

Aayush Bansal

✉ ab.nsit@gmail.com

www.aayushbansal.xyz

☎ (773).562.9927

Education

Carnegie Mellon University, The Robotics Institute, Dec 2020
Pittsburgh, USA.

Ph.D. in Robotics (Aug 2015 - Dec 2020).

Thesis: **Unsupervised Learning of the 4D Audio-Visual World from Sparse Unconstrained Real-World Samples.**

Chair(s): Deva Ramanan and Yaser Sheikh.

Committee: Martial Hebert, David Forsyth and Alexei Efros.

Carnegie Mellon University, The Robotics Institute, Dec 2014
Pittsburgh, USA.

M.S. in Robotics (Aug 2013 - Dec 2014).

Advisor: Abhinav Gupta.

University of Delhi, Netaji Subhas Institute of Technology, June 2011
New Delhi, India.

B.E. in Electronics & Communication Engineering (Aug 2007 - June 2011).

First Class with Distinction.

Research Interests

Computer Vision & Graphics, Machine Learning, HCI, Robotics, and Psychology.

Honors and Awards

- Outstanding reviewer recognition for European Conference on Computer Vision. 2022
- Outstanding reviewer recognition for Computer Vision and Pattern Recognition. 2021
- Top-10% high-scoring reviewers for Neural Information and Processing Systems. 2020
- Open4D appeared as the front-page spotlight story on www.cmu.edu. 2020
- Best Paper Award Finalist (top-1% papers), Computer Vision and Pattern Recognition. 2019
- Outstanding reviewer recognition for Computer Vision and Pattern Recognition. 2019
- Snap Research Fellowship. Award: 10,000 USD. 2019
- Recycle-GAN appeared as the front-page spotlight story on www.cmu.edu. 2018
- International Conference on Learning Representations (ICLR) travel award. 2018
- Qualcomm Innovation Fellowship. Award: 100,000 USD. 2017
- Uber Presidential Fellowship. Award: Tuition and Stipend. 2016
- Presidential Fellow at Carnegie Mellon University. 2016
- 2nd place (out of 6 shortlisted teams) in CANINE Robotics competition hosted by U.S. Army Tank Automotive Research Development and Engineering Center (TARDEC). 2012
- Microsoft Research India travel award. 2010
- University of Delhi merit scholarship. 2007-2010

