

Andrew Owens

CONTACT INFORMATION	Website: http://andrewowens.com Email: andrew.owens@cornell.edu	Office: Bloomberg 368 Cornell Tech
EXPERIENCE	Cornell Tech <i>Associate Professor of Computer Science</i>	Jul. 2025 – present
	University of Michigan <i>Assistant Professor of Electrical Engineering and Computer Science</i>	Jan. 2020 – Jun. 2025
	UC Berkeley <i>Postdoctoral Researcher</i> Advisors: Alexei Efros, Jitendra Malik	Sep. 2016 – Dec. 2019
	Microsoft Research , Redmond, WA <i>Research Intern</i> Advisor: Rick Szeliski	Summer 2014
	Google , Seattle, WA <i>Software Engineering Research Intern</i> Advisor: Sameer Agarwal	Summer 2011
EDUCATION	Massachusetts Institute of Technology Ph.D., Electrical Engineering and Computer Science Advisors: William Freeman, Antonio Torralba Thesis: <i>Learning Visual Models from Paired Audio-Visual Examples</i>	2013 – 2016
	Massachusetts Institute of Technology M.S., Electrical Engineering and Computer Science Advisors: William Freeman, Antonio Torralba	2010 – 2013
	Cornell University B.A., Computer Science Advisors: Daniel Huttenlocher, Noah Snaveley	2006 – 2010
HONORS	Sloan Research Fellowship	2025
	NSF CAREER	2024
	UMich 1938E Award (highest award for junior faculty in College of Engineering)	2024
	Outstanding Reviewer Award, ICASSP 2023	2023
	UMich EECS Outstanding Achievement Award	2022
	Sony Research Award 2021	2021
	Best Paper Award, Honorable Mention. WACV 2022	2022
	Outstanding Reviewer Award, NeurIPS 2021	2021
	RA-L Best Paper Award Finalist	2018
	Best Reviewer Award, ICLR 2018	2018
	Microsoft Research Fellowship	2015
	NSF Graduate Research Fellowship (declined)	2012
	NDSEG Fellowship	2011

