Carl Vondrick

Email: vondrick@cs.columbia.edu — Web: http://cs.columbia.edu/~vondrick
Address: 530 West 120th St, Columbia University, New York, NY 10027.

Fields of Specialization

- Computer Vision
- Machine Perception
- Machine Learning

Education

- Ph.D. in Computer Science – Massachusetts Institute of Technology (MIT), 2017.
  Advisor: *Antonio Torralba;* Title: *Predictive Vision.*
  Minor: Cognitive Science.
- M.Sc. in Computer Science – Massachusetts Institute of Technology (MIT), 2013.
  Advisor: *Antonio Torralba*; Title: *Visualizing Object Detection Features.*
- B.Sc. in Computer Science – University of California at Irvine (UCI), 2011.
  Advisor: *Deva Ramanan*; Title: *Crowdsourced Video Annotation.*

Appointments and Employment

- Columbia University
  - YM Associate Professor of Computer Science, September 2023 - now.
  - Associate Professor of Computer Science, January 2023 - August 2023.
  - Assistant Professor of Computer Science, July 2018 - December 2022.
- Cruise, Visiting AI Faculty, February 2022 - 2024.
- Snap, Research Design Consultant, December 2022 - 2023.
- Google, Research Scientist, June 2017 - April 2019.

Awards and Honors

- 2024 Amazon Research Award ($70k).
- 2021 National Science Foundation Early Career Development (CAREER) Award ($550k).
- 2021 Toyota Research Institute Young Faculty Award ($750k).

- 2018 Amazon Research Award ($100k).
- 2018 Best Paper Finalist at CVPR.
- 2015 Google Ph.D. Fellowship in Machine Perception.
- 2011 National Science Foundation Graduate Research Fellowship.

Awards and Honors of Lab Members

- 2024 Apple Ph.D. Fellowship to Tendulkar (2 years).
- 2023 National Science Foundation Graduate Research Fellowship to Geng (3 yrs).
- 2022 Microsoft Ph.D. Fellowship to Surís (2 yrs).
- 2022 National Science Foundation Graduate Research Fellowship to Sudhakar (3 yrs).
- 2022 Paul & Daisy Soros Fellowship to Epstein (1 yr).
- 2021 Amazon CAIT Ph.D. Fellowship to Chiquier (2 yrs).
- 2021 National Science Foundation Graduate Research Fellowship to Mani (3 yrs).
- 2020 National Science Foundation Graduate Research Fellowship to Menon (3 yrs).
- 2020 CRA Honorable Mention to Epstein.

Publications

H-index: 36; Citations: 13,557 (retrieved from Google Scholar on 3/14/24).

Google Scholar: https://scholar.google.com/citations?user=3MzhkFIAAAAJ

The trainees from my group are underlined. The authorship convention in my field is to order by decreasing contribution with the advisor often appearing last. In computer science, the primary publishing venue are conferences, which are ranked higher than journals. The conference H5-index below was retrieved from Google Scholar in February 2022.

Conference Papers (Peer Reviewed)

1. Ege Ozguroglu, Ruoshi Liu, Dídac Surís, Dian Chen, Achal Dave, Pavel Tokmakov, **Carl Vondrick**. "pix2gestalt: Amodal Segmentation by Synthesizing Wholes." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2024.

2. Utkarsh Mall, Cheng Perng Phoo, Meilin Liu, **Carl Vondrick**, Bharath Hariharan, Kavita Bala. "Remote Sensing Vision-Language Foundation Models without Annotations via Ground Remote Alignment." *International Conference on Learning Representations* (ICLR), 2024.

3. Chengzhi Mao, **Carl Vondrick**, Hao Wang, Junfeng Yang. "Raidar: geneRative AI Detection viA Rewriting." *International Conference on Learning Representations* (ICLR), 2024.

4. Haozhe Chen, Junfeng Yang, **Carl Vondrick**, Chengzhi Mao. "Interpreting and Controlling Vision Foundation Models via Text Explanations." *International Conference on Learning Representations* (ICLR), 2024.

5. Rundi Wu, Ruoshi Liu, **Carl Vondrick**, Changxi Zheng. "Sin3DM: Learning a Diffusion Model from a Single 3D Textured Shape." *International Conference on Learning Representations* (ICLR), 2024.

