Angjoo Kanazawa

University of California, Berkeley
Berkeley Way West 8014, Berkeley, CA 94720
kanazawa@berkeley.edu
http://www.people.eecs.berkeley.edu/~kanazawa/

Academic and Research Employment

Assistant Professor, EECS, University of California, Berkeley — July 2020–

Research Scientist, Google Research — September 2019-August 2021

Postdoctoral Scholar, EECS, University of California, Berkeley — 2017-2019

Education

Ph.D., Computer Science, University of Maryland, College Park, USA — July 2017
Thesis: *Single-View 3D Reconstruction of Animals*
Advisor: David Jacobs

BA., Computer Science and Mathematics, New York University, USA — May 2011
Magna Cum Laude

Top Five Cited Papers (as of July 31 2023)

1. **A. Kanazawa**, M. J. Black, D. W. Jacobs, and J. Malik, "End-to-end recovery of human shape and pose," in *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2018 Citation: 1524
2. F. Bogo*, **A. Kanazawa***, C. Lassner, P. Gehler, J. Romero, and M. J. Black, "Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image," in *European Conference on Computer Vision (ECCV)*, 2016 Citation: 1423
3. S. Saito*, Z. Huang*, R. Natsume*, S. Morishima, **A. Kanazawa**, and H. Li, "PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization," in *International Conference on Computer Vision (ICCV)*, October 2019 Citation: 927
4. A. Yu, V. Ye, M. Tancik, and **A. Kanazawa**, "pixelNeRF: Neural radiance fields from one or few images," in *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2021 Citation: 663
5. **A. Kanazawa***, S. Tulsiani*, A. A. Efros, and J. Malik, "Learning category-specific mesh reconstruction from image collections," in *European Conference on Computer Vision (ECCV)*, 2018 Citation: 528

Publications in Reverse Chronological Order

1. S. Goel, G. Pavlakos, J. Rajasegaran, **A. Kanazawa***, and J. Malik*, "Humans in 4D: Reconstructing and tracking humans with transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023
2. J. Kerr*, C. M. Kim*, K. Goldberg, **A. Kanazawa**, and M. Tancik, "Lerf: Language embedded radiance fields," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023
3. A. Haque, M. Tancik, A. Efros, A. Holynski, and **A. Kanazawa**, "Instruct-nerf2nerf: Editing 3d scenes with instructions," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023

4. R. Li, H. Gao, M. Tancik, and **A. Kanazawa**, "Nerfacc: Efficient sampling accelerates nerfs.," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023
5. F. Warburg*, E. Weber*, M. Tancik, A. Hoyski, and **A. Kanazawa**, "Nerfbusters: Removing ghostly artifacts from casually captured nerfs," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023
6. M. Tancik*, E. Weber*, E. Ng*, R. Li, B. Yi, J. Kerr, T. Wang, A. Kristoffersen, J. Austin, K. Salahi, A. Ahuja, D. McAllister, and **A. Kanazawa**, "Nerfstudio: A modular framework for neural radiance field development," in *ACM SIGGRAPH 2023 Conference Proceedings*, (New York, NY, USA), Association for Computing Machinery, 2023
7. V. Ye, G. Pavlakos, J. Malik, and **A. Kanazawa**, "Decoupling human and camera motion from videos in the wild," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023
8. Sara Fridovich-Keil and Giacomo Meanti, F. R. Warburg, B. Recht, and **A. Kanazawa**, "K-planes: Explicit radiance fields in space, time, and appearance," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023
9. J. Rajasegaran, G. Pavlakos, **A. Kanazawa**, C. Feichtenhofer, and J. Malik, "On the benefits of 3d pose and tracking for human action recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023
10. J. Kerr, L. Fu, H. Huang, Y. Avigal, M. Tancik, J. Ichnowski, **A. Kanazawa**, and K. Goldberg, "Evo-nerf: Evolving nerf for sequential robot grasping of transparent objects," in *6th Annual Conference on Robot Learning*, 2022
11. H. Gao, R. Li, S. Tulsiani, B. Russell, and **A. Kanazawa**, "Monocular dynamic view synthesis: A reality check," in *Conference on Neural Information Processing Systems (NeurIPs)*, 2022
12. G. Pavlakos*, E. Weber*, M. Tancik, and **A. Kanazawa**, "The one where they reconstructed 3d humans and environments in tv shows," in *European Conference on Computer Vision (ECCV)*, 2022
13. Z. Li, Q. Wang, N. Snavely, and **A. Kanazawa**, "InfiniteNature-Zero: Learning perpetual view generation of natural scenes from single images," in *European Conference on Computer Vision (ECCV)*, 2022
14. R. Li, J. Tanke, M. Vo, M. Zollhoefer, J. Gall, **A. Kanazawa**, and C. Lassner, "Tava: Template-free animatable volumetric actors," in *European Conference on Computer Vision (ECCV)*, 2022
15. V. H. Maluleke, N. Thakkar, T. Brooks, E. Weber, T. Darrell, A. A. Efros, **A. Kanazawa**, and D. Guillory, "Studying bias in gans through the lens of race," in *European Conference on Computer Vision (ECCV)*, 2022
16. Z. Weng, K.-C. Wang, **A. Kanazawa**, and S. Yeung, "Domain adaptive 3d pose augmentation for in-the-wild human mesh recovery," in *2022 International Conference on 3D Vision (3DV)*, pp. 261–270, IEEE, 2022
17. S. Fridovich-Keil*, A. Yu*, M. Tancik, Q. Chen, B. Recht, and **A. Kanazawa**, "Plenoxels: Radiance fields without neural networks," in *Computer Vision and Pattern Recognition (CVPR)*, 2022
18. V. Ye, Z. Li, R. Tucker, **A. Kanazawa**, and N. Snavely, "Deformable sprites for unsupervised video decomposition," in *Computer Vision and Pattern Recognition (CVPR)*, June 2022
19. J. Rajasegaran, G. Pavlakos, **A. Kanazawa**, and J. Malik, "Tracking people by predicting 3d appearance, location and pose," in *Computer Vision and Pattern Recognition (CVPR)*, pp. 2740–2749, 2022
Best Paper Finalist
20. E. Ng, H. Joo, L. Hu, H. Li, , T. Darrell, **A. Kanazawa**, and S. Ginosar, "Learning to listen: Modeling non-deterministic dyadic facial motion," in *Computer Vision and Pattern Recognition (CVPR)*, 2022
21. G. Pavlakos, J. Malik, and **A. Kanazawa**, "Human mesh recovery from multiple shots," in *Computer Vision and Pattern Recognition (CVPR)*, 2022