	Best Paper Award, Honorable Mention. CVPR 2011	2011
	CRA Outstanding Undergraduate Researcher Award – Finalist	2010
FUNDING	Sloan Research Fellowship (\$75,000)	2025 – ongoing
	NSF CAREER Award (\$599,778)	2024 – ongoing
	Title: <i>Career: Learning Multimodal Representations of the Physical World</i>	
	DARPA Grant (\$633,195)	2020 – 2024
	Title: <i>Semantic Information Defender</i>	
	Toyota Research Institute (\$125,254)	2022 – 2023
	Sony Research Award (\$100,000)	2022 – 2023
	Cisco Systems:	
	Title: <i>Learning Audio-Visual Grouping</i> (\$149,999)	2021 – 2022
	Title: <i>Learning Correspondence-based Measures of Image Similarity</i> (\$149,999)	2022 – 2023
	Title: <i>Gift funding</i> (\$100,000)	2023
	Adobe gift funding (\$40,000 total)	2022 – 2025
PROFESSIONAL ACTIVITIES	Co-organizer UMich AI Symposium (2024)	
	CVPR Workshop Chair (2024)	
	Lead organizer, <i>Sight and Sound</i> workshop at CVPR 2018-2025.	
	Co-organizer, <i>AV4D: Visual Learning of Sounds in Spaces</i> workshop, ECCV 2022, ICCV 2023.	
	Co-organizer, <i>Open World Vision</i> workshop, CVPR 2021-2025.	
	Co-organizer, <i>Embodied Multimodal Learning</i> workshop at ICLR 2021.	
	Reviewer: CVPR (2015-2020, 2022, 2026), ICCV (2015, 2017, 2019, 2021), ECCV (2016, 2018, 2020, 2022), SIGGRAPH (2020, 2024), SIGGRAPH Asia (2024), ICLR (2018, 2019, 2021, 2022), ICRA (2019, 2020, 2024), ICML (2017), NeurIPS (2017, 2019, 2021, 2022, 2025), CHI (2018), UIST (2019), ACL (2022), CoRL (2022), ICASSP (2023), RA-L (2025)	
	Area Chair: CVPR (2021, 2023, 2024, 2025), NeurIPS (2023), NeurIPS Dataset and Benchmarks (2022), WACV (2023), ICCV (2023, 2025), ECCV (2024), ICLR (2026)	
	NSF Panelist (2023, 2024, 2025)	
PHD STUDENTS SUPERVISED	Daniel Geng	UMich PhD student, 2020 - 2025 → OpenAI
	NSF Graduate Research Fellow	
	Ziyang Chen	UMich PhD student, 2022 - 2025 → Luma AI
	Ayush Shrivastava	UMich PhD student, 2021 - present
	Jeongsoo Park	UMich PhD student, 2023 - present
	Yiming Dou	UMich PhD student, 2023 - present
	Samanta Rodriguez	UMich PhD student, 2024 - present
	Co-advised with Nima Fazeli	
	NSF Graduate Research Fellow	
	Chao Feng	UMich PhD student, 2024 - present
	Xuanchen Lu	UMich PhD student, 2024 - present
OTHER ADVISING	Xixi Hu	UMich MS, 2020 - 2021 → UT Austin CS PhD
	Jing Zhu	UMich undergrad, 2020 - 2021 → UMich CS PhD
	Max Hamilton	UMich MS, 2021 - 2022 → UMass Amherst CS PhD
	Zhangxing Bian	UMich MS, 2020 - 2021 → Johns Hopkins PhD
	Yuexi Du	UMich undergrad, 2021 - 2022 → Yale CS PhD
	Rui Guo	UMich MS, 2021 → Xmotors.ai
	Oscar de Lima, 2020	UMich MS → Neato Robotics
	Fengyu Yang	UMich undergrad, 2021 - 2023 → Yale PhD / Robotics startup founder

CRA Outstanding Undergraduate Award Runner-up, 2023

Jiacheng Zhang	UMich undergrad, 2022 - 2023 → UMich PhD
Chenhao Zheng	UMich undergrad 2022 - 2024 → University of Washington PhD
<i>CRA Outstanding Undergraduate Award Honorable Mention, 2024</i>	
Chenyang Ma	UMich undergrad 2021 - 2022 → Cambridge MS
Zhaoying Pan	UMich MS, 2022 - 2023 → Purdue PhD
Yuqing Luo	UMich BS, 2024 - 2025 → CMU MS
Wonseok Oh	UMich MS, 2024 - 2025 → University of Colorado Boulder PhD
Zixuan Pan	UMich undergrad/MS, 2022 - 2025 → Simon Fraser University PhD
Zihao Wei	UMich undergrad/MS, 2022 - 2025 → University of Maryland PhD
Inbum Park	UMich MS, 2023 - 2025 → University of Maryland PhD