Publications

Conferences

1. L. Xu, V. Agrawal, W. Laney, T. Garcia, **A. Bansal**, C. Kim, S.R. Bulo, L. Porzi, P. Kotschieder, A. Bozic, D. Lin, M. Zollhoefer and C. Richardt. *VR-NeRF: High-Fidelity Virtualized Walkable Spaces*. ACM SIGGRAPH Asia Conference Proceedings, 2023.
2. R. Khirodkar, **A. Bansal**, L. Ma, R. Newcombe, M. Vo and K. Kitani. *EgoHumans: An Egocentric 3D Multi-Human Benchmark*. International Conference on Computer Vision (ICCV), 2023. <https://rawalkhirodkar.github.io/egohumans/>.
3. **A. Bansal** and M. Zollhoefer. *Neural Pixel Composition for 3D-4D View Synthesis from Multi-Views*. Computer Vision and Pattern Recognition (CVPR), 2023. www.aayushbansal.xyz/npc/.
4. M. Mihajlovic, **A. Bansal**, M. Zollhoefer, S. Tang, and S. Saito. *KeypointNeRF: Generalizing Image-based Volumetric Avatars using Relative Spatial Encoding of Keypoints*. European Conference on Computer Vision (ECCV), 2022. <https://markomih.github.io/KeypointNeRF/>.
5. M. Mihajlovic, S. Saito, **A. Bansal**, M. Zollhoefer, and S. Tang. *COAP: Compositional Articulated Occupancy of People*. Computer Vision and Pattern Recognition (CVPR), 2022. <https://neuralbodies.github.io/COAP/>.
6. J. Chibane, **A. Bansal**, V. Lazova, and G. Pons-Moll. *Stereo Radiance Fields: Learning View Synthesis for Sparse Views of Novel Scenes*. Computer Vision and Pattern Recognition (CVPR), 2021. <https://virtualhumans.mpi-inf.mpg.de/srf/>.
7. K. Deng, **A. Bansal**, and D. Ramanan. *Unsupervised Audiovisual Synthesis via Exemplar Autoencoders*. International Conference on Learning Representations (ICLR), 2021. www.cs.cmu.edu/~exemplar-ae/.
8. **A. Bansal**, M. Vo, Y. Sheikh, D. Ramanan, and S. Narasimhan. *4D Visualization of Dynamic Events from Unconstrained Multi-View Videos*. Computer Vision and Pattern Recognition (CVPR), 2020. www.aayushbansal.xyz/Open4D/.
9. **A. Bansal**, Y. Sheikh, and D. Ramanan. *Shapes and Context: In-the-wild Image Synthesis & Manipulation*. Computer Vision and Pattern Recognition (CVPR), 2019. **Best Paper Award Finalist**. www.aayushbansal.xyz/OpenShapes/.
10. **A. Bansal**, S. Ma, D. Ramanan, and Y. Sheikh. *Recycle-GAN: Unsupervised Video Retargeting*. European Conference on Computer Vision (ECCV), 2018. www.aayushbansal.xyz/Recycle-GAN/.
11. **A. Bansal**, Y. Sheikh, and D. Ramanan. *PixelNN: Example-based Image Synthesis*. International Conference on Learning Representations (ICLR), 2018. www.aayushbansal.xyz/pixelNN/.
12. **A. Bansal**, B. Russell, and A. Gupta. *Marr Revisited: 2D-3D Alignment via Surface Normal Prediction*. Computer Vision and Pattern Recognition (CVPR), 2016. www.aayushbansal.xyz/marrRevisited/.
13. **A. Bansal**, A. Farhadi, and D. Parikh. *Towards Transparent Systems: Semantic Characterization of Failure Modes*. European Conference on Computer Vision (ECCV), 2014.
14. **A. Bansal**, H. Badino, and D. Huber. *Understanding How Camera Configuration and Environmental Conditions Affect Appearance-based Localization*. Intelligent Vehicles (IV), 2014.
15. B. A. Stancil, J. Hyams, J. Shelly, K. Babu, H. Badino, **A. Bansal**, D. Huber, and P. Batavia. *CANINE: A Robotic Mine Dog*. In proceedings of the IS&T Conference on Electronic Imaging (SPIE), 2013.

Pre-prints, Workshops, and Technical Reports

1. K. Wang, D. Ramanan, and **A. Bansal**. *Video-Specific Autoencoders for Exploring, Editing, and Transmitting Videos*. www.aayushbansal.xyz/Video-ViSA/.
2. Z. Lin, D. Ramanan, and **A. Bansal**. *A Simple Baseline for Task-Agnostic Self-Training*.

3. V. Fragaso, C. Liu, **A. Bansal**, and D. Ramanan. *Patch-Correspondences for Interpreting Pixel-level CNNs*. Tech Report, arXiv 2018.
4. **A. Bansal**, X. Chen, B. Russell, A. Gupta, and D. Ramanan. *PixelNet: Representation of the pixels, by the pixels, and for the pixels*. Tech Report, arXiv 2017. ([> 200 citations](#)).
5. **A. Bansal**, A. Shrivastava, C. Doersch, and A. Gupta. *Mid-level Elements for Object Detection*. Tech Report, arXiv 2015.
6. **A. Bansal**. *See the reasons why I fail: Interactively Discovering Failure Modes*. Workshop on Parts & Attributes, European Conference on Computer Vision (ECCV), 2014.
7. **A. Bansal**, and K. Singh. *Storytelling Patches: Predicting Tourist Spots in a City*. Workshop on Storytelling with Images and Videos, European Conference on Computer Vision (ECCV), 2014.
8. **A. Bansal**, H. Badino, and D. Huber. *Analysis of the CMU Localization Algorithm Under Varied Conditions*. RI-CMU Technical Report, 2014.
9. **A. Bansal**, A. Kowdle, D. Parikh, A. C. Gallagher, and C. L. Zitnick. *Which Edges Matter?*. Workshop on 3D Representation and Recognition, International Conference on Computer Vision (ICCV), 2013.

