual Data with Minimal Human Supervision”
 - [UC Davis Best Graduate Researcher in Computer Science Award \(Honorable Mention\) 2019](#)
 - Next position: Research Scientist at Adobe Research
- Maheen Rashid (PhD student), 2015-2021
 - Thesis: “Towards Automatic Visual Recognition of Horse Pain”
 - Next position: Computer Vision Scientist at Univrse AB
- Xueyan Zou (PhD student), 2018-2024
 - Thesis: “Towards Generalized Multimodal Foundation Models”

- Next position: Postdoc at UCSD
- Utkarsh Ojha (PhD student), 2019-2024
 - Thesis: “Understanding Generative Models and Detecting their Images”
 - Next position: Postdoc at Carnegie Mellon University
- Haotian Liu (PhD student), 2019-2024
 - Thesis: “Steerable Visual Intelligence”
 - Next position: Research Scientist at xAI
- Yuheng Li (PhD student), 2019-2024
 - Thesis: “Controllable Image Generative Models”
 - Next position: Research Scientist at Adobe Research
- Zhuoran Yu (PhD student), 2021-Present
- Mu Cai (PhD student), 2021-Present
- Moniek Smink (Undergraduate student), 2022-Present
 - [Hilldale Research Fellowship 2023-2024](#)
- Zeyi Huang (PhD student), 2022-Present
- Thao Nguyen (PhD student), 2022-Present
- Anirudh Sundara Rajan (MS student), 2023-Present
- Rafael A. Rivera-Soto (MS student), 2020-2021
 - Next position: PhD student at Johns Hopkins University
- Zhongzheng (Jason) Ren (MS student), 2016-2018
 - Next position: PhD student at UIUC
- Wenjian Hu (MS student / Physics PhD), 2017-2018
 - Next position: Research Scientist at Facebook
- Leonardo Ferrer (MS student), 2017-2018
 - Next position: Software Engineer at Google
- Chong Zhou (MS student), 2019-2020
 - Next position: PhD student at Nanyang Technological University
- Yuheng Li (MS student), 2019-2020
 - Next position: PhD student at UC Davis
- Wei-Pang (Tyler) Jan (MS student), 2019-2020
 - Next position: Software Engineer at Amazon
- Antonia Creswell (undergraduate student), 2014-2015
 - Next position: PhD student at Imperial College London
- Xie Zhou (undergraduate student), 2017-2018

- Next position: MS student at UC Berkeley
- Yi Mang (Terry) Yang (undergraduate student), 2017-2018
 - Next position: Software Engineer at Amazon
- Daniel Bolya (undergraduate student), 2018-2019
 - [Chancellor's Award for Excellence in Undergraduate Research \(Honorable Mention\) 2019](#)
 - Next position: PhD student at Georgia Tech
- Aron Sarmasi (undergraduate student), 2018-2019
 - Next position: MS student at UC Davis
- Waiyu Lam (undergraduate student), 2019-2020
 - Next position: MS student at Cornell
- Qi Zhu (undergraduate intern), 2015 summer
 - Next position: MS student at UIUC
- Xiuye Gu (undergraduate intern), 2016 summer
 - Next position: MS student at Stanford
- Haolin Fu (undergraduate intern), 2016 summer
 - Next position: MS student at Yale
- Haotian Liu (undergraduate intern), 2018 summer
 - Next position: PhD student at UC Davis
- Hao Yu (undergraduate intern), 2018 summer
 - Next position: PhD student at Boston University
- Gautam Pradeep (high school student), 2016-2019
 - [1st place, National Environmentors Science Fair 2017](#)
 - [2nd place, National Environmentors Science Fair 2018](#)
 - Next position: undergraduate student at Stanford