PHD THESIS COMMITTEES	Mandela Patrick (Oxford; chair: Andrea Vedaldi)	2021
	Ryan Szeto (UMich; chair: Jason Corso)	2021
	Wonhui Kim (UMich; chair: Matt Johnson-Roberson)	2021
	Junming Zhang (UMich; chair: Matt Johnson-Roberson)	2022
	Yizhen Zhang (UMich; chair: Zhongming Liu)	2021
	Moitreya Chaterjee (UIUC; chair: Narendra Ahuja)	2022
	Haozhu Wang (UMich; chair: Jay Guo)	2022
	Madan Ganesh (UMich; chair: Jason Corso)	2022
	Oana Ignat (UMich; chair: Rada Mihalcea)	2022
	Shurjo Banerjee (UMich; chair: Jason Corso)	2022
	Rodrigo Mira (Imperial College London; chair: Björn Schuller)	2023
	Yu Chen (UMich; chair: Hun-Seok Kim)	2023
	Santiago Castro (UMich; chair: Rada Mihalcea)	2024
	Nathan Louis (UMich; chair: Jason Corso)	2024
	Ekdeep Singh Lubana (UMich; chair: Robert Dick)	2024
	Nilesh Kulkarni (UMich; chairs: David Fouhey and Justin Johnson)	2024
	Karan Desai (UMich; chair: Justin Johnson)	2023
	Mohamed El Banani (UMich; chair: Justin Johnson)	2023
	Didac Surís (Columbia; chair: Carl Vondrick)	2024
	Asiegbu Miracle Kanu-Asiegbu (UMich; chairs: Xiaoxiao Du and Ram Vasudevan)	2025
Jing Zhu (UMich; chair: Danai Koutra)	2025	
Mingyu Yang (UMich; chair: Hun-Seok Kim)	ongoing	
Linyi Jin (UMich; chair: David Fouhey)	ongoing	
Christopher Rockwell (UMich; chairs: David Fouhey and Justin Johnson)	ongoing	
Yash Kant (University of Toronto; chair: Igor Gilitschenski)	ongoing	
Etai Sella (Tel Aviv University; chair: Hadar Elor)	ongoing	
Jacob Chalk (University of Bristol; chair: Dima Damen)	ongoing	

SELECTED PRESS COVERAGE OF MY WORK	<i>In Motion</i> , an art exhibit based on our motion sculpture work. MIT Museum, 2019
	MIT Develops a Novel Camouflaging Algorithm That Hides Eyesores. <i>Wired</i> , 2014.
	MIT researchers built an AI that predicts what the world sounds like. <i>Quartz</i> , 2016.
	This computer is selecting sound effects for silent videos that seem so real humans can't tell they're fake. <i>Washington Post</i> , 2016.
	Creating 3D sculptures from 2D video and other news. <i>BBC</i> , 2018.
	New algorithm can help spot faked photos before they go viral. <i>New Scientist</i> , 2018.
	Fooling You, Fooling Me. <i>Communications of the ACM</i> , 2025.

PRESS
COVERAGE
AS THIRD-PARTY
EXPERT

Teaching artificial intelligence to connect senses like vision and touch. *MIT News*, 2019.
Is technology spying on you? New AI could prevent eavesdropping. *Science*. 2022.
Paparazzi Photos Were the Scourge of Celebrities. Now, It's AI. *Wall Street Journal*, 2023.

OUTREACH

Mentor for doctoral consortium
At CVPR 2021, CVPR 2024, CVPR 2025
Participant in doctoral mentor session
At ICCV 2021, CVPR 2022, ECCV 2022, CVPR 2023
Instructor at Michigan Institute for Data & AI in Society High School Summer Camp, 2023

PUBLICATIONS

Students from my group (at the time of doing the work) are indicated with the following colors:
blue for PhD students, green for MS students, and red for undergraduate students.

Conference and journal publications:

- [1] **Ayush Shrivastava**, **Sanyam Mehta**, **Daniel Geng**, **Andrew Owens**. Point Prompting: Counterfactual Tracking with Video Diffusion Models. *International Conference on Learning Representations (ICLR)*, 2026.
- [2] **Ayush Shrivastava**, Connelly Barnes, Xuaner Zhang, Lingzhi Zhang, **Andrew Owens**, Sohrab Amirghodsi, Eli Shechtman. Fine-grained Defocus Blur Control for Generative Image Models. *Winter Conference on Applications of Computer Vision (WACV)*, 2026.
- [3] **Samanta Rodriguez**, **Yiming Dou**, Miquel Oller, **Andrew Owens**, Nima Fazeli. Cross-Sensor Touch Generation. *Conference on Robot Learning (CoRL)*, 2025.
- [4] Sikai Li, **Samanta Rodriguez**, **Yiming Dou**, **Andrew Owens**, Nima Fazeli. Tactile Functasets: Neural Implicit Representations of Tactile Datasets. *International Conference on Robotics and Automation (ICRA)*, 2025.
- [5] **Chao Feng**, **Ziyang Chen**, Aleksander Holynski, Alexei A. Efros, **Andrew Owens**. GPS as a Control Signal for Image Generation. *Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [6] **Ziyang Chen**, Prem Seetharaman, Bryan Russell, Oriol Nieto, David Bourgin, **Andrew Owens**, Justin Salamon. Video-Guided Foley Sound Generation with Multimodal Controls. *Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [7] **Jeongsoo Park**, **Andrew Owens**. Community Forensics: Using Thousands of Generators to Train Fake Image Detectors. *Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [8] **Daniel Geng**, Charles Herrmann, Junhwa Hur, Forrester Cole, Chen Sun, Oliver Wang, Tobias Pfaff, Tatiana Lopez-Guevara, Carl Doersch, Yusuf Aytar, Michael Rubinstein, **Andrew Owens**, Deqing Sun. Motion Prompting: Controlling Video Generation with Motion Trajectories. *Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [9] **Yiming Dou**, **Wonseok Oh**, **Yuqing Luo**, Antonio Loquercio, **Andrew Owens**. Hearing Hands: Generating Sounds from Physical Interactions in 3D Scenes. *Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [10] Anna Min, **Ziyang Chen**, Hang Zhao, **Andrew Owens**. Supervising Sound Localization using In-the-wild Egomotion. *Computer Vision and Pattern Recognition (CVPR)*, 2025.

- [11] [Ayush Shrivastava](#), [Andrew Owens](#). Self-Supervised Spatial Correspondence Across Modalities. *Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [12] [Samanta Rodriguez](#), [Yiming Dou](#), William van den Bogert, Miquel Oller, Kevin So, [Andrew Owens](#), Nima Fazeli. Contrastive Touch-to-Touch Pretraining. *International Conference on Robotics and Automation (ICRA)*, 2025.
- [13] [Ziyang Chen](#), [Daniel Geng](#), [Andrew Owens](#). Images that Sound: Composing Images and Sounds on a Single Canvas. *Neural Information Processing Systems (NeurIPS)*, 2024.
- [14] [Daniel Geng](#), [Inbum Park](#), [Andrew Owens](#). Factorized Diffusion: Perceptual Illusions by Noise Decomposition. *European Conference on Computer Vision (ECCV)*, 2024.
- [15] [Ayush Shrivastava](#), [Andrew Owens](#). Self-Supervised Any-Point Tracking by Contrastive Random Walks. *European Conference on Computer Vision (ECCV)*, 2024.
- [16] Tingle Li, Renhao Wang, Po-Yao Huang, [Andrew Owens](#), Gopala Krishna Anumanchipalli. Self-Supervised Audio-Visual Soundscape Stylization. *European Conference on Computer Vision (ECCV)*, 2024.
- [17] [Yiming Dou](#), [Fengyu Yang](#), [Yi Liu](#), Antonio Loquercio, [Andrew Owens](#). Tactile-Augmented Radiance Fields. *Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [18] [Daniel Geng](#), [Inbum Park](#), [Andrew Owens](#). Visual Anagrams: Generating Multi-View Optical Illusions with Diffusion Models. *Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [19] [Ziyang Chen](#), Israel D. Gebru, Christian Richardt, Anurag Kumar, William Laney, [Andrew Owens](#), Alexander Richard. Real Acoustic Fields: An Audio-Visual Room Acoustics Dataset and Benchmark. *Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [20] [Fengyu Yang](#), [Chao Feng](#), [Ziyang Chen](#), Hyoungseob Park, Daniel Wang, [Yiming Dou](#), Ziyao Zeng, Xien Chen, Rit Gangopadhyay, [Andrew Owens](#), Alex Wong. Binding Touch to Everything: Learning Unified Multimodal Tactile Representations. *Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [21] [Zihao Wei](#), [Zixuan Pan](#), [Andrew Owens](#). Efficient Vision-Language Pre-training by Cluster Masking. *Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [22] [Daniel Geng](#), [Andrew Owens](#). Motion Guidance: Diffusion-Based Image Editing with Differentiable Motion Estimators. *International Conference on Learning Representations (ICLR) 2024*, 2024.
- [23] [Zhaoying Pan](#), [Daniel Geng](#), [Andrew Owens](#). Self-Supervised Motion Magnification by Backpropagating Through Optical Flow. *Neural Information Processing Systems (NeurIPS)*, 2023.
- [24] [Ziyang Chen](#), Shengyi Qian, [Andrew Owens](#). Sound Localization from Motion: Jointly Learning Sound Direction and Camera Rotation. *International Conference on Computer Vision (ICCV)*, 2023.
- [25] [Fengyu Yang](#), [Jiacheng Zhang](#), [Andrew Owens](#). Generating Visual Scenes from Touch. *International Conference on Computer Vision (ICCV)*, 2023.
- [26] Lukas Höllein, Ang Cao, [Andrew Owens](#), Justin Johnson, Matthias Nießner. Text2Room: Extracting Textured 3D Meshes from 2D Text-to-Image Models. *International Conference on Computer Vision (ICCV)*, 2023.

