Andrew Owens

CONTACT INFORMATION

Website: http://andrewowens.com
Email: ahowens@umich.edu

EECS 4231
University of Michigan

EDUCATION

Massachusetts Institute of Technology 2013 – 2016
Ph.D., Electrical Engineering and Computer Science
Advisors: William Freeman and Antonio Torralba
Thesis: *Learning Visual Models from Paired Audio-Visual Examples*

Massachusetts Institute of Technology 2010 – 2013
M.S., Electrical Engineering and Computer Science
Advisors: William Freeman and Antonio Torralba

Cornell University 2006 – 2010
B.A., Computer Science
Advisor: Daniel Huttenlocher

EXPERIENCE

University of Michigan 2020 – present
Assistant Professor
Electrical Engineering and Computer Science

UC Berkeley 2016 – 2019
Postdoctoral Researcher
Advisors: Alexei Efros and Jitendra Malik

Microsoft Research, Redmond, WA Summer 2014
Research Intern
Advisor: Rick Szeliski

Google, Seattle, WA Summer 2011
Software Engineering Research Intern
Advisor: Sameer Agarwal

HONORS

Outstanding Reviewer Award, ICASSP 2023

UMich EECS Outstanding Achievement Award, 2022

Sony Research Award 2021

Best Paper Award, Honorable Mention. WACV 2022

Outstanding Reviewer Award, NeurIPS 2021

RA-L Best Paper Award Finalist, 2018

Best Reviewer Award, ICLR 2018

Microsoft Research Fellowship, 2015 - 2016

NSF Graduate Research Fellowship, 2012 (declined)

NDSEG Fellowship, 2011 - 2014

Best Paper Award, Honorable Mention. CVPR 2011

CRA Outstanding Undergraduate Researcher Award – Finalist, 2010

FUNDING

DARPA Grant, subcontractor for Kitware, Inc. ($633,195), 2020 - 2024
Title: Semantic Information Defender

Toyota Research Institute ($125,254), 2022-2023
Title: Meta-Learning Compositional Representations for 3D Video Understanding

Sony ($100,000) 2022-2023
Title: Learning auditory scene analysis for complex environments through audio-visual cycle consistency

Cisco Systems
Learning Audio-Visual Grouping, ($149,999) 2021-2022
Learning Correspondence-based Measures of Image Similarity ($149,999), 2022-2023
Gift funding ($100,000 pending), 2023

Adobe gift ($10,000), 2022

Publications

Publications:

[1] Ziyang Chen, Shengyi Qian, Andrew Owens. Sound Localization from Motion: Jointly Learning Sound Direction and Camera Rotation. *International Conference on Computer Vision (ICCV)*, 2023.

[2] Fengyu Yang, Jiacheng Zhang, Andrew Owens. Generating Visual Scenes from Touch. *International Conference on Computer Vision (ICCV)*, 2023.

[3] 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.

[4] Jiatian Sun, Longxiulin Deng, Triantafyllos Afouras, Andrew Owens, Abe Davis. Eventfulness for Interactive Video Alignment. *Proceedings of ACM SIGGRAPH*, 2023.

[5] 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.

[6] Rui Guo, Jasmine Collins, Oscar de Lima, Andrew Owens. GANmouflage: 3D Object Nondetection with Texture Fields. *Computer Vision and Pattern Recognition (CVPR)*, 2023.

[7] Chao Feng, Ziyang Chen, Andrew Owens. Self-Supervised Video Forensics by Audio-Visual Anomaly Detection. *Computer Vision and Pattern Recognition (CVPR)*, 2023.

[8] 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.

[9] 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.

[10] 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.

[11] Ziyang Chen, David F. Fouhey, Andrew Owens. Sound Localization by Self-Supervised Time Delay Estimation. *European Conference on Computer Vision (ECCV)*, 2022.

[12] Artem Abzaliev, Andrew Owens, Rada Mihalcea. Towards Understanding the Relation between Gestures and Language. *International Conference On Computational Linguistics (COLING)*, 2022.

[13] Tingle Li, Yichen Liu, Andrew Owens, Hang Zhao. Learning Visual Styles from Audio-Visual Associations. *European Conference on Computer Vision (ECCV)*, 2022.

[14] Zhangxing Bian, Allan Jabri, Alexei A. Efros, Andrew Owens. Learning Pixel Trajectories with Multiscale Contrastive Random Walks. *Computer Vision and Pattern Recognition (CVPR)*, 2022.

[15] Daniel Geng, Max Hamilton, Andrew Owens. Comparing Correspondences: Video Prediction with Correspondence-wise Losses. *Computer Vision and Pattern Recognition (CVPR)*, 2022.

[16] Xixi Hu, Ziyang Chen, Andrew Owens. Mix and Localize: Localizing Sound Sources in Mixtures. *Computer Vision and Pattern Recognition (CVPR)*, 2022.

[17] 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.

[18] Ziyang Chen, Xixi Hu, Andrew Owens. Structure from Silence: Learning Scene Structure from Ambient Sound. *Conference on Robot Learning (CoRL)*, 2021.

[19] Linyi Jin, Shengyi Qian, Andrew Owens, David F. Fouhey. Planar Surface Reconstruction from Sparse Views. *International Conference on Computer Vision (ICCV)*, 2021.

[20] Allan Jabri, Andrew Owens, Alexei A. Efros. Space-Time Correspondence as a Contrastive Random Walk. *Neural Information Processing Systems (NeurIPS)*, 2020.

[21] 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.

[22] 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.