6. Dídac Surís, Sachit Menon, **Carl Vondrick**. "ViperGPT: Visual Inference via Python Execution for Reasoning." *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), to appear, 2023.
   Oral Presentation (5% acceptance rate)
   Conference H5-index: 184

7. Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, **Carl Vondrick**. "Zero-1-to-3: Zero-shot One Image to 3D Object." *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), to appear, 2023.
   Conference H5-index: 184

8. Arjun Mani, Ishaan Preetam Chandratreya, **Carl Vondrick**, Richard Zemel. "SURFSUP: Learning Fluid Simulation for Novel Surfaces." *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), to appear, 2023.
   Conference H5-index: 184

9. Mia Chiquier, **Carl Vondrick**. "Muscles in Action." *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), to appear, 2023.
   Conference H5-index: 184

10. Ruoshi Liu, Chengzhi Mao, Purva Tendulkar, Hao Wang, **Carl Vondrick**. "Landscape Learning for Neural Network Inversion." *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), to appear, 2023.
    Conference H5-index: 184

11. Hongge Chen, Zhao Cheng, Greg Meyer, Dennis Park, **Carl Vondrick,** Ashish Shrivastava, Yuning Chai. "SHIFT3D: Synthesizing Hard Inputs for Tracking 3D Detectors." *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), to appear, 2023.
    Conference H5-index: 184

12. Ruoshi Liu, **Carl Vondrick**. "Humans as Light Bulbs: 3D Scene Reconstruction with Thermal Reflection." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), to appear, 2023.
    Conference H5-index: 356

13. Ruoshi Liu, Sachit Menon, Chengzhi Mao, Dennis Park, Simon Stent, **Carl Vondrick**. "What You Can Reconstruct from a Shadow." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), to appear, 2023.
    Conference H5-index: 356

14. Basile Van Hoorick, Pavel Tokmakov, Simon Stent, Jie Li, **Carl Vondrick**. "Tracking through Containers and Occluders in the Wild." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), to appear, 2023.
    Conference H5-index: 356

15. Purva Tendulkar, Didac Suris, **Carl Vondrick**. "FLEX: Full-Body Grasping Without Full-Body Grasps." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), to appear, 2023.
    Conference H5-index: 356

16. Chengzhi Mao, Revant Teotia, Amrutha Sundar, Sachit Menon, Junfeng Yang, Xin Wang, **Carl Vondrick**. "Doubly Right Object Recognition: A Why Prompt for Visual Rationales." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), to appear, 2023.
    Conference H5-index: 356

17. Chengzhi Mao, Lingyu Zhang, Abhishek Vaibhav Joshi, Hao Wang, Junfeng Yang, **Carl Vondrick**. "Robust Perception through Equivariance." *Proceedings of the International Conference on Machine Learning* (ICML)*,* to appear, 2023.
    Conference H5-index: 237

18. Sachit Menon, **Carl Vondrick**. "Visual Classification via Description from Large Language Models." *Proceedings of the International Conference on Learning Representations* (ICLR), 2023.
    Oral Presentation (5% acceptance rate)
    Conference H5-index: 253

19. Chengzhi Mao, Scott Geng, Junfeng Yang, Xin Wang, **Carl Vondrick**. "Understanding Zero-shot Adversarial Robustness for Large-Scale Models." *Proceedings of the International Conference on Learning Representations* (ICLR), 2023.
    Conference H5-index: 253

20. Hui Lu, Mia Chiquier, **Carl Vondrick**. "Private Multiparty Perception for Navigation." *Proceedings for Neural Information Processing Systems* (NeurIPS), pp. 3318-3328, 2022.
Conference H5-index: 245

21. Dídac Surís, **Carl Vondrick**. "Representing Spatial Trajectories as Distributions." *Proceedings for Neural Information Processing Systems* (NeurIPS), pp. 13731-13744, 2022.
Conference H5-index: 245

22. Sachit Menon, David Blei, **Carl Vondrick**. "Forget-me-not! Contrastive Critics for Mitigating Posterior Collapse." *Proceedings for the Conference on Uncertainty in Artificial Intelligence* (UAI), pp. 1360-1370, 2022.