22. S. Zhu, S. Ebrahimi, **A. Kanazawa**, and T. Darrell, "Differentiable gradient sampling for learning implicit 3d scene reconstructions from a single image," in *International Conference on Learning Representations (ICLR)*, 2021
23. J. Rajasegaran, G. Pavlakos, **A. Kanazawa**, and J. Malik, "Tracking people with 3d representations," *Advances in Neural Information Processing Systems*, vol. 34, pp. 23703–23713, 2021
24. R. Li, S. Yang, D. A. Ross, and **A. Kanazawa**, "Ai choreographer: Music conditioned 3d dance generation with aist++," in *International Conference on Computer Vision (ICCV)*, 2021
25. A. Yu, R. Li, M. Tancik, H. Li, R. Ng, and **A. Kanazawa**, "PlenOctrees for real-time rendering of neural radiance fields," in *International Conference on Computer Vision (ICCV)*, 2021
26. A. Liu*, R. Tucker*, V. Jampani, A. Makadia, N. Snavely, and **A. Kanazawa**, "Infinite nature: Perpetual view generation of natural scenes from a single image," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021
27. Z. Cao*, I. Radosavovic*, **A. Kanazawa**, and J. Malik, "Reconstructing hand-object interactions in the wild," in *International Conference on Computer Vision (ICCV)*, 2021
28. X. B. Peng, Z. Ma, P. Abbeel, S. Levine, and **A. Kanazawa**, "AMP: Adversarial motion priors for stylized physics-based character control," *ACM Transactions on Graphics, (ACM SIGGRAPH 2021 issue)*, vol. 40, July 2021
29. S. Wu, A. Makadia, J. Wu, N. Snavely, R. Tucker, and **A. Kanazawa**, "De-rendering the world's revolutionary artefacts," in *Computer Vision and Pattern Recognition (CVPR)*, 2021
30. T. Jakab, R. Tucker, A. Makadia, J. Wu, N. Snavely, and **A. Kanazawa**, "Keypointdeformer: Unsupervised 3d keypoint discovery for shape control," in *Computer Vision and Pattern Recognition (CVPR)*, 2020
31. A. Yu, V. Ye, M. Tancik, and **A. Kanazawa**, "pixelNeRF: Neural radiance fields from one or few images," in *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2021
32. J. Levinson, C. Esteves, K. Chen, N. Snavely, **A. Kanazawa**, A. Rostamizadeh, and A. Makadia, "An analysis of SVD for deep rotation estimation," in *Advances in Neural Information Processing Systems*, 2020
33. J. Y. Zhang*, S. Pepose*, H. Joo, D. Ramanan, J. Malik, and **A. Kanazawa**, "Perceiving 3d human-object spatial arrangements from a single image in the wild," in *European Conference on Computer Vision (ECCV)*, 2020
34. S. Goel, **A. Kanazawa**, , and J. Malik, "Shape and viewpoints without keypoints," in *European Conference on Computer Vision (ECCV)*, 2020
35. S. Zuffi, **A. Kanazawa**, T. Berger-Wolf, and M. J. Black, "Three-d safari: Learning to estimate zebra pose, shape, and texture from images "In the Wild"," in *International Conference on Computer Vision (ICCV)*, Oct. 2019
36. S. Saito*, Z. Huang*, R. Natsume*, S. Morishima, **A. Kanazawa**, and H. Li, "PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization," in *International Conference on Computer Vision (ICCV)*, October 2019
37. J. Y. Zhang, P. Felsen, **A. Kanazawa**, and J. Malik, "Predicting 3D human dynamics from video," in *International Conference on Computer Vision (ICCV)*, 2019
38. **A. Kanazawa***, J. Y. Zhang*, P. Felsen*, and J. Malik, "Learning 3D human dynamics from video," in *Computer Vision and Pattern Recognition (CVPR)*, 2019
39. X. B. Peng, **A. Kanazawa**, S. Toyer, P. Abbeel, and S. Levine, "Variational discriminator bottleneck: Improving imitation learning, inverse RL, and GANs by constraining information flow," in *International Conference on Learning Representations (ICLR)*, 2018
40. X. B. Peng, **A. Kanazawa**, J. Malik, P. Abbeel, and S. Levine, "SFV: Reinforcement learning of physical skills from videos," *ACM Transactions on Graphics, (ACM SIGGRAPH ASIA 2018 issue)*, vol. 37, Nov. 2018

