

Carl Vondrick

Assistant Professor
Department of Computer Science
Columbia University

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

Appointments

- | | | |
|------------------------|---------------------|---------------------|
| 1. Columbia University | Assistant Professor | Jul 2018 - Present |
| 2. Google Research | Research Scientist | Jul 2017 - Jul 2018 |

Education

- | | | |
|--|---------------------------|---------------------|
| 1. Massachusetts Institute of Technology | Ph.D. in Computer Science | Sep 2013 - Jun 2017 |
| 2. Massachusetts Institute of Technology | M.S. in Computer Science | Sep 2011 - Jun 2013 |
| 3. University of California, Irvine | B.S. in Computer Science | Sep 2008 - Jun 2011 |

Journal Publications (peer reviewed)

1. Moments in Time Dataset: one million videos for event understanding
Mathew Monfort, et al.
Transactions on Pattern Analysis and Machine Intelligence (PAMI). 2019.
2. Cross-Modal Scene Networks.
Yusuf Aytar, Lluís Castrejon, Carl Vondrick, Hamed Pirsiavash, Antonio Torralba.
Transactions on Pattern Analysis and Machine Intelligence (PAMI). 2017.
3. Visualizing Object Detection Features.
Carl Vondrick, Aditya Khosla, Hamed Pirsiavash, Tomasz Malisiewicz, Antonio Torralba.
International Journal of Computer Vision (IJCV). 2016.
4. Do we need more training data?
Xiangxin Zhu, Carl Vondrick, Charless Fowlkes, Deva Ramanan.
International Journal of Computer Vision (IJCV). 2015.
5. Efficiently Scaling Up Crowdsourced Video Annotation.
Carl Vondrick, Donald Patterson, Deva Ramanan.
International Journal of Computer Vision (IJCV). 2012.

Conference Publications (peer reviewed)

1. Metric Learning for Adversarial Robustness
Chengzhi Mao, Ziyuan Zhong, Junfeng Yang, Carl Vondrick, Baishakhi Ray
Neural Information Processing Systems (NeurIPS). 2019.
2. VideoBERT: A Joint Model for Video and Language Representation Learning.
Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, Cordelia Schmid
International Conference on Computer Vision (ICCV). 2019.
3. Relational Action Forecasting
Chen Sun, Abhinav Shrivastava, Carl Vondrick, Rahul Sukthankar, Kevin Murphy, Cordelia

- Schmid.
Computer Vision and Pattern Recognition (CVPR). 2019.
4. Multi-level Multimodal Common Semantic Space for Image-Phrase Grounding.
Hassan Akbari, Svebor Karaman, Surabhi Bhargava, Brian Chen, Carl Vondrick, Shih-Fu Chang.
Computer Vision and Pattern Recognition (CVPR) 2019.
 5. Tracking Emerges by Colorizing Videos.
Carl Vondrick, Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, Kevin Murphy.
European Conference on Computer Vision (ECCV). 2018.
 6. The Sound of Pixels.
Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, Antonio Torralba.
European Conference on Computer Vision (ECCV). 2018.
 7. Actor-centric Relation Network.
Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, Cordelia Schmid.
European Conference on Computer Vision (ECCV). 2018.
 8. AVA: A Video Dataset of Spatio-temporally Localized Atomic Visual Actions.
Chunhui Gu, et al.
Computer Vision and Pattern Recognition (CVPR). 2018.
 9. Following Gaze in Video.
Adria Recasens, Carl Vondrick, Aditya Khosla, Antonio Torralba.
International Conference on Computer Vision (ICCV). 2017.
 10. Generating the Future with Adversarial Transformers.
Carl Vondrick, Antonio Torralba.
Computer Vision and Pattern Recognition (CVPR). 2017.
 11. Generating Videos with Scene Dynamics.
Carl Vondrick, Hamed Pirsiavash, Antonio Torralba.
Neural Information Processing Systems (NIPS). 2016.
 12. SoundNet: Learning Sound Representations from Unlabeled Video.
Carl Vondrick, Yusuf Aytar, Antonio Torralba.
Neural Information Processing Systems (NIPS). 2016.
 13. Anticipating Visual Representations with Unlabeled Video.
Carl Vondrick, Hamed Pirsiavash, Antonio Torralba.
Computer Vision and Pattern Recognition (CVPR). 2016.
 14. Predicting Motivations Behind Actions by Leveraging Text.
Carl Vondrick, Deniz Oktay, Hamed Pirsiavash, Antonio Torralba.
Computer Vision and Pattern Recognition (CVPR). 2016.

