

EXPERIENCE

Google DeepMind , Staff Research Scientist	2024-
Columbia University , Assistant Professor	2024-
University of California, Berkeley , Postdoctoral Scholar (w/ Alexei Efros, Angjoo Kanazawa)	2022 - 2025
Google DeepMind , Senior Research Scientist	2023 - 2024
Google Research , Research Scientist	2022 - 2023
Tooploox AI , Research Advisor	2018 - 2022
Facebook , Research Intern w/ Richard Szeliski. (See Publications [5])	2018
Facebook Visiting Researcher, Computational Photography Group	2017 - 2018
Facebook , Research Intern w/ Johannes Kopf. (See Publications [3])	2017
Google , Research Intern w/ Carlos Hernandez, Changchang Wu	2016
Qualcomm Research & Development , Research Intern Image Processing Algorithms	2014
Qualcomm Innovation Center , Intern, Digital Signal Processing	2013
Qualcomm Inc. , Intern, Camera Sensor Drivers & Algorithms	2012

EDUCATION

University of Washington Ph.D. in Computer Science and Engineering Advised by Steve Seitz, Richard Szeliski, Brian Curless	2014 - 2022
University of Illinois at Urbana-Champaign B.S. in Computer Science with High Honors Advisors: Robin Kravets, Svetlana Lazebnik	2011 - 2014

PUBLICATIONS

- [38] G. Luo, J. Granskog, **A. Holynski**, T. Darrell, "Dual-Process Image Generation" in IEEE International Conference on Computer Vision (ICCV), 2025
- [37] S. Szymanowicz, J. Zhang, P. Srinivasan, R. Gao, A. Brussee, **A. Holynski**, R. Brualla, J. Barron, P. Henzler, "Bolt3D: Generating 3D Scenes in Seconds" in IEEE International Conference on Computer Vision (ICCV), 2025
- [36] C. Feng, Z. Chen, **A. Holynski**, A.A. Efros, A. Owens, "GPS as a Control Signal for Image Generation" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2025
- [35] Q. Wang*, Y. Zhang*, **A. Holynski**, A. Efros, A. Kanazawa, "Continuous 3D Perception Model with Persistent State" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2025, (*Oral*)

- [34] L. Jin, R. Tucker, Z. Li, D. Fouhey, N. Snavely*, **A. Holynski***, "Stereo4D: Learning How Things Move in 3D from Internet Stereo Videos" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2025, (*Oral*)
- [33] Z. Li, R. Tucker, F. Cole, Q. Wang, L. Jin, V. Ye, A. Kanazawa, **A. Holynski**, N. Snavely, "MegaSaM: Accurate, Fast and Robust Structure and Motion from Casual Dynamic Videos" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2025, (*Oral, Best Paper, Honorable Mention*)
- [32] A. Trevithick, R. Paiss, P. Henzler, D. Verbin, R. Wu, H. Alzayer, R. Gao, B. Poole, J. Barron, **A. Holynski**, R. Ramamoorthi, P. Srinivasan, "SimVS: Simulating World Inconsistencies for Robust View Synthesis" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2024
- [31] R. Wu, R. Gao, B. Poole, A. Trevithick, C. Zheng, J. Barron, **A. Holynski**, "CAT4D: Create Anything in 4D with Multi-View Video Diffusion Models" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2024, (*Oral*)
- [30] X. Wang, B. Zhou, B. Curless, I. Kemelmacher, **A. Holynski**, S. Seitz, "Generative Inbetweening: Adapting Image-to-Video Models for Keyframe Interpolation" in International Conference on Learning Representations (ICLR), 2025
- [29] J. Ma, E. Lu, R. Paiss, S. Zada, **A. Holynski**, T. Dekel, B. Curless, M. Rubinstein, F. Cole, "VidPanos: Generative Panoramic Videos from Casual Panning Videos" in SIGGRAPH Asia, 2024
- [28] I. Siglidis, **A. Holynski**, A.A. Efros, M. Aubry, S. Ginosar, "Diffusion Models as Data Mining Tools" in European Conference on Computer Vision (ECCV), 2024
- [27] D. McAllister, S. Ge, J.B. Huang, D.W. Jacobs, A.A. Efros, **A. Holynski**, A. Kanazawa, "Rethinking Score Distillation as a Bridge Between Image Distributions" in Advances in Neural Information Processing Systems (NeurIPS), 2024
- [26] Y. Wang, **A. Holynski**, B. Curless, S. Seitz, "Infinite Texture: Text-guided High Resolution Diffusion Texture Synthesis", in arXiv, 2024
- [25] R. Gao*, **A. Holynski***, P. Henzler, A. Brussee, R. Martin-Brualla, P. Srinivasan, J. Barron, B. Poole*, "CAT3D: Create Anything in 3D With Multi-View Diffusion Models", in Advances in Neural Information Processing Systems (NeurIPS), 2024, (*Oral*)
- [24] Z. Li, R. Tucker, N. Snavely, **A. Holynski**, "Generative Image Dynamics" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2024, (*Oral, Best Paper Award*)
- [23] S. Jain, D. Watson, **A. Holynski**, E. Tabellion, B. Poole, J. Kontkanen, "Video Interpolation with Diffusion Models" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2024
- [22] D. Epstein, B. Poole, B. Mildenhall, A. Efros, **A. Holynski**, "Disentangled 3D Scene Generation with Layout Learning", in International Conference on Machine Learning (ICML), 2024
- [21] M.L. Shih, W.C. Ma, L. Boyce, **A. Holynski**, F. Cole, B. Curless, J. Kontkanen, "ExtraNeRF: Visibility-Aware View Extrapolation of Neural Radiance Fields with Diffusion Models" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2024
- [20] R. Wu*, B. Mildenhall*, P. Henzler, R. Gao, K. Park, D. Watson, P.P. Srinivasan, D. Verbin, J.T. Barron, B. Poole, **A. Holynski***, "ReconFusion: 3D Reconstruction with Diffusion Priors" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2024
- [19] G. Luo, T. Darrell, O. Wang, DB Goldman, **A. Holynski**, "Readout Guidance: Learning Control from Diffusion Features" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2024, (*Highlight*)
- [18] E. Weber, **A. Holynski**, V. Jampani, S. Saxena, N. Snavely, A. Kar, A. Kanazawa, "NeRFiller: Completing Scenes via Generative 3D Inpainting" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2024
- [17] X. Wang, J. Kontkanen, B. Curless, S. Seitz, I. Kemelmacher, B. Mildenhall, P.P. Srinivasan, D. Verbin, **A. Holynski**, "Generative Powers of Ten" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2024, (*Highlight*)