41. **A. Kanazawa***, S. Tulsiani*, A. A. Efros, and J. Malik, "Learning category-specific mesh reconstruction from image collections," in *European Conference on Computer Vision (ECCV)*, 2018
42. **A. Kanazawa**, M. J. Black, D. W. Jacobs, and J. Malik, "End-to-end recovery of human shape and pose," in *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2018
43. S. Sengupta, **A. Kanazawa**, C. D. Castillo, and D. W. Jacobs, "Sfsnet: Learning shape, refectance and illuminance of faces "in the wild"," in *Computer Vision and Pattern Regognition (CVPR)*, 2018
44. S. Zuffi, **A. Kanazawa**, and M. J. Black, "Lions and tigers and bears: Capturing non-rigid, 3D, articulated shape from images," in *Computer Vision and Pattern Recognition (CVPR)*, pp. 3955–3963, IEEE Computer Society, 2018
45. Y. Huang, F. Bogo, C. Lassner, **A. Kanazawa**, P. V. Gehler, J. Romero, I. Akhter, and M. J. Black, "Towards accurate marker-less human shape and pose estimation over time," in *International Conference on 3D Vision (3DV)*, pp. 421–430, IEEE, 2017
46. S. Zuffi, **A. Kanazawa**, D. Jacobs, and M. J. Black, "3D menagerie: Modeling the 3D shape and pose of animals," in *Computer Vision and Pattern Recognition (CVPR)*, 2017
47. F. Bogo*, **A. Kanazawa***, C. Lassner, P. Gehler, J. Romero, and M. J. Black, "Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image," in *European Conference on Computer Vision (ECCV)*, 2016
48. **A. Kanazawa**, D. W. Jacobs, and M. Chandraker, "Warpnet: Weakly supervised matching for single-view reconstruction," in *Computer Vision and Pattern Recognition (CVPR)*, pp. 3253–3261, 2016
49. **A. Kanazawa**, S. Kovalsky, R. Basri, and D. Jacobs, "Learning 3d deformation of animals from 2d images," in *Computer Graphics Forum (EUROGRAPHICS 2016 issue)*, 2016 **Best Paper Award**
50. J. Liu, **A. Kanazawa**, D. Jacobs, and P. Belhumeur, "Dog breed classification using part localization," in *European Conference on Computer Vision (ECCV)*, pp. 172–185, Springer, 2012