15. Learning Aligned Cross-Modal Representations from Weakly Aligned Data.
Lluís Castrejon, Yusuf Aytar, Carl Vondrick, Hamed Pirsiavash, Antonio Torralba.
Computer Vision and Pattern Recognition (CVPR). 2016.
16. Learning Visual Biases from Human Imagination.
Carl Vondrick, Hamed Pirsiavash, Aude Oliva, Antonio Torralba.
Neural Information Processing Systems (NIPS). 2015.
17. Where are they looking?
Adria Recasens, Aditya Khosla, Carl Vondrick, Antonio Torralba.
Neural Information Processing Systems (NIPS). 2015.
18. Assessing the Quality of Actions.
Hamed Pirsiavash, Carl Vondrick, Antonio Torralba.
European Conference on Computer Vision (ECCV). 2014.
19. HOGgles: Visualizing Object Detection Features.
Carl Vondrick, Aditya Khosla, Tomasz Malisiewicz, Antonio Torralba.
International Conference on Computer Vision (ICCV). 2013.
20. Do We Need More Training Data or Better Models for Object Detection?
Xiangxin Zhu, Carl Vondrick, Deva Ramanan, Charles Fowlkes.
British Machine Vision Conference (BMVC). 2012.
21. Video Annotation and Tracking with Active Learning.
Carl Vondrick and Deva Ramanan.
Neural Information Processing Systems (NIPS). 2011.
22. A Large-scale Benchmark Dataset for Event Recognition in Surveillance Video.
Sangmin Oh, et al.
Computer Vision and Pattern Recognition (CVPR). 2011.
23. Efficiently Scaling Up Video Annotation with Crowdsourced Marketplaces.
Carl Vondrick, Deva Ramanan, Donald Patterson.
European Conference on Computer Vision (ECCV). 2010.

Technical Reports

1. See, Hear, and Read: Deep Aligned Representations
Yusuf Aytar, Carl Vondrick, Antonio Torralba.
arXiv. 2017.

Grants Awarded

- | | |
|--|-----------|
| 1. Amazon Research Gift
Anticipating Human Goals in Video
Principal Investigator, \$100k | 2018-2019 |
| 2. National Science Foundation, CRII
Learning Predictive Representations from Unlabeled Video
Principal Investigator, \$175k | 2019-2021 |

3. National Science Foundation, NRI 2019-2022
Learning Visual Dynamics from Interaction
Principal Investigator, \$750k
4. Defense Advanced Research Projects Agency, GAILA 2019-2020
Learning Visually Grounded Language Models
Principal Investigator, \$1M

Selected Awards and Honors

1. Best Paper Finalist at CVPR 2019
2. Google Ph.D Fellowship in Machine Perception 2015 - 2017
3. National Science Foundation Graduate Fellowship 2011 - 2014
4. Outstanding Reviewer Award for ECCV, CVPR 2015 - 2016
5. UCI Chancellor's Award for Undergraduate Research 2011