- [16] R. Po, W. Yifan, V. Golyanik, K. Aberman, J. T. Barron, A. Bermano, E. Chan, T. Dekel, **A. Holynski**, A. Kanazawa, C. K. Liu, L. Liu, B. Mildenhall, M. Niessner, B. Ommer, C. Theobalt, P. Wonka, G. Wetzstein, "State of the Art on Diffusion Models for Visual Computing" in Eurographics, 2024
- [15] L. Tang, N. Ruiz, Q. Chu, Y. Li, **A. Holynski**, D.E. Jacobs, B. Hariharan, Y. Pritch, N. Wadhwa, K. Aberman, M. Rubinstein, "RealFill: Reference-Driven Generation for Authentic Image Completion" in SIGGRAPH, 2024
- [14] Q. Wang, Y. Chang, R. Cai, Z. Li, B. Hariharan, **A. Holynski**, N. Snavely, "Tracking Everything Everywhere All at Once" in IEEE International Conference on Computer Vision (ICCV), 2023, *(Oral, Best Student Paper Award)*,
- [13] D. Epstein, A. Jabri, B. Poole, A.A. Efros, **A. Holynski**, "Diffusion Self-Guidance for Controllable Image Generation" in Advances in Neural Information Processing Systems (NeurIPS), 2023
- [12] G. Luo, L. Dunlap, S.D.H. Park, **A. Holynski**, T. Darrell, "Diffusion Hyperfeatures: Searching Through Time and Space for Semantic Correspondence" in Advances in Neural Information Processing Systems (NeurIPS), 2023
- [11] A. Haque, M. Tancik, A.A. Efros, **A. Holynski**, A. Kanazawa, "Instruct-NeRF2NeRF: Editing 3D Scenes with Instructions" in IEEE International Conference on Computer Vision (ICCV), 2023, *(Oral)*
- [10] F. Warburg*, E. Weber*, M. Tancik, **A. Holynski**, A. Kanazawa, "Nerfbusters: Removing Ghostly Artifacts from Casually Captured NeRFs" in IEEE International Conference on Computer Vision (ICCV), 2023
- [9] J. Karras, **A. Holynski**, T. Wang, I. Kemelmacher, "DreamPose: Fashion Image-to-Video Synthesis via Stable Diffusion" in IEEE International Conference on Computer Vision (ICCV), 2023
- [8] T. Brooks*, **A. Holynski***, A.A. Efros, "InstructPix2Pix: Learning to Follow Image Editing Instructions" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2023, *(Highlight)*
- [7] Y. Wang, **A. Holynski**, X. Zhang, X. Zhang, "SunStage: Portrait Reconstruction and Relighting using the Sun as a Light Stage" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2023
- [6] **A. Holynski**, B. Curless, S.M. Seitz, R. Szeliski, "Animating Pictures with Eulerian Motion Fields" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2021, *(Oral)*
- [5] **A. Holynski**, D. Geraghty, J.M. Frahm, C. Sweeney, R. Szeliski, "Reducing Drift in Structure from Motion using Extended Features" in International Conference on 3D Vision (3DV), 2020, *(Oral)*
- [4] J.J. Park, **A. Holynski**, S.M. Seitz, "Seeing the World in a Bag of Chips" in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020, *(Oral)*
- [3] **A. Holynski**, J. Kopf, "Fast Depth Densification for Occlusion-aware Augmented Reality" in ACM Transactions on Graphics (SIGGRAPH Asia), 2018
- [2] C. Sweeney, **A. Holynski**, B. Curless, S.M. Seitz, "Structure from Motion for Panorama-Style Videos" in arXiv, 2018
- [1] A. Khosrowpour, I. Fedorov, **A. Holynski**, J.C. Niebles, M. Golparvar-Fard, "Automated Worker Activity Analysis in Indoor Environments for Direct Work Rate Improvement from long sequences of RGBD Images" in Construction Research Congress (CRC), 2014