- [27] Jiatian Sun, Longxiulin Deng, Triantafyllos Afouras, **Andrew Owens**, Abe Davis. Eventfulness for Interactive Video Alignment. *Proceedings of ACM SIGGRAPH*, 2023.
- [28] **Chenhao Zheng**, **Ayush Shrivastava**, **Andrew Owens**. EXIF as Language: Learning Cross-Modal Associations Between Images and Camera Metadata. *Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [29] **Rui Guo**, Jasmine Collins, **Oscar de Lima**, **Andrew Owens**. GANmouflage: 3D Object Nondetection with Texture Fields. *Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [30] **Chao Feng**, **Ziyang Chen**, **Andrew Owens**. Self-Supervised Video Forensics by Audio-Visual Anomaly Detection. *Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [31] **Yuexi Du**, **Ziyang Chen**, Justin Salamon, Bryan Russell, **Andrew Owens**. Conditional Generation of Audio from Video via Foley Analogies. *Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [32] Kim Sung-Bin, Arda Senocak, Hyunwoo Ha, **Andrew Owens**, Tae-Hyun Oh. Sound to Visual Scene Generation by Audio-to-Visual Latent Alignment. *Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [33] **Fengyu Yang**, **Chenyang Ma**, **Jiacheng Zhang**, **Jing Zhu**, Wenzhen Yuan, **Andrew Owens**. Touch and Go: Learning from Human-Collected Vision and Touch. *Neural Information Processing Systems (NeurIPS) - Datasets and Benchmarks Track*, 2022.
- [34] **Ziyang Chen**, David F. Fouhey, **Andrew Owens**. Sound Localization by Self-Supervised Time Delay Estimation. *European Conference on Computer Vision (ECCV)*, 2022.
- [35] Artem Abzaliev, **Andrew Owens**, Rada Mihalcea. Towards Understanding the Relation between Gestures and Language. *International Conference On Computational Linguistics (COLING)*, 2022.
- [36] Tingle Li, Yichen Liu, **Andrew Owens**, Hang Zhao. Learning Visual Styles from Audio-Visual Associations. *European Conference on Computer Vision (ECCV)*, 2022.
- [37] **Zhangxing Bian**, Allan Jabri, Alexei A. Efros, **Andrew Owens**. Learning Pixel Trajectories with Multiscale Contrastive Random Walks. *Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [38] **Daniel Geng**, **Max Hamilton**, **Andrew Owens**. Comparing Correspondences: Video Prediction with Correspondence-wise Losses. *Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [39] **Xixi Hu**, **Ziyang Chen**, **Andrew Owens**. Mix and Localize: Localizing Sound Sources in Mixtures. *Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [40] Medhini Narasimhan, Shiry Ginosar, **Andrew Owens**, Alexei A. Efros, Trevor Darrell. Strumming to the Beat: Audio-Conditioned Contrastive Video Textures. *Winter Conference on Applications of Computer Vision (WACV)*, 2022.
- [41] **Ziyang Chen**, **Xixi Hu**, **Andrew Owens**. Structure from Silence: Learning Scene Structure from Ambient Sound. *Conference on Robot Learning (CoRL)*, 2021.
- [42] Linyi Jin, Shengyi Qian, **Andrew Owens**, David F. Fouhey. Planar Surface Reconstruction from Sparse Views. *International Conference on Computer Vision (ICCV)*, 2021.
- [43] Allan Jabri, **Andrew Owens**, Alexei A. Efros. Space-Time Correspondence as a Contrastive Random Walk. *Neural Information Processing Systems (NeurIPS)*, 2020.