Industry Experience

Reality Labs Research , Pittsburgh, USA. Meta Platforms, Inc. (formerly Facebook). Host: Yaser Sheikh, Michael Zollhoefer. Position: Postdoctoral Research Scientist.	June 2021 - June 2023
Facebook Reality Labs , Pittsburgh, USA. Host: Yaser Sheikh, Shugao Ma. Position: Research Intern.	May 2017 - Dec 2017
Adobe Research , San Francisco, USA. Host: Bryan Russell. Position: Research Intern.	May 2015 - Aug 2015

Consulting

In the past, I have provided confidential consulting services to prominent **production houses**:

- BBC Studios, London.
Naomi Austin, Series Producer/Director.
- Full frontal with Samantha Bee (TBS), New York.
Razan Ghalayini , Co-Producer/Director.
- L'Hôtel du temps (The Time Hotel).
Thierry Ardisson, French TV Producer/Host.
- NBC News, New York.
Arleen Aguasvivas, Producer Editor.
- France TV, Paris.
Louis Milano-Dupont, Journaliste/Grand Reporter.
- CBS News, Pittsburgh.
David Highfield, Co-Host/Executive Producer "Your Pittsburgh".

Teaching at Carnegie Mellon University

- 16:822 - Geometry-based Methods in Computer Vision. Fall 2017
Teaching Assistant with **Martial Hebert**.
- 16:720 - Graduate Computer Vision. Fall 2017
Guest Lecture on Image Synthesis & Generative Adversarial Networks.
Instructors: Yaser Sheikh, Simon Lucey, Srinivasa Narasimhan.
- 16:720 - Graduate Computer Vision. Spring 2015
Teaching Assistant with **Srinivasa Narasimhan**.

Service

- **Workshops**

Organizer, AI-driven content creation at Hollywood Professional Association Tech Retreat. 2021

- **Reviewer/Program Committee**

Journals

IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI). 2017-

International Journal of Computer Vision (IJCV). 2018-

ACM Transaction on Graphics (TOG, SIGGRAPH). 2019-

IEEE Transactions on Visualization and Computer Graphics (TVCG). 2019-2020

Computer Vision and Image Understanding (CVIU). 2016-2018

Machine Vision Applications (MVA). 2016-2018

Conferences

IEEE/CVF Computer Vision and Pattern Recognition (CVPR). 2018-

IEEE/CVF International Conference on Computer Vision (ICCV). 2019-

European Conference on Computer Vision (ECCV). 2018-

International Conference on Learning Representations (ICLR). 2018-

Neural Information Processing Systems (NeurIPS). 2019-

International Conference on Machine Learning (ICML). 2019-

Computer Human Interaction (CHI). 2020-

Eurographics. 2019-2020

Pacific Graphics (PG). 2019-2020

AAAI Conference on Artificial Intelligence. 2018-2020

IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). 2015-2017

IEEE International Conference on Robotics and Automation (ICRA). 2015-2017

- **Carnegie Mellon University**

Student Lead, Artificial Intelligence Seminar Series. Fall 2019 - Fall 2020

Student Lead, VASC Seminar. Fall 2019 - Fall 2020

Master's Thesis Committee Member

Gines Hidalgo Martinez, *OpenPose: Whole-Body Pose Estimation*. 2019

Donglai Xiang, *Monocular Total Capture: Posing Face, Body, and Hands in the Wild*. 2019

Ishan Nigam, *Learning from Auxiliary Supervision*. 2018

Zhe Cao, *Realtime Multi-person 2D Pose Estimation using Part Affinity Fields*. 2017

Ph.D. Research Qualifier Committee

- Martin Li, *Resource-Constrained Learning and Inference for Visual Perception.* 2020
Dinesh Reddy, *Reconstruction of dynamic vehicles from multiple unsynchronized cameras.* 2020

• Netaji Subhas Institute of Technology

- Students' Placement Coordinator. 2010-2011
President, Entrepreneurship Cell. 2010-2011
Student Head, Alumni Relationship Cell. 2010-2011
Students' Representative, Board of Students' Affairs. 2009-2010
General Secretary, Innovation (Annual Technical Symposium). 2009-2010

- **Media Interview:** NBC News, WQED, CBS Pittsburgh, 90.5 WESA, France TV, Deutschlandfunk Kultur, Journalist (Das Medienmagazine), and many other print and electronic media.

