

Akshay Paruchuri

✉ akshaypa@stanford.edu | 🌐 akshayparuchuri.com/ | 📧 yahskapar | 📄 in akshayparuchuri

Summary

My research interests are at the intersection of machine learning, multimodal agents, and interpretability, with a focus on health AI applications and a commitment to AI safety. Currently, I'm working toward a future where multimodal agents can safely and reliably integrate into healthcare systems in order to improve the entire patient journey, from advanced diagnostic imaging and surgical support to all-day health monitoring and management, with the aim to achieve better therapeutic outcomes for cancer and aging-related diseases. I've published in leading venues on topics such as remote health sensing (WACV, NeurIPS), 3D reconstruction (ECCV, MICCAI), LLM-based conversational agents for personal health (EMNLP, Nature Communications), and energy-efficient operation of smart glasses (ISMAR). I'm generally interested in opportunities that would allow me to continue to deepen my research expertise while leading and working on projects with meaningful, positive real-world impact, especially with respect to areas such as healthcare, environmental sustainability, and public policy.

Education

University of North Carolina at Chapel Hill

PH.D. IN COMPUTER SCIENCE

- Advisor: Henry Fuchs
- Thesis Title: *Toward Perceptual and Ubiquitous Intelligence for Healthcare*

Chapel Hill, NC, USA

Aug 2021 - Dec 2025

North Carolina State University

B.S. IN ELECTRICAL AND COMPUTER ENGINEERING

Raleigh, NC, USA

Aug 2014 - Dec 2019

Experience

Stanford University

POSTDOCTORAL SCHOLAR (ADVISOR: EHSAN ADELI)

- Research at the intersection of machine learning, multimodal agents, and interpretability.

Palo Alto, CA, USA

Jan 2026 - Present

Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA)

VISITING RESEARCHER (ADVISOR: PIOTR DIDYK)

- Research involving computer vision and human perception, with a focus on multimodal language models (MLMs) and interpretability.

Lugano, Ticino, CH

Jun 2025 - Jan 2026

Google

STUDENT RESEARCHER (ADVISOR: ISHAN CHATTERJEE)

- Research involving context-aware, multimodal systems, with a focus on realizing AR smart glasses for all-day use.

Seattle, WA, USA

Feb 2025 - Jan 2026

University of North Carolina at Chapel Hill

GRADUATE RESEARCH ASSISTANT (ADVISOR: HENRY FUCHS)

- Research at the intersection of health AI, computer vision, and machine learning.

Chapel Hill, NC, USA

Aug 2021 - Dec 2025

Google

STUDENT RESEARCHER (ADVISORS: XIN LIU AND DANIEL McDUFF)

- Research at the intersection of healthcare and language models.

Seattle, WA, USA

Mar 2024 - Aug 2024

Kitware

RESEARCH AND DEVELOPMENT INTERN (ADVISOR: BRIAN CLIPP)

- Various research projects involving person re-identification, object detection, segmentation, and tracking.

Carrboro, NC, USA

Apr 2023 - Jul 2023

Nike

EMBEDDED SYSTEMS ENGINEER (MANAGER: VIKRAM MALHOTRA)

- Developed hardware, algorithms, and software toward novel, wearable consumer devices for experiences involving physical fitness.

Beaverton, OR, USA

Jan 2020 - Jul 2021

- Prototyped a feature-rich, non-form factor PCB to characterize power consumption in unique contexts and developed software toward meaningful gesture recognition using adaptive, self-lacing shoes.