- [44] Triantafyllos Afouras, **Andrew Owens**, Joon Son Chung, Andrew Zisserman. Self-Supervised Learning Of Audio-Visual Objects From Video. *European Conference on Computer Vision (ECCV)*, 2020.
- [45] Sheng-Yu Wang, Oliver Wang, Richard Zhang, **Andrew Owens**, Alexei A. Efros. CNN-generated images are surprisingly easy to spot... for now. *Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [46] Tianfan Xue, **Andrew Owens**, Daniel Scharstein, Michael Goesele, Richard Szeliski. Multi-frame stereo matching with edges, planes, and superpixels. *Image and Vision Computing*, 2019.
- [47] Sheng-Yu Wang, Oliver Wang, **Andrew Owens**, Richard Zhang, Alexei A. Efros. Detecting Photoshopped Faces by Scripting Photoshop. *International Conference on Computer Vision (ICCV)*, 2019.
- [48] Shiry Ginosar, Amir Bar, Gefen Kohavi, Caroline Chan, **Andrew Owens**, Jitendra Malik. Learning Individual Styles of Conversational Gesture. *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [49] **Andrew Owens**, Alexei A. Efros. Audio-Visual Scene Analysis with Self-Supervised Multi-sensory Features. *European Conference on Computer Vision (ECCV)*, 2018.
- [50] Minyoung Huh, Andrew Liu, **Andrew Owens**, Alexei A. Efros. Fighting Fake News: Image Splice Detection via Learned Self-Consistency. *European Conference on Computer Vision (ECCV)*, 2018.
- [51] Roberto Calandra, **Andrew Owens**, Dinesh Jayaraman, Justin Lin, Wenzhen Yuan, Jitendra Malik, Edward H. Adelson, Sergey Levine. More Than a Feeling: Learning to Grasp and Regrasp using Vision and Touch. *Robotics and Automation Letters (RA-L)*, 2018.
- [52] Xiuming Zhang, Tali Dekel, Tianfan Xue, **Andrew Owens**, Qiurui He, Jiajun Wu, Stefanie Mueller, William T. Freeman. MoSculp: Interactive Visualization of Shape and Time. *User Interface Software and Technology (UIST)*, 2018.
- [53] **Andrew Owens**, Jiajun Wu, Josh McDermott, William T. Freeman, Antonio Torralba. Learning Sight From Sound: Ambient Sound Provides Supervision for Visual Learning. *International Journal of Computer Vision (IJCV)*, 2018.
- [54] Roberto Calandra, **Andrew Owens**, Manu Upadhyaya, Wenzhen Yuan, Justin Lin, Edward H. Adelson, Sergey Levine. The Feeling of Success: Does Touch Sensing Help Predict Grasp Outcomes?. *Conference on Robot Learning (CoRL)*, 2017.
- [55] Wenzhen Yuan, Chenzhuo Zhu, **Andrew Owens**, Mandayam Srinivasan, Edward H. Adelson. Shape-independent Hardness Estimation Using Deep Learning and a GelSight Tactile Sensor. *International Conference on Robotics and Automation (ICRA)*, 2017.
- [56] **Andrew Owens**, Jiajun Wu, Josh McDermott, William T. Freeman, Antonio Torralba. Ambient Sound Provides Supervision for Visual Learning. *European Conference on Computer Vision (ECCV)*, 2016.
- [57] **Andrew Owens**, Phillip Isola, Josh McDermott, Antonio Torralba, Edward H. Adelson, William T. Freeman. Visually Indicated Sounds. *Computer Vision and Pattern Recognition (CVPR)*, 2016.