Selected News

- **System Combines Smartphone Videos to Create 4D Visualizations**, Communications of the ACM. July 2020
www.cacm.acm.org/news/246062-system-combines-smartphone-videos-to-create-4d-visualizations
- **3D capture with normal smartphone cameras**, ZDNet. July 2020
www.zdnet.com/article/3d-capture-with-normal-smartphone-cameras/
- **As COVID Slows Hollywood Pipeline, Future Tech Could Help**, CMU. May 2020
www.cmu.edu/news/stories/archives/2020/may/covid-hollywood-pipeline.html
- **Just a Minute: Recycle-GAN**, WQED, Pittsburgh. July 2019
www.youtube.com/watch?v=NuXIL0S8dKA
- **How concerned should we be that deepfakes will mess with 2020?**, NBC News. June 2019
www.youtube.com/watch?v=63QI1GByMqA
- **Deepfakes, why we should not totally freak out and how journalists should deal with it?**, Journalist. June 2019
- **Artificial imagination: CMU lab automates video “fakes”**, Tartans, CMU. Feb 2019
<https://thetartan-assets.s3.amazonaws.com/edition/pdf/423/Vol113Issue15.pdf>
- **Journalistes : mais à quoi servent-ils?**, France TV. Feb 2019
www.youtube.com/watch?v=X7-ZN6gg8wY
- **Deep Fakes: How they are made and how they can be detected**, NBC News. Nov 2018
www.youtube.com/watch?v=C8F00P2a3dA
- **This Tech Could Resurrect Charlie Chaplin & Better Autonomous Car Tests...**, 90.5 WESA. Nov 2018
- **Beyond Deep Fakes**, Communications of the ACM. Sep 2018
www.cacm.acm.org/news/231081-beyond-deep-fakes

Interns and Collaborators

Past: Researchers

- Rawal Khirodkar, Meta Ph.D. Intern with Minh Vo. 2022
Ph.D. student with Kris Kitani at CMU.
Lead author of work on EgoHumans (under submission).
- Christian Diller, Meta Ph.D. Intern with Minh Vo. 2022
Ph.D. student with Angela Dai at TU Munich.
- Sean Liu, Ph.D. student with Maneesh Agrawala at Stanford University. 2022

Vihaan Mishra, Undergraduate Intern (Junior) from NSUT Delhi. Next Position: Ph.D. student, Robotics, CMU.	2022
Anand Bhattad, Ph.D. student with David Forsyth at UIUC. <i>Co-authored DURA (under submission).</i>	2022
Jerry Liu, Undergraduate Student with Cynthia Rudin at Duke. Next Position: Ph.D. student, Stanford.	2021
Bhrij Patel, Undergraduate Student with Cynthia Rudin at Duke. Next Position: Ph.D. student, University of Maryland at College Park.	2021
Marko Mihajlovic, Meta Ph.D. Intern with Michael Zollhoefer and Shunsuke Saito. Ph.D. student at ETH Zurich with Siyu Tang. <i>Lead author of COAP (CVPR 2022) and KeypointNeRF (ECCV 2022).</i>	2021
Songyou Peng, Meta Ph.D. Intern with Michael Zollhoefer and Shunsuke Saito. Ph.D. student at ETH Zurich with Marc Pollefeys and Andreas Geiger.	2021
Zhiqiu Lin, Ph.D. student with Deva Ramanan at RI, CMU. <i>Lead author of work on Task-Agnostic Self-Training.</i>	2020
Julian Chibane, Ph.D. student with Gerard Pons-Moll at University of Tubingen. <i>Lead author of CVPR 2021 paper on Stereo Radiance Fields.</i>	2020
Kevin Wang, Undergraduate Student (Junior), CSD, CMU. <i>Lead author of work on Video