Selected Press Coverage

Television and Radio

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

Newspaper and Magazine

6. Associated Press How Do You Teach Human Interaction to a Robot? Lots of TV
7. NBC Deep Learning: Teaching Computers to Predict the Future
8. Newsweek Artificial Intelligence Algorithms Predicts the Future
9. Forbes MIT Computers Binge-Watch To Learn About Hugs
10. ABC News New AI Can Predict When Two People Will Kiss
11. Fox News New Artificial Intelligence Can Predict When You Will Kiss
12. Wired This AI learned to predict the future by watching loads of TV
13. Popular Science Algorithm Binge Watches TV to Predict Human Behavior
14. Scientific American Artificial Intelligence Can Predict How Scenes Will Play Out
15. New Scientist Binge-watching videos teaches computers to recognise sounds
16. New Scientist AI learns to predict the future by watching 2 million videos
17. Vice Magazine This Algorithm Taught Itself to Animate a Still Photo
18. The Verge Machine Learning's Next Trick is Generating Videos from Photos
19. The Week Junior A machine that learns by listening (children's magazine)
20. Technology Review Image Experiment Reveals The Building Blocks of Imagination

Invited Talks

Learning from Unlabeled Video

1. University of Massachusetts, Amherst Nov 2019

- 2. Butterfly Network Nov 2018
- 3. International Computer Vision Summer School Jul 2018
- 4. CVPR Tutorial Jun 2018

Predictive Vision

- 5. University of Maryland, College Park Mar 2018
- 6. University of Pennsylvania Nov 2017
- 7. Snapchat Research Nov 2017
- 8. University of Southern California Nov 2017
- 9. Workshop on Video Frontiers Nov 2017
- 10. Rework Summit May 2017
- 11. University of California, San Deigo Apr 2017
- 12. Cornell University Apr 2017
- 13. University of Texas, Austin Mar 2017
- 14. Columbia University Mar 2017
- 15. Google Research Mar 2017
- 16. Adobe Research Mar 2017
- 17. OpenAI Mar 2017
- 18. Brown University Feb 2017
- 19. University of California, Los Angeles Feb 2017
- 20. NVidia Feb 2017
- 21. Rework Summit Nov 2016
- 22. Twitter Oct 2016
- 23. TTI Chicago Sep 2016
- 24. Massachusetts Institute of Technology Sep 2016
- 25. Apple Aug 2016
- 26. University of California, Berkeley Aug 2016
- 27. Stanford University Aug 2016
- 28. Boston University Mar 2016
- 29. University of Massachusetts, Boston Mar 2016

Visualizing Object Detection Features

- 30. University of Massachusetts, Boston Mar 2016
- 31. Massachusetts Institute of Technology Sep 2015
- 32. Brown University Nov 2013

Efficient Video Annotation

- 33. CVPR Workshop Jun 2013
- 34. CVPR Workshop Jun 2011

Teaching

- 1. W4731 Computer Vision, Columbia Engineering Fall 2019
- 2. E6998 Advanced Computer Vision, Columbia Engineering Spring 2019

3. W4731 Computer Vision, Columbia Engineering

Fall 2018

Professional Service

1. Senior Program Committee (Area Chair), CVPR 2020
2. Senior Program Committee (Area Chair), ICLR 2020
3. Senior Program Committee (Area Chair), NeurIPS 2019
4. Senior Program Committee (Area Chair), CVPR 2019
5. Senior Program Committee (Area Chair), ICML 2019
6. Senior Program Committee (Area Chair), CVPR 2018
7. Organizer, Workshop on Self-supervised Learning, ICML 2019
8. Organizer, Workshop on Learning from Unlabeled Video, CVPR 2019
9. Organizer, Tutorial on Unsupervised Visual Learning, CVPR 2018
10. Program Committee, Action and Anticipation Workshop, CVPR 2016
11. Program Committee, Action and Anticipation Workshop, CVPR 2017
12. Program Committee, Action and Anticipation Workshop, CVPR 2018
13. Program Committee, Workshop on Human Computation for Image Analysis, CVPR 2016
14. Reviewer for CVPR, ICCV, ECCV, NIPS, ICML, IJCV, PAMI, 2011 to 2020

Departmental Service

1. Admissions Committee 2019
2. Distinguished Lectures Committee 2018-2019