- [58] **Andrew Owens**, Connelly Barnes, Alex Flint, Hanumant Singh, William T. Freeman. Camouflaging an Object from Many Viewpoints. *Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [59] David Crandall, **Andrew Owens**, Noah Snavely, Dan Huttenlocher. SfM with MRFs: Discrete-Continuous Optimization for Large-Scale Structure from Motion. *Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2013.
- [60] **Andrew Owens**, Jianxiong Xiao, Antonio Torralba, William T. Freeman. Shape Anchors for Data-Driven Multi-view Reconstruction. *International Conference on Computer Vision (ICCV)*, 2013.
- [61] Jianxiong Xiao, **Andrew Owens**, Antonio Torralba. SUN3D: A Database of Big Spaces Reconstructed using SfM and Object Labels. *International Conference on Computer Vision (ICCV)*, 2013.
- [62] David Crandall, **Andrew Owens**, Noah Snavely, Dan Huttenlocher. Discrete-Continuous Optimization for Large-Scale Structure from Motion. *Computer Vision and Pattern Recognition (CVPR)*, 2011.

Preprints:

- [1] Ege Ozguroglu, Junbang Liang, Ruoshi Liu, Mia Chiquier, Michael DeTienne, Wesley Wei Qian, Alexandra Horowitz, **Andrew Owens**, Carl Vondrick. New York Smells: A Large Multimodal Dataset for Olfaction. *arXiv*, 2025.

Theses:

- [1] Andrew Owens. Learning Visual Models from Paired Audio-Visual Examples. *Ph.D. Thesis, Massachusetts Institute of Technology*, 2016.
- [2] Andrew Owens. Combining Recognition and Geometry for Data-Driven 3D Reconstruction. *M.S. Thesis, Massachusetts Institute of Technology*, 2013.

TALKS

Point Prompting: Counterfactual Tracking with Video Diffusion Models
 ICCV Curated Data for Efficient Learning Workshop — October 2025

Moving beyond video in, audio out
 Generative AI for Audio-Visual Content Creation — October 2025

Generating Sounds for Physical Interactions in 3D Scenes
 ICCV Audio-Visual Generation and Learning Workshop — October 2025

Video generation for tracking & tracking for video generation
 Video AI Symposium 2025 — September 2025

Multimodal Learning from the Bottom Up
 Stanford Vision Seminar — July 2025

CVPR “What is Next in Multimodal Foundation Models?” workshop – June 2025
MIT Embodied Intelligence Seminar – March 2025
Google Vision and Graphics Seminar – March 2025
UIUC SINE Seminar – March 2025
Cornell University – December 2024
UPenn GRASP SFI Seminar – February 2024
Guest lecture: CMU Seminar on Multimodal Foundation Models – November 2023
Video AI Symposium 2023 at Google DeepMind – October 2023
Stanford University, Jiajun Wu’s group – March 2023
Adobe Research – March 2023
UC Berkeley, BAIR – March 2023