PRODUCT IMPACT

Veo 3: Google DeepMind's state-of-the-art video generation model, now with audio!	2025
Veo 2: Google DeepMind's state-of-the-art video generation model	2024
Google Search 360-Spin: AI-generated 360-degree spins around products on Google Search	2024

AWARDS & HONORS

CVPR Best Paper, Honorable Mention for “MegaSAM: [...]”	2025
CVPR Oral (top 0.4% of submissions) for “Stereo4D: [...]”	2025
CVPR Oral (top 0.4% of submissions) for “MegaSAM: [...]”	2025
CVPR Oral (top 0.4% of submissions) for “Continuous 3D [...]”	2025
CVPR Oral (top 0.4% of submissions) for “CAT4D: [...]”	2025
NeurIPS Oral (top 0.4% of submissions) for “CAT3D: [...]”	2024
CVPR Best Paper Award for “Generative Image Dynamics”	2024
CVPR Oral (top 0.8% of submissions) for “Generative Image Dynamics”	2024
CVPR Highlight (top 2.8% of submissions) for “Readout Guidance: ...”	2024
CVPR Highlight (top 2.8% of submissions) for “Generative Powers of Ten”	2024
ICCV Best Student Paper Award for “Tracking Everything Everywhere All at Once”	2023
ICCV Oral (top 1.8% of submissions) for “Tracking Everything Everywhere All at Once”	2023
ICCV Oral (top 1.8% of submissions) for “Instruct-NeRF2NeRF: ...”	2023
CVPR Highlight (top 2.6% of submissions) for “InstructPix2Pix: ...”	2023
CVPR Oral (top 4.1% of submissions) for “Animating Pictures ...”	2021
3DV Oral (top 13% of submissions) for “Reducing Drift ...”	2020
CVPR Oral (top 5.7% of submissions) for “Seeing the World in a Bag of Chips”	2020
University of Washington Reality Lab Fellowship	2018 - 2022
Runner-up, Pacific Northwest ACM ICPC	2014
Leach/Winokur Endowed Fellowship in Computer Science & Engineering	2014 - 2015
Achievement Rewards for College Scientists (ARCS) Fellowship	2014 - 2016
University of Illinois Edmund J. James Scholar	2011 - 2014
University of Illinois College of Engineering Dean’s List	2011 - 2014

TEACHING

CS 180 Computer Vision & Computational Photography, UC Berkeley (Guest Lecture)	2024
CS184/284A Computer Graphics and Imaging, UC Berkeley (Guest Lecture)	2024
CS 280 Graduate Computer Vision, UC Berkeley (Guest Lecture)	2024
CS 194-173 Learning for 3D Vision, UC Berkeley (Guest Lecture)	2023
CSE576 Computer Vision, University of Washington (Guest Lecture)	2020
CSE576 Computer Vision, University of Washington (Graduate Student Instructor)	2020
CSE481V AR/VR Capstone, University of Washington (Graduate Student Instructor)	2019
CS398 Computer Architecture, University of Illinois (Teaching Assistant)	2013 - 2014
CS125 Intro. to Computer Science, University of Illinois (Teaching Assistant)	2012