[23] Tianfan Xue, Andrew Owens, Daniel Scharstein, Michael Goesele, Richard Szeliski. Multi-frame stereo matching with edges, planes, and superpixels. *Image and Vision Computing*, 2019.

[24] 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.

[25] 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.

[26] Andrew Owens, Alexei A. Efros. Audio-Visual Scene Analysis with Self-Supervised Multisensory Features. *European Conference on Computer Vision (ECCV)*, 2018.

[27] 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.

[28] 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.

[29] 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.

[30] 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.

[31] 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.

[32] 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.

[33] 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.

[34] Andrew Owens, Phillip Isola, Josh McDermott, Antonio Torralba, Edward H. Adelson, William T. Freeman. Visually Indicated Sounds. *Computer Vision and Pattern Recognition (CVPR)*, 2016.

[35] Andrew Owens, Connelly Barnes, Alex Flint, Hanumant Singh, William T. Freeman. Camouflaging an Object from Many Viewpoints. *Computer Vision and Pattern Recognition (CVPR)*, 2014.

[36] 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.

[37] Andrew Owens, Jianxiong Xiao, Antonio Torralba, William T. Freeman. Shape Anchors for Data-Driven Multi-view Reconstruction. *International Conference on Computer Vision (ICCV)*, 2013.

[38] 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.

[39] David Crandall, Andrew Owens, Noah Snavely, Dan Huttenlocher. Discrete-Continuous Optimization for Large-Scale Structure from Motion. *Computer Vision and Pattern Recognition (CVPR)*, 2011.

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

Learning by Audio-Visual Analogy
Keynote Address, DCASE Workshop — September 21, 2023

Multimodal Learning from the Bottom Up
AI Video Symposium at Google DeepMind — October 1, 2023
Stanford University, Jiajun Wu's group — March 2023
Adobe Research — March 2023
UC Berkeley, BAIR — March 2023

Learning Multimodal Models of the Physical World
Oxford Visual Geometry Group (VGG) — September 26, 2023
Caltech Vision Group — August 2023
Notre Dame — August 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 2017

Guest Lecture, CS194-26, UC Berkeley — October 2016

Professional Activities

Lead organizer, *Sight and Sound* workshop at CVPR 2018-2023.

Co-organizer, *AV4D: Visual Learning of Sounds in Spaces* workshop, ECCV 2022, ICCV 2023.

Co-organizer, *Open World Vision* workshop, CVPR 2021-2023.

Co-organizer, *Embodied Multimodal Learning* workshop at ICLR 2021.

Reviewer: CVPR (2015-2022), ICCV (2015, 2017, 2019, 2021), ECCV (2016, 2018, 2020, 2022), SIGGRAPH (2020), ICLR (2018, 2019, 2021, 2022), ICRA (2019, 2020), ICML (2017), NeurIPS

(2017, 2019, 2021, 2022), CHI (2018), UIST (2019), ACL (2022), CoRL (2022), ICASSP (2023)

Area Chair: CVPR (2021, 2023), NeurIPS (2023), NeurIPS Dataset and Benchmarks (2022), WACV (2023), ICCV (2023)

NSF Panelist (2023)

CVPR Workshop Chair (2024)

PhD Students Supervised

Daniel Geng. UMich PhD student, 2020 - present

Ayush Shrivastava. UMich PhD student, 2021 - present

Ziyang Chen. UMich PhD student, 2022 - present

Other Advising

Xixi Hu. UMich MS, 2020 - 2021 → UT Austin CS PhD

Jing Zhu. UMich undergrad, 2020 - 2021 → UMich CS PhD

Chenhao Zheng. UMich undergrad, 2022 - ongoing → 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

Chenyang Ma, 2021 - 2022. UMich undergrad → Cambridge MS

Fengyu Yang, UMich undergrad 2021 - ongoing

Jiacheng Zhang, UMich undergrad, 2022 - ongoing

Sheng-Yu Wang. UC Berkeley Undergrad, 2018 - 2019 → CMU PhD

Minyoung Huh. UC Berkeley Undergrad, 2017 → MIT 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

Yizhen Zhang (UMich; chair: Zhongming Liu), 2021

Moitreya Chaterjee (UIUC; chair: Narendra Ahuja), 2022

Junming Zhang (UMich; chair: Johnson-Roberson), 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), ongoing

Yu Chen (UMich; chair: Hun-Seok Kim), 2023

Santiago Castro (UMich; chair: Rada Mihalcea), ongoing

Mingyu Yang (UMich; chair: Hun-Seok Kim), ongoing

Nathan Louis (UMich; chair: Jason Corso), ongoing

Ekdeep Singh Lubana (UMich; chair: Robert Dick), ongoing

Nilesh Kulkarni (UMich; chairs: David Fouhey and Justin Johnson), ongoing

Asiegbu Miracle Kanu-Asiegbu (UMich; chairs: Xiaoxiao Du and Ram Vasudevan), ongoing

Mohamed El Banani (UMich; chair: Justin Johnson), ongoing

Selected Press Coverage of My Work

In Motion, an art exhibit based on our motion sculpture work. MIT Museum, 2019

MIT Develops a Novel Camouflaging Algorithm That Hides Eyesores. *Wired*, 2014.

MIT researchers built an AI that predicts what the world sounds like. *Quartz*, 2016.

This computer is selecting sound effects for silent videos that seem so real humans can't tell they're fake. *Washington Post*, 2016.

Creating 3D sculptures from 2D video and other news. *BBC*, 2018.

New algorithm can help spot faked photos before they go viral. *New Scientist*, 2018.

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.