23. Basile Van Hoorick, Purva Tendulkar, Dídac Surís, Dennis Park, Simon Stent, **Carl Vondrick**. "Revealing Occlusions with 4D Neural Fields." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 3011-3021, 2022.
Oral Presentation (3% acceptance rate)
Conference H5-index: 356

24. Dídac Surís, Dave Epstein, **Carl Vondrick**. "Globetrotter: Connecting Languages by Connecting Images." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 16474-16484, 2022.
Oral Presentation (3% acceptance rate)
Conference H5-index: 356

25. Dídac Surís, **Carl Vondrick**, Bryan Russell, Justin Salamon. "It's Time for Artistic Correspondence in Music and Video." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 10564-10574, 2022.
Conference H5-index: 356

26. Chengzhi Mao, Kevin Xia, James Wang, Hao Wang, Junfeng Wang, Elias Bareinboim, **Carl Vondrick**. "Causal Transportability for Visual Recognition." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 7521-7531, 2022.
Conference H5-index: 356

27. Will Price, **Carl Vondrick**, Dima Damen. "UnweaveNet: Unweaving Activity Stories." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 13770-13779, 2022.
Conference H5-index: 356

28. Mia Chiquier, Chengzhi Mao, **Carl Vondrick**. "Real-time Neural Voice Camouflage." *Proceedings of the International Conference on Learning Representations* (ICLR), 2022.
Oral Presentation (1% acceptance rate)
Conference H5-index: 253

29. Chengzhi Mao, Lu Jiang, Mostafa Dehghani, **Carl Vondrick**, Rahul Sukthankar, Irfan Essa. "Discrete Representations Strengthen Vision Transformer Robustness." *Proceedings of the International Conference on Learning Representations* (ICLR), 2022.
Conference H5-index: 253

30. Xingyu Fu, Ben Zhou, Ishaan Chandratreya, **Carl Vondrick**, Dan Roth. "There's a Time and Place for Reasoning Beyond the Image." *Proceedings of the Association for Computational Linguistics* (ACL), pp. 1138–1149*,* 2022.
Conference H5-index: 157

31. Boyuan Chen, Mia Chiquier, Hod Lipson, **Carl Vondrick**. "The Boombox: Visual Reconstruction from Acoustic Vibrations." *Proceedings of the Conference on Robot Learning* (CoRL), pp. 1067-1077, 2021.

32. Chengzhi Mao, Mia Chiquier, Hao Wang, Junfeng Yang, **Carl Vondrick**. "Adversarial Attacks are Reversible with Natural Supervision." *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), pp. 661-671, 2021.
Conference H5-index: 184

33. Basile Van Hoorick, **Carl Vondrick**. "Dissecting Image Crops." *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), pp. 9741-9750, 2021.
Conference H5-index: 184

34. Dídac Surís, Ruoshi Liu, **Carl Vondrick**. "Learning the Predictability of the Future." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 12607-12617, 2021.
Conference H5-index 356

35. Chengzhi Mao, Amogh Gupta, Augustine Cha, Hao Wang, Junfeng Yang, **Carl Vondrick**. "Generative Interventions for Causal Learning." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 3947-3956, 2021.
Conference H5-index: 356

36. Dave Epstein, **Carl Vondrick**. "Learning Goals from Failure." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 11194-11204, 2021.
Conference H5-index: 356

37. Ruilin Xu, Rundi Wu, Yuko Ishiwaka, **Carl Vondrick,** Changxi Zheng. "Listening to Sounds of Silence for Speech Denoising." *Proceedings for the Advances in the Neural Information Processing Systems* (NeurIPS), pp. 9633-9648, 2020.
Conference H5-index: 245

38. Chengzhi Mao, Amogh Gupta, Vikram Nitin, Baishakhi Ray, Shuran Song, Junfeng Yang, **Carl Vondrick**. "Multitask Learning Strengthens Adversarial Robustness." *Proceedings of the European Conference on Computer Vision* (ECCV), pp. 158-174, 2020.
Oral Presentation (2% acceptance rate)
Conference H5-index: 197

39. Dídac Surís, Dave Epstein, Heng Ji, Shih-Fu Chang, **Carl Vondrick**. "Learning to Learn Words from Visual Scenes." *Proceedings of the European Conference on Computer Vision* (ECCV), pp. 434-452, 2020.
Conference H5-index: 197

40. Alex Andonian, Camilo Fosco, Mathew Monfort, Allen Lee, Rogerio Feris, **Carl Vondrick**, Aude Oliva. "We Have So Much In Common: Modeling Semantic Relational Set Abstractions in Videos." *Proceedings of the European Conference on Computer Vision* (ECCV), pp. 18-34, 2020.
Conference H5-index: 197

41. Dave Epstein, Boyuan Chen, **Carl Vondrick**. "Oops! Predicting Unintentional Action in Video." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 919-929, 2020.
Conference H5-index: 365