Image Forensics as Anomaly Detection

CVPR Visual Anomaly and Novelty Detection Workshop – June 2025

Controlling Diffusion Models with Motion and Geolocation

Adobe GenTech Seminar – April 17, 2025

Images that Sound

ECCV AVGenL Workshop – September 25, 2024

Connecting Sight, Sound, and Touch in 3D

CVPR Multimodalities for 3D Scenes (M3DS) Workshop – June 17, 2024

Learning Multimodal Models of the Physical World

CVPR Visual Perception and Learning in an Open World Workshop – June 18, 2024

Generating Multi-view Visual Illusions

UC San Diego, Pixel Cafe – October 31, 2025
Guest lecture, TTIC 31270 (Host: Shiry Ginosar) – October 2025
Cornell Tech, Learning Machines Seminar – October 24, 2024
University of Washington, GRAIL Seminar – October 18, 2024
Midwest Computer Vision Workshop – September 15, 2024

Tactile-Augmented Radiance Fields

CompVision meeting, UC Berkeley – February 28, 2024

Learning Multimodal Models of the Physical World

Oxford Visual Geometry Group (VGG) – September 26, 2023
Caltech Vision Group – August 2023
Notre Dame – August 2023

Learning by Audio-Visual Analogy

Keynote Address, DCASE Workshop – September 2023

Sound Localization from Motion, paper talk

ICCV AV4D workshop – October 2023

Image Forensics as Open World Perception
CVPR Visual Perception and Learning in an Open World Workshop — June 2023

Cross-modal synthesis from sight, sound, and touch
AAAI Creative AI Across Modalities Workshop — February 2023

Learning Visual, Audio, and Cross-Modal Correspondences
CMU VASC Seminar — November 2022

Learning Correspondences with Contrastive Random Walks
ECCV “What is Motion For?” Workshop — October 24, 2022

Sound Localization by Self-Supervised Time Delay Estimation, paper talk
ECCV AV4D workshop — October 23, 2022

Learning to Represent and Synthesize Motion
University of Rochester – Computer Vision Seminar — April 2021

Learning Image Forensics
Google Computational Imaging Workshop — March 2020

Learning Audio-Visual Objects
ECCV Multi-Modal Video Analysis Workshop — August 2020

Learning Sight from Sound
Oxford University — September 2019
Facebook AI Video Summit — June 2019
CVPR Multimodal Learning and Applications Workshop — June 2019
Google Machine Perception Workshop — October 2018
RSS Workshop on Multi-Modal Perception and Control — May 2018
Toyota Technological Institute Chicago — March 2018

Audio-Visual Scene Analysis with Self-Supervised Multisensory Features
Oral presentation, ECCV 2018 — September 2018

Self-Supervising Sight, Sound, and Image Forensics
CVPR Workshop, Beyond Supervised Learning — May 2018
University of Southern California — October 2018

Visually Indicated Sounds
Oral presentation, CVPR 2016 — June 2016

Ambient Sound Provides Supervision for Visual Learning
Oral presentation, ECCV 2016 — October 2016

Sound Provides Supervision for Visual Learning

CMU Robotics Institute – April 2016

Camouflaging an Object From Many Viewpoints
Oral presentation, CVPR 2014 – June 2014

Guest Lecture, CS194-26, UC Berkeley – October 2016 and 2017

TEACHING

Computer Vision (UMich EECS 442)	Fall 2024
Advanced Topics in Computer Vision (UMich EECS 542)	Winter 2024
Computer Vision (UMich EECS 442)	Fall 2023
Computer Vision (UMich EECS 442, shared lecture with EECS 504)	Fall 2022
Foundations of Computer Vision (UMich EECS 504, shared lecture with EECS 504)	Fall 2022
Advanced Topics in Computer Vision (UMich EECS 542)	Winter 2022
Computer Vision (UMich EECS 442)	Fall 2021
Unsupervised Visual Learning (EECS 598-012)	Winter 2021
Computer Vision (UMich EECS 442)	Fall 2020
Foundations of Computer Vision (EECS 504)	Winter 2020