Publications

- 13** Tianwen Zhou, **Akshay Paruchuri**, Josef Spjut, Kaan Akşit. Editing Physiological Signals in Videos Using Latent Representations. *arXiv preprint:2509.25348*, 2025. In submission.
- 12** A. Ali Heydari*, Ken Gu*, Vidya Srinivas*, Hong Yu*, Zhihan Zhang, Yuwei Zhang, **Akshay Paruchuri**, Qian He, Hamid Palangi, Nova Hammerquist, Ahmed A. Metwally, Brent Winslow, Yubin Kim, Kumar Ayush, Yuzhe Yang, Girish Narayanswamy, Maxwell A. Xu, Jake Garrison, Amy Aremnto Lee, Jenny Vafeiadou, Ben Graef, Isaac R. Galatzer-Levy, Erik Schenck, Andrew Barakat, Javier Perez, Jacqueline Shreibati, John Hernandez, Anthony Z. Faranesh, Javier L. Prieto, Connor Heneghan, Yun Liu, Jiening Zhan, Mark Malhotra, Shwetak Patel, Tim Althoff, Xin Liu[†], Daniel McDuff[†], and Xuhai "Orson" Xu[†]. The Anatomy of a Personal Health Agent. *arXiv preprint arXiv:2508.20148*, 2025. In submission.
- 11** **Akshay Paruchuri**, Maryam Aziz, Rohit Vartak, Ayman Ali, Best Uchehara, Xin Liu, Ishan Chatterjee, Monica Agrawal. "What's Up, Doc?": Analyzing How Users Seek Health Information in Large-Scale Conversational AI Datasets. *arXiv preprint arXiv:2506.21532*, 2025. Accepted to EMNLP 2025 (Findings).
[Spotlight Talk at Machine Learning for Health (ML4H) 2025 Conference]
- 10** Ken Gu, Zhihan Zhang, Kate Lin, Yuwei Zhang, **Akshay Paruchuri**, Hong Yu, Mehran Kazemi, Kumar Ayush, A. Ali Heydari, Maxwell A. Xu, Yun Liu, Ming-Zher Poh, Yuzhe Yang, Mark Malhotra, Shwetak Patel, Hamid Palangi, Xuhai Xu, Daniel McDuff, Tim Althoff, Xin Liu. RADAR: Benchmarking Language Models on Imperfect Tabular Data. *arXiv preprint arXiv:2506.08249*, 2025. Accepted to NeurIPS 2025 Datasets and Benchmarks Track.
- 9** **Akshay Paruchuri**, Sinan Hersek, Lavisha Aggarwal, Qiao Yang, Xin Liu, Achin Kulshrestha, Andrea Colaço, Henry Fuchs, and Ishan Chatterjee. EgoTrigger: Toward Audio-Driven Image Capture for Human Memory Enhancement in All-Day Energy-Efficient Smart Glasses. *arXiv preprint arXiv:2508.01915*, 2025. Accepted to ISMAR 2025. **[TVCG Journal Paper, Top 8%, 60 of 762 submissions]**
- 8** **Akshay Paruchuri**, Jake Garrison, Shun Liao, John Hernandez, Jacob Sunshine, Tim Althoff, Xin Liu, and Daniel McDuff. What Are the Odds? Language Models Are Capable of Probabilistic Reasoning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing* (pp. 11712-11733), November 2024.
- 7** **Akshay Paruchuri***, Mike A. Merrill*, Naghmeh Rezaei, Geza Kovacs, Javier Perez, Yun Liu, Erik Schenck, Nova Hammerquist, Jake Sunshine, Shyam Tailor, Kumar Ayush, Hao-Wei Su, Qian He, Cory McLean, Mark Malhotra, Shwetak Patel, Jiening Zhan, Tim Althoff[†], Daniel McDuff[†], and Xin Liu[†]. Transforming wearable data into personal health insights using large language model agents. *Nat Commun* 17, 1143 (2026).
<https://doi.org/10.1038/s41467-025-67922-y>
- 6** Shuxian Wang, **Akshay Paruchuri**, Zhaoxi Zhang, Sarah McGill, Roni Sengupta. Structure-Preserving Image Translation for Depth Estimation in Colonoscopy. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 667-677). Cham: Springer Nature Switzerland, October 2024. **[Oral]**
- 5** **Akshay Paruchuri**, Samuel Ehrenstein, Shuxian Wang, Inbar Fried, Stephen M. Pizer, Marc Niethammer, and Roni Sengupta. (2024). Leveraging Near-Field Lighting for Monocular Depth Estimation from Endoscopy Videos. In *European Conference on Computer Vision* (pp. 473-491). Springer, Cham.
- 4** Xin Liu, **Akshay Paruchuri***, Girish Narayanswamy*, Xiaoyu Zhang, Jiankai Tang, Yuzhe Zhang, Roni Sengupta, Shwetak Patel, Yuntao Wang, and Daniel McDuff. rPPG-Toolbox: Deep Remote PPG Toolbox. *Advances in Neural Information Processing Systems*, vol. 36, 2024. **[700 stars on GitHub as of July, 2025]**
- 3** **Akshay Paruchuri**, Xin Liu, Yulu Pan, Shwetak Patel, Daniel McDuff, and Soumyadip Sengupta. Motion Matters: Neural Motion Transfer for Better Camera Physiological Measurement. *Proceedings of the IEEE/CVF Winter*

Conference on Applications of Computer Vision (WACV), January 2024, pp. 5933-5942.