Awards and Honors

Sloan Fellow 2023
— $75,000 over two years, awarded to 22 early career researchers across Computer Science
Spark Award, Bakar Fellows Program 2022
— $225,000 over three years
Hellman Fellow, Hellman Foundation 2022
— $60,000 Google Research Scholar Program Award 2021
— Awarded to 7 early career faculty globally in machine vision.
Rising Stars in EECS 2018
— Awarded to 80 women EECS graduate and postdoctoral women.
Best Paper Award, Rank Prize 2018
— Symposium on Geometry and Uncertainty in Deep Learning for Computer Vision.
Best Paper Award, EUROGRAPHICS 2016
— "Learning 3D Articulation and Deformation using 2D Images"
Google Anita Borg Scholar 2011
— Awarded to 25 students in the US
NYU Computer Science Prize for Academic Excellence and Service to the Department 2011

Teaching

CS194-26/294-26: Intro to Computer Vision & Computational Photography
Co-Instructor UC Berkeley, Fall 2021, Fall 2022
CS184/284a: Foundations of Computer Graphics UC Berkeley, Spring 2021
Co-Instructor

CS294-173 Learning for 3D Vision UC Berkeley, Fall 2020, Fall 2023
Instructor

CMSC828L Deep Learning University of Maryland, College Park, Fall 2016
TA for Prof. David Jacobs
CMSC421 Introduction to Artificial Intelligence University of Maryland, College Park, Spring 2012
TA for Prof. Hal Daumé III
CMSC 131 Object-Oriented Programming I University of Maryland, College Park, Fall 2011
Teaching Assistant
CSCI-UA.0101 Introduction to Computer Science I New York University, Fall 2008
Teaching Assistant
CSCI-UA.0103 Introduction to Computer Science II New York University, Spring 2009
Teaching Assistant

Synergistic Activities

Professional Services:

Area Chair: CVPR 2020-2023, ICCV 2021, ECCV 2022, WACV 2021, ACCV 2020,

Diversity Chair: ICCV 2023, ACCV 2020

Program Committee: ACM SIGGRAPH (2021, 2022), ACM SIGGRAPH Asia 2023

Reviewer: CVPR, ICCV, ECCV, NeurIPS, PAMI, 3DV, ICRA, SIGGRAPH, SIGGRAPH Asia

Invited Talks and Interviews:

1. **Keynote** Korean Conference on Computer Vision August 8, 2023
   From Videos to 4D Worlds and Beyond
2. Computer Vision for Mixed Reality, CVPR'23 Workshop June 18, 2023
   Capturing and Interacting with the 3D World
3. End-to-End Autonomous Driving, CVPR'23 Workshop June 18, 2023
   Perceiving 4D People in the World: Progress on Human Mesh Recovery — from 2018 to 2023
4. Generative Models for Computer Vision, CVPR'23 Workshop June 18, 2023
   Editing 3D Scenes & Modeling 3D Social Interaction with Generative Models
5. Stanford Invited Talk May 14, 2023
   From Videos to 4D Worlds & Beyond
6. Neural Fields across Fields, ICLR'23 workshop May 4, 2023
   Editing 3D Scenes with Instructions & beyond
7. CMU Invited Talk April 11, 2023
   From Videos to 4D Worlds & Beyond
8. UMD Invited Talk April 10, 2023
   From Videos to 4D Worlds & Beyond
9. NVIDIA GTC March 2023
   Nerfstudio: A Modular Framework for Neural Radiance Field Development
10. SXSW Panel: Understanding the Role of AI in Reshaping the Film TV Industry March 2023
    website
11. Behavioral Digital Twins for Smart Cities, WACV'23 Workshop January 6, 2023
    Perceiving People in the 3D World: Next Steps