42. Chengzhi Mao, Ziyuan Zhong, Junfeng Yang, **Carl Vondrick**, Baishakhi Ray. "Metric Learning for Adversarial Robustness." *Proceedings for the Advances in the Neural Information Processing Systems* (NeurIPS), pp. 480-491, 2019.
Conference H5-index: 245

43. Chen Sun, Austin Myers, **Carl Vondrick**, Kevin Murphy, Cordelia Schmid. "VideoBERT: A Joint Model for Video and Language Representation Learning." *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), pp. 7464-7473, 2019.
Conference H5-index: 184

44. Chen Sun, Abhinav Shrivastava, **Carl Vondrick**, Rahul Sukthankar, Kevin Murphy, Cordelia Schmid. "Relational Action Forecasting." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 273-283, 2019.
Conference H5-index: 365

45. Hassan Akbari, Svebor Karaman, Surabhi Bhargava, Brian Chen, **Carl Vondrick**, Shih-Fu Chang. "Multi-level Multimodal Common Semantic Space for Image-Phrase Grounding." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 12476-12486, 2019.
Conference H5-index: 365

46. **Carl Vondrick**, Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, Kevin Murphy. "Tracking Emerges by Colorizing Videos." *Proceedings of the European Conference on Computer Vision* (ECCV), pp. 391-408, 2018.
Conference H5-index: 197

47. Hang Zhao, Chuang Gan, Andrew Rouditchenko, **Carl Vondrick**, Josh McDermott, Antonio Torralba. "The Sound of Pixels." *Proceedings of the European Conference on Computer Vision* (ECCV), pp. 570-586, 2018.
Conference H5-index: 197

48. Chen Sun, Abhinav Shrivastava, **Carl Vondrick**, Kevin Murphy, Rahul Sukthankar, Cordelia Schmid. "Actor-centric Relation Network." *Proceedings of the European Conference on Computer Vision* (ECCV), pp. 318-334, 2018.
Conference H5-index: 197

49. Chunhui Gu, Chen Sun, David A. Ross, **Carl Vondrick**, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, Jitendra Malik. "AVA: A Video Dataset of Spatio-temporally Localized Atomic Visual Actions." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 6047-6056, 2018.
Conference H5-index: 365

50. Adria Recasens, **Carl Vondrick**, Aditya Khosla, Antonio Torralba. "Following Gaze in Video." *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), pp. 1435-1443, 2017.
Conference H5-index: 184

51. **Carl Vondrick**, Antonio Torralba. "Generating the Future with Adversarial Transformers." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 1020-1028, 2017.
Conference H5-index: 365

52. **Carl Vondrick**, Hamed Pirsiavash, Antonio Torralba. "Generating Videos with Scene Dynamics." *Proceedings for the Advances in the Neural Information Processing Systems* (NeurIPS), pp. 613-621, 2016.
Conference H5-index: 245

53. **Carl Vondrick**, Yusuf Aytar, Antonio Torralba. "SoundNet: Learning Sound Representations from Unlabeled Video." *Proceedings for the Advances in the Neural Information Processing Systems* (NeurIPS), pp. 892-900, 2016.
    Conference H5-index: 245

54. **Carl Vondrick**, Hamed Pirsiavash, Antonio Torralba. "Anticipating Visual Representations with Unlabeled Video." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 98-106, 2016.
    Conference H5-index: 356

55. **Carl Vondrick**, Deniz Oktay, Hamed Pirsiavash, Antonio Torralba. "Predicting Motivations Behind Actions by Leveraging Text." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 2997-3005, 2016.
    Conference H5-index: 356

56. Lluis Castrejon, Yusuf Aytar, **Carl Vondrick**, Hamed Pirsiavash, Antonio Torralba. "Learning Aligned Cross-Modal Representations from Weakly Aligned Data." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 2940-2949, 2016.
    Conference H5-index: 356

57. **Carl Vondrick**, Hamed Pirsiavash, Aude Oliva, Antonio Torralba. "Learning Visual Biases from Human Imagination." *Proceedings for the Advances in the Neural Information Processing Systems* (NeurIPS), pp. 289-297, 2015.
    Conference H5-index: 245

58. Adria Recasens, Aditya Khosla, **Carl Vondrick**, Antonio Torralba. "Where are they looking?" *Proceedings for the Advances in the Neural Information Processing Systems* (NeurIPS), pp. 199-207, 2015.
    Conference H5-index: 245

59. Hamed Pirsiavash, **Carl Vondrick**, Antonio Torralba. "Assessing the Quality of Actions." *Proceedings of the European Conference on Computer Vision* (ECCV), pp. 556-571, 2014.
    Conference H5-index: 197

60. **Carl Vondrick**, Aditya Khosla, Tomasz Malisiewicz, Antonio Torralba. "HOGgles: Visualizing Object Detection Features." *Proceedings of the IEEE International Conference on Computer Vision* (ICCV), pp. 1-8, 2013.
    Oral Presentation (3% acceptance rate)
    Conference H5-index: 184

61. Xiangxin Zhu, **Carl Vondrick**, Deva Ramanan, Charless Fowlkes. "Do We Need More Training Data or Better Models for Object Detection?" *Proceedings of the British Machine Vision Conference* (BMVC), vol. 3, no. 5, 2012.