ACADEMIC SERVICE

Volunteer:

ICCV DEI Committee 2023

Reviewer:

SIGGRAPH, SIGGRAPH Asia, TOG, TVCG, CVPR, ECCV, ICCV, 3DV 2014-

Area Chair:

ICLR 2025

Organizer:

Workshop on “4D Vision”, CVPR 2025

Workshop on “The Future of Generative Visual Art”, CVPR 2024

STUDENT COLLABORATORS & INTERNS

Haian Jin @ Google DeepMind. 2025

Kyle Sargent @ Google DeepMind. 2025

Ale Escontrela @ Google DeepMind. 2024

Stan Szymanowicz @ Google DeepMind. 2024

David Charatan @ Google DeepMind. 2024

Linyi Jin, Research Intern @ Google DeepMind. See Publications [34]. 2024

Alex Trevithick @ Google DeepMind. See Publications [32]. 2024

David McAllister @ UC Berkeley. See Publications [27]. 2024

Rundi Wu, Research Intern @ Google Research. See Publications [20,31]. 2023, 2024

Ethan Weber @ Google Research. See Publications [18]. 2023

Luming Tang @ Google Research. See Publications [15]. 2023

Dave Epstein, Research Intern @ Google Research. See Publications [13,23]. 2023

Grace Luo, Research Intern @ Google Research. See Publications [19]. 2023

Xiaojuan Wang, Research Intern @ Google Research. See Publications [17]. 2023

Qianqian Wang, Research Intern @ Google Research. See Publications [14]. 2022

Meng-Li Shih, Research Intern @ Google Research. See Publications [21]. 2022

Yifan Wang @ University of Washington. See Publications [7]. 2022

Johanna Karras @ University of Washington. See Publications [9]. 2022

Jingwei Ma @ University of Washington, Google Research. 2022

GRADUATE COMMITTEES SERVED

Rundi Wu, PhD, Columbia University May 2025

Orr Avrech, Masters, Columbia University Apr 2025

Yifan Wang, PhD, University of Washington Jun 2024

DEPARTMENTAL SERVICE

Graduate Admissions Committee, UC Berkeley 2023-2025

INVITED TALKS

“How I Learned to Stop Worrying and Love the (3D) Data Monster” Stanford University	Oct 2024
“CAT3D: Create Anything in 3D with Multi-View Diffusion Models” CVPR 2024 Workshop on AI for 3D Generation	Jun 2024
“How I Learned to Stop Worrying and Love the Data Monster” University of British Columbia	Apr 2024
UT Austin	Mar 2024
Columbia University	Mar 2024
TTI Chicago	Mar 2024
ELLIS Institute Tübingen	Mar 2024
University of California, Berkeley	Feb 2024
“Generative Models As Creative Tools: What’s Next?” NeurIPS 2023 Workshop on Machine Learning for Creativity and Design	Dec 2023
“Generative Models for Visual Creation” Warsaw Institute of Technology, Poland	Oct 2023
“Diffusion Priors for Creation and Editing” Bay Area Computer Vision Day, Stanford	Sep 2023
“2D Generative Priors for 3D” Summit for Understanding the Real World, Mountain View, CA	Aug 2023
“Language-grounded Visual Manipulation” CVPR 2023 6th Multimodal Learning and Applications Workshop	Jun 2023
“Tracking Everything Everywhere All At Once” Google Computer Vision Talk Series	May 2023
“InstructPix2Pix: Learning to Follow Image Editing Instructions” VISTEC Graduate Institute, Thailand	Mar 2023
Google Vision-Language Seminar	Feb 2023
“Augmenting Visual Memories” NVIDIA	Nov 2021
Massachusetts Institute of Technology	Oct 2021
Adobe Research	Oct 2021
Apple	Oct 2021
Facebook	Oct 2021
Google Research	Sep 2021
UC Berkeley	Aug 2021

ETC

Languages:

English (fluent, technical), Spanish (fluent, technical), Farsi (fluent), French (conversational), Polish (conversational)

Programming:

C/C++, Python, OpenGL, PyTorch, Jax

Misc Projects:

Holoscanner: Gamified 3D Scanning ([link](#))