[Oral, Top 2.6%, 53 of 2042 submissions]

- 2 Qian Zhang, **Akshay Paruchuri**, Young-Woon Cha, Jia-Bin Huang, Jade Kandel, Howard Jiang, Adrian Ilie, Andrei State, Danielle Szafir, Daniel Szafir, and Henry Fuchs. Reconstruction of Human Body Pose and Appearance Using Body-Worn IMUs and a Nearby Camera View for Collaborative Egocentric Telepresence. *2023 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, Shanghai, China, 2023, pp. 96-97, doi: 10.1109/VRW58643.2023.00025.
- 1 Angelos Angelopoulos, Austin Hale, Husam Shaik, **Akshay Paruchuri**, Ken Liu, Randal Tuggle, and Daniel Szafir. Drone Brush: Mixed Reality Drone Path Planning. *Late-Breaking Reports at the IEEE/ACM International Conference on Human-Robot Interaction (HRI 2022)*.

Teaching

STANFORD

AI4ALL Pre-Collegiate Program (Summer 2026, Graduate Mentor)

UNC CHAPEL HILL

2D Computer Graphics (Fall 2024, Teaching Assistant)

Mentoring

Jessica Brown (Stanford CS PhD, Winter 2026-Present)
Favour Nerrise (Stanford EE PhD, Winter 2026-Present)
Fangrui Huang (Stanford CS PhD, Winter 2026-Present)
Gary Liu (Stanford SYMSYS BS, Winter 2026-Present)
Xiaoyue Fan (UCL CS MS, Winter 2026-Present)
Tianwen Zhou (UCL CS MS, Summer 2025-Fall 2025, now at Huawei as an AI Researcher)
Ashley Neall (UNC CS BS, Fall 2024-Present)
Ray Shealy (UNC CS BS, Fall 2024-Fall 2025, now at USAA as a SWE)
Yulu Pan (UNC CS BS, Fall 2022-Spring 2023, now at UNC Chapel Hill as a CS PhD)
Mingxuan Li (UNC CS BS, Spring 2022, now at CMU as a CS BS)

Presentations

From Sensing to Understanding: Building All-Day Wearable Systems for Personal Health Management

Keynote Presentation, Everyday Wearables for Personalized Health and Well-Being Workshop at CHI, Barcelona, ES (Spring 2026)

"What's Up, Doc?": Analyzing How Users Seek Health Information in Large-Scale Conversational AI Datasets

Spotlight Talk, ML4H, San Diego, CA (Fall 2025)

EgoTrigger: Toward Audio-Driven Image Capture for Human Memory Enhancement in All-Day Energy-Efficient Smart Glasses

Journal Paper Presentation, IEEE ISMAR (Fall 2025)

Toward All-Day Ambient and Wearable Intelligence for Healthcare

Research Talk, Stanford Translation AI (STAI) Lab, Remote (Summer 2025)

Research Talk, NYU Immersive Computing Lab, Remote (Fall 2025)

Research Talk, UCL Computational Light Laboratory, Remote (Fall 2025)

Your life's co-pilot. In your glasses.

Research Talk, NVIDIA Research, Durham, NC (Summer 2025)

Motion Matters: Neural Motion Transfer for Better Camera Physiological Sensing

Poster Presentation, International Conference on Computational Photography (Summer 2023)

Poster Presentation, UNC Data Science Day (Fall 2023)

Oral + Poster Presentation, IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (Winter 2024)

Proposals

NIH SCH: An Augmented Reality Neurorehabilitation System for Monitoring and Management of Motor Symptoms of Parkinson's Disease. Project Number: 1R01HD111074 (\$1,186,393 across 4 years). Role: Student lead. I wrote significant portions of the proposal under the supervision of Professor Henry Fuchs and I helped in various aspects of the proposal submission process.

Outreach & Academic Service

Outreach

UNC-CH Computer Science Student Association Officer, Summer 2023 - Spring 2024
UNC-CH Computer Science Student Association President, Fall 2022 - Summer 2023
UNC CS Fellowship Panel Organizer, Fall 2022
Decoding Graduate Programs in CS Panel Member, Fall 2022
UNC CS Middle School/High School Open House Volunteer, Spring 2023
UNC CS Vision Seminar Organizer, Spring 2023
Summer Geometry Initiative (SGI) Mentor, Summer 2024

Academic Service

IEEE VR 2023, Reviewer
Information Systems Frontiers 2024, Reviewer
NeurIPS 2024 Datasets and Benchmarks, Reviewer
CSCW 2025, Reviewer
CHI 2025, Reviewer
CVPR 2025 - Workshop on Computer Vision for Mixed Reality, Reviewer
ISMAR 2025, Reviewer
UIST 2025, Reviewer
ML4H 2025, Reviewer and Junior Research Roundtable Chair
EuroGraphics 2026, Reviewer
Pattern Recognition 2026, Reviewer
CHI 2026, Reviewer
SIGGRAPH 2026, Reviewer