62. **Carl Vondrick**, Deva Ramanan. "Video Annotation and Tracking with Active Learning." *Proceedings for the Advances in the Neural Information Processing Systems* (NeurIPS), pp. 28-36, 2011.
Conference H5-index: 245

63. Sangmin Oh, Anthony Hoogs, Amitha Perera, Naresh Cuntoor, Chia-Chih Chen, Jong Taek Lee, Saurajit Mukherjee, J. K. Aggarwal, Hyungtae Lee, Larry Davis, Eran Swears, Xioyang Wang, Qiang Ji, Kishore Reddy, Mubarak Shah, **Carl Vondrick**, Hamed Pirsiavash, Deva Ramanan, Jenny Yuen, Antonio Torralba, Bi Song, Anesco Fong, Amit Roy-Chowdhury, Mita Desai. "A Large-scale Benchmark Dataset for Event Recognition in Surveillance Video." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pp. 3153-3160, 2011.
Conference H5-index: 356

64. **Carl Vondrick**, Deva Ramanan, Donald Patterson. "Efficiently Scaling Up Video Annotation with Crowdsourced Marketplaces." *Proceedings of the European Conference on Computer Vision* (ECCV), pp. 610-623, 2010.
Conference H5-index: 197

Journal Papers (Peer Reviewed)

65. Boyuan Chen, Robert Kwiatkowski, **Carl Vondrick**, Hod Lipson. "Full-Body Visual Self-Modeling of Robot Morphologies," *Science Robotics*, vol. 8, issue 68, 2022.

66. Boyuan Chen, **Carl Vondrick**, Hod Lipson. "Visual Behavior Modelling for Robotic Theory of Mind." *Nature Scientific Reports*, vol. 11, pp. 1-14, 2021.
Journal H5-index: 200

67. Mathew Monfort, **Carl Vondrick**, Aude Oliva, Alex Andonian, Bolei Zhou, Kandan Ramakrishnan, Sarah Adel Bargal, Tom Yan, Lisa M. Brown, Quanfu Fan, Dan Gutfreund. "Moments in Time Dataset: one million videos for event understanding." *Transactions on Pattern Analysis and Machine Intelligence* (PAMI), vol. 42, pp. 502-508, 2019.
Journal H5-index: 149

68. Yusuf Aytar, Lluis Castrejon, **Carl Vondrick**, Hamed Pirsiavash, Antonio Torralba. "Cross-Modal Scene Networks." *Transactions on Pattern Analysis and Machine Intelligence* (PAMI), vol. 40, pp. 2303-2314, 2017.
Journal H5-index: 149

69. **Carl Vondrick**, Aditya Khosla, Hamed Pirsiavash, Tomasz Malisiewicz, Antonio Torralba. "Visualizing Object Detection Features." *International Journal of Computer Vision* (IJCV), vol. 119, pp. 145-158, 2016.

70. Xiangxin Zhu, **Carl Vondrick**, Charless Fowlkes, Deva Ramanan. "Do we need more training data?" *International Journal of Computer Vision* (IJCV), vol. 119, pp. 76-92, 2015.

71. **Carl Vondrick**, Donald Patterson, Deva Ramanan. "Efficiently Scaling Up Crowdsourced Video Annotation." *International Journal of Computer Vision* (IJCV), vol. 101, pp. 184-204, 2012.

Teaching

Semester	Course	#Responses/ Enrollment	Course Quality†	Instructor Quality†
Spring 2024	COMS W4732 Computer Vision II	-	-	-
Spring 2023	COMS W4732 Computer Vision II	119 / 159	4.35	4.55
Fall 2022	COMS E6998 Rep. Learning	43 / 44	4.79	4.79
Spring 2022	COMS W4732 Computer Vision II	315 / 472	4.30	4.45
Fall 2021	COMS E6998 Rep. Learning	11 / 32	4.69	4.77
Sum. 2021‡	COMS W4732 Computer Vision II	58 / 91	4.19	4.50
Fall 2020	COMS E6998 Rep. Learning	25 / 26	4.96	4.96
Fall 2019	COMS W4731 Computer Vision I	119 / 209	4.11	4.19
Spring 2019	COMS E6998 Adv. Computer Vision	11 / 31	4.27	4.55
Fall 2018	COMS W4731 Computer Vision I	67 / 110	4.27	4.45

† Scale: 5 excellent, 4 very good, 3 good, 2 fair, 1 poor.
‡ Due to the COVID-19 pandemic, the spring 2021 course was rescheduled to summer 2021.

Additional Teaching and Tutorials

- Lecture on AI, Computer Vision, and Machine Leraning.
  Columbia+ and Columbia Engineering Executive Education. École Polytechnique. 2023.
- Lecture on Computational Imaging and Vision
  Dual MBA/Executive MS: Engineering and Applied Science. 2022, 2023.
  Co-taught with Shree Nayar.

- Computer Vision and Robotics.
  The Online AI Program from Columbia Engineering. 2022.
  Co-taught with Hod Lipson.
- Tutorial on Unsupervised Visual Learning.
  International Computer Vision Summer School (ICVSS). 2018.
- Tutorial on Unsupervised Visual Learning.
  IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2018.
  Co-taught with Anelia Angelova and Pierre Sermanet.

Mentoring and Team

Postdoctoral Researchers

- Utkarsh Mall, Fall 2023 - present
- Chengzhi Mao, Fall 2023 - present (until he starts his faculty position at McGill)

PhDs Completed

- Chengzhi Mao, Fall 2018 - Summer 2023
  Title: Robust Machine Learning via Integrating Context
  Committee: Richard Zemel (moderator), David Blei, Hao Wang, Junfeng Yang, Carl Vondrick
  Next Position: Assistant Professor of Computer Engineering, McGill University
  Co-advised with Junfeng Yang

PhDs in Progress

- Dídac Surís, Fall 2019 - Spring 2024 (expected)
  Title (tentative): Geometric Embeddings for Visual Prediction
  2022 Microsoft Ph.D. Fellowship
- Mia Chiquier, Fall 2020 - Spring 2025 (expected)
  2021 Amazon CAIT Ph.D. Fellowship
- Sachit Menon, Fall 2020 - Spring 2025 (expected)
  National Science Foundation Graduate Research Fellowship (NSF GRFP)
  Columbia University Presidential Fellowship
- Basile Van Hoorick, Spring 2021 - Spring 2025 (expected)
- Ruoshi Liu, Fall 2021 - Spring 2026 (expected)

- Arjun Mani, Fall 2021 - Spring 2026 (expected)
  National Science Foundation Graduate Research Fellowship (NSF GRFP)
  Co-advised with Richard Zemel

- Purva Tendulkar, Fall 2021 - Spring 2026 (expected)
  Columbia University Presidential Fellowship
  2024 Apple Ph.D. Fellowship

- Sruthi Sudhakar, Fall 2022 - Spring 2027 (expected)
  National Science Foundation Graduate Research Fellowship (NSF GRFP)
  Co-advised with Richard Zemel

Undergraduates

- Hui Lu, Fall 2019 - 2022
  Topic: Secure Multiparty Visual Recognition
  Next Position: Facebook

- Ishaan Chandratreya, Fall 2019 - Spring 2023
  Topic: Predictive Models for Robotics
  Next Position: PhD in Computer Science at Massachusetts Institute of Technology

- Scott Geng, Spring 2022 - Spring 2023
  Topic: Affective Human Activity Understanding
  National Science Foundation Graduate Research Fellowship (NSF GRFP)
  Next Position: PhD in Computer Science at University of Washington

- Jillian Ross, Fall 2020 - Spring 2021
  Topic: Visually Grounded Language Models
  Next Position: PhD in Computer Science at Massachusetts Institute of Technology

- Ruoshi Liu, Fall 2019 - Spring 2021
  Topic: Hyperbolic Machine Learning
  Next Position: PhD in Computer Science at Columbia University

- Dave Epstein, Fall 2018 - Spring 2020
  Topic: Recognizing Human Activity and Goals from Video
  Next Position: PhD in Computer Science at University of California, Berkeley
  CRA Honorable Mention
  Paul & Daisy Soros Ph.D Fellowship

Masters

- Revant Teotia. Fall 2021 - now.
  Topic: Explainable Visual Recognition.
  Next Position: PhD in Computer Science at New York University.

- Amogh Gupta, Fall 2019 - Spring 2021
  Topic: Robust Visual Representations
  Next Position: Amazon Research

Ph.D. Committee (excluding my own students)

- Alireza Zareian (advisor: Shih-Fu Chang, defense: January 2021)
- Iretiayo Akinola (advisor: Peter Allen, defense: March 2021)
- Chang Xiao (advisor: Changxi Zheng, defense: April 2021)
- Yuchi Tian (advisor: Baishakhi Ray, defense: July 2021)
- Boxi Xian (advisor: Hod Lipson, defense: October 2021)
- Hassan Akbari (advisor: Shih-Fu Chang, defense: January 2022)
- Boyuan Chen (advisor: Hod Lipson, defense: April 2022)
- Rob Kwiatkowski (advisor: Hod Lipson, defense: May 2022)
- Simone Fobi (advisor: Vijay Modi, defense: June 2022)
- Terry Conlon (advisor: Vijay Modi, defense: June 2022)
- Sai Saketh Rambhatla (advisor: Abhinav Shrivastava, defense: December 2022)
- Emily Hannigan (advisor: Matei Ciocarlie, defense: February 2023)
- Brian Chen (advisor: Shih-Fu Chang, defense: February 2023)
- Xudong Lin (advisor: Shih-Fu Chang, defense: November 2023)

Patents

Granted

1. Visual Tracking by Colorization.
   Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, Kevin Murphy, **Carl Vondrick**.
   Date Granted: December 2021.

Pending

2. System and Method for 3D Reconstruction from Shadows.
   Ruoshi Liu, Sachit Menon, Chengzhi Mao, Dennis Park, Simon Stent, **Carl Vondrick**.

3. A Neural Network Model for Speech Denoising.
   Changxi Zheng, Ruilin Xu, Rundi Wu, **Carl Vondrick**, Yuko Ishikawa.

4. Action Localization using Relational Features.
   Chen Sun, Abhinav Shrivastava, Cordelia Schmid, Rahul Sukthankar, Kevin Murphy, **Carl Vondrick**.

Invited Talks and Visiting Lectures

- Multimodal Learning from Pixels to People
  - Stanford University, March 2024
  - Cornell, November 2023
  - Cornell Tech, November 2023
  - Carnegie Mellon University, November 2023
- The Rise of Visual Skills in Large Models
  - Stanford University, May 2023
  - University of California, Berkeley, June 2023
  - T4V: Transformers for Vision, at CVPR, June 2023
  - Holistic Video Understanding Workshop, at CVPR, June 2023
  - GeoNet: Unsupervised Adaptation across Geographies, at ICCV, October 2023
  - International Challenge on Compositional and Multimodal Perception, at ICCV, October 2023
  - BigMAC: Big Model Adaptation for Computer Vision, at ICCV, October 2023
  - Workshop on Large-scale Video Object Segmentation, at ICCV, October 2023
- Emergent 3D Vision

- Visual Object Tracking Challenge, at ICCV, October 2023
- Frontiers of Monocular 3D Perception, at ICCV, October 2023

- Connecting Vision, Language, and Code
  - Data Science Day, at Columbia University, April 2023
- Visual Recognition by Reading
  - Workshop on Out of Distribution Generalization, at ECCV, October 2022
  - Workshop on Compositional and Multimodal Perception, at ECCV, October 2022
- Audio Privacy
  - Workshop on Sound and Vision, at CVPR, June 2022
- Inverting the Neural Network
  - Workshop on Visual Perception and Learning in an Open World, at CVPR, June 2022
  - Massachusetts Institute of Technology, March 2022
  - Workshop on Dealing with the Novelty in Open Worlds, at WACV, December 2021
- Learning the Predictability of the Future
  - Samsung AI Research, July 2022
  - AiBee Research, July 2022
  - Workshop on Robustness in Sequential Data, at CVPR, June 2022
  - Google Research, February 2022
  - Workshop on Applications of Signal Processing to Audio and Acoustics, October, 2021
  - Cruise, September 2021
  - Stanford University, February 2021
  - Massachusetts Institute of Technology, February 2021
  - Princeton University, February 2021
  - International Business Machines (IBM), Feb 2021
- Learning from Unlabeled Video
  - New York University, April 2020
  - Carnegie Mellon University, March 2020
  - University of Pittsburgh, March 2020
  - University of Massachusetts, Amherst, November 2019
  - Butterfly Network, November 2018

- University of Maryland, College Park, March 2018

- Predictive Vision
  - University of Pennsylvania, November 2017
  - Snapchat Research, November 2017
  - University of Southern California, November 2017
  - Workshop on Video Frontiers, November 2017
  - Rework Summit, May 2017
  - University of California, San Deigo, April 2017
  - Cornell University, April 2017
  - University of Texas, Austin, March 2017
  - Columbia University, March 2017
  - Google Research, March 2017
  - Adobe Research, March 2017
  - OpenAI, March 2017
  - Brown University, February 2017
  - University of California, Los Angeles, February 2017
  - NVidia, February 2017
  - Rework Summit, November 2016
  - Twitter, October 2016
  - TTI Chicago, September 2016
  - Massachusetts Institute of Technology, September 2016
  - Apple, August 2016
  - University of California, Berkeley, August 2016
  - Stanford University, August 2016
  - Boston University, March 2016
  - University of Massachusetts, Boston, March 2016
- Visualizing Object Detection Features
  - University of Massachusetts, Boston, March 2016
  - Massachusetts Institute of Technology, September 2015
  - Brown University, November 2013

- Efficient Video Annotation
  - CVPR Workshop, June 2013
  - CVPR Workshop, June 2011

Service

Department Service

- Hiring Committee, Computer Science
  - Member: 2023 - now
- PhD Admissions Committee, Computer Science
  - Vice Chair: 2021 - 2023
  - Member: 2019 - 2021
- Distinguished Lectures Committee, Computer Science
  - Member: 2018 - now

University Service

- Reviewer, Adhoc Committees for Fellowship Selection: 2019, 2021
- Speaker, Data Science Institute Council Annual Meeting: 2021
- Speaker, NIH Workshop at Mailman School of Public Health: 2021

Service to Discipline

- Area Chair and Senior Program Committee
  - ECCV: 2022
  - CVPR: 2018, 2019, 2021, 2022, 2024
  - ICCV: 2021, 2023
  - NeurIPS: 2019, 2020, 2023
  - ICLR: 2020
- Review Panelist, National Science Foundation
  - NRI: 2020
  - IIS: 2020, 2021, 2022, 2023
  - SBIR: 2020
- Reviewer (Conferences and Journals)
  - CVPR: 2011 - 2023

  - ICCV: 2011 - 2022
  - ECCV: 2011 - 2020
  - NeurIPS: 2015 - 2020
  - ICML: 2018
  - IJCV: 2014 - now
  - TPAMI: 2015 - now
- Organizer, Workshop on Open-world Vision, CVPR: 2021, 2022
- Organizer, Workshop on Self-supervised Learning
  - ICML: 2019
  - CVPR: 2018
- Organizer, Workshop on Learning from Unlabeled Video, CVPR: 2019, 2020, 2021

Outreach

- Teaching a Machine to See.
  Artificial Intelligence for All (AI4All), Summer 2019.
  High-school summer camp for New York City metropolitan area.

Media Coverage

Television and Radio

- NPR, Algorithms Identify Audio through Video Footage
- NPR, Computer Binge-Watched TV And Learned To Predict
- CNN, New AI Can Predict When Two People Will Kiss
- CBC, Teaching Software to Predict Handshakes, Hugs, and Kisses
- Stephen Colbert, Television clip on human action prediction

Newspapers and Magazines

- Science, Is Technology Spying on You? New AI Could Prevent Eavesdropping
- Associated Press, How Do You Teach Human Interaction to a Robot? Lots of TV
- NBC, Deep Learning: Teaching Computers to Predict the Future
- Newsweek, Artificial Intelligence Algorithms Predicts the Future
- Forbes, MIT Computers Binge-Watch To Learn About Hugs
- ABC News, New AI Can Predict When Two People Will Kiss

- Fox News, New Artificial Intelligence Can Predict When You Will Kiss
- Wired, This AI learned to predict the future by watching loads of TV
- Popular Science, Algorithm Binge Watches TV to Predict Human Behavior
- Scientific American, Artificial Intelligence Can Predict How Scenes Will Play Out
- New Scientist, Binge-watching videos teaches computers to recognise sounds
- New Scientist, AI learns to predict the future by watching 2 million videos
- Vice Magazine, This Algorithm Taught Itself to Animate a Still Photo
- The Verge, Machine Learning's Next Trick is Generating Videos from Photos
- The Week Junior, A machine that learns by listening (children's magazine)
- Technology Review, Image Experiment Reveals The Building Blocks of Imagination