12. **Keynote** Korea AI Summit: Visual AI 2022 — December 13, 2022
    Towards Capturing Reality: Scenes and 3D People: recording
13. What is Motion For? ECCV'22 workshop — October 2022
    Motion is Everything
14. Neural Geometry and Rendering, ECCV'22 workshop — October 2022
    Dynamic View Synthesis: A Reality Check: recording
15. **Keynote** International Conference on 3D Vision — September 15 2022
    Towards 4D Reality Capture
16. Computational Cameras and Displays (CCD), CVPR '22 workshop — June 2022
    The Role of Neural Networks in Radiance Fields
17. WiGRAPH Path Tracing Interview: link — May 2022
18. Motion Planning with Implicit Neural Representations of Geometry, ICRA'22 Workshop — May 2022
    Towards Practical Reality Capture
19. Adobe Tech Summit, Digital Humans, Keynote — May 2022
    Perceiving People and Places in the 3D World
20. Netflix Workshop on Machine Learning for Content Creation — May 2022
21. Stanford 2022 HAI Spring Conference, Speaker and Panelist — April 12, 2022
    Towards Capturing Reality
22. Columbia University Vision Seminar — April 7, 2022
    Towards Capturing Reality
23. Stanford CS231A Invited Lecture — March 9, 2022
    Practical Modeling of the Plenoptic Function
24. NVIDIA Invited Talk — February 4, 2022
    Practical Modeling of the Plenoptic Function
25. AI for Creative Video Editing and Understanding, ICCV'21 Workshop — October 2021
    Infinite Nature: Perpetual View Generation of Natural Scenes from a Single Image, video link
26. Human Trajectory and Pose Dynamics Forecasting in the Wild, ICCV'21 Workshop — October, 2021
    Learning to Dance! Music Conditioned 3D Human Motion Generation
27. More Exploration, Less Exploitation, ICCV'21 Workshop — October, 2021
28. Deep Learning for Geometric Computing, ICCV'21 Workshop — October, 2021
29. Unsupervised 3D Learning In the Wild, ICCV'21 Workshop — October, 2021
    Real-time rendering of NeRFs with PlenOctrees, video link
30. Brown Visual Computing Seminar — September 27, 2021
    Infinite Nature: Perpetual View Generation of Natural Scenes from a Single Image
31. Frontiers of Monocular Perception, CVPR'21 workshop — June 2021
    Predicting Scenes from One or Few Images
32. Media Forensics, CVPR'21 workshop — June 2021
    Towards Relighting and Material Recovery from Image Collections
33. 3D Scene Understanding for Vision, Graphics, and Robotics, CVPR'21 workshop — June 2021
    Perceiving 3D Human Interaction in the Wild.
34. Learning from Unlabeled Video, CVPR'21 workshop — June 2021
    Infinite Nature Perpetual View Generation of Natural Scenes from a Single Image
35. SMPL-made-simple Tutorial, CVPR'21 workshop — June 2021
    Visual Imitation with SMPL
36. Stanford Imaging SCIEN Seminar — May 2021
    Pushing the Boundaries of Novel View Synthesis

37. ML Collective — April 2021
Infinite Nature: Perpetual View Generation of Natural Scenes from a Single Image

38. UCSD Invited Lecture — March 2021

39. MIT Vision Seminar — February 2021
On Novel and Perpetual View Synthesis

40. UIUC Vision Group — February 2021
On Novel and Perpetual View Synthesis

41. Google Research MobileVision team — January 2021

42. **Keynote**, ACM SIGGRAPH European Conference on Visual Media Production (CVMP) — Dec 2020
Perceiving Humans, Animals and Objects in the 3D World

43. MITxHarvard Women in AI Interview — October 2020
https://www.youtube.com/watch?v=MGo_Vca29m0

44. TUM AI Lecture Series — July 2020
Perceiving Humans and Objects in the 3D World

45. Learning 3D Representations for Shape and Appearance, ECCV'20 workshop — August 2020
Learning Morphable Shape Models from Image Collections

46. Tracking and its many guises, ECCV'20 workshop — August 2020
Challenges in perceiving 3D humans in videos

47. Women in Machine Learning (WiML) Panel ICML'20 — July 2020

48. AI for Content Creation CVPR'20 Workshop — June 2020

49. Compositionality in Computer Vision CVPR'20 Workshop — June 2020

Workshops/Tutorials Organized:

Scholars & Big Models: How Can Academics Adapt? — CVPR 2023
website

Tutorial: Neural Volumetric Rendering for Computer Vision — ECCV 2022

CV4Animals: Computer Vision for Animal Behavior Tracking and Modeling — CVPR 2021-2023
website

Artificial Social Intelligence Workshop — CVPR 2022, 2023

AI for Content Creation — CVPR 2021

3D Scene Understanding for Vision, Graphics, and Robotics — CVPR 2021

3D Poses in the Wild Challenge — ECCV 2020

Sensing, Understanding and Synthesizing Humans — ICCV 2019

Women in Computer Vision — ECCV 2018

Service Activities:

Other Engagements:

1. UCB Society of Women Engineers (SWE) Professor's Luncheon — March 2022
2. EECS Women's History Month, Faculty Lunch Host — March 2021
3. UCB Society of Women Engineers (SWE) townhall, Facilitator — March 2021

Faculty Sponsor, BAIR Research Experience for Undergraduates (REU) — Summer 2021-

3DGV: Seminar on 3D Geometry and Vision, Organizer — September 2020-2021

Graduate Admissions Committee, UC Berkeley — 2018, 2020-2022

Mentor, EECS Peers, UC Berkeley	2018
Computer Vision Student Seminar Organizer, University of Maryland, College Park	2012-2015
President of Women in Computing, New York University	2009-2011
Vice President of ACM, New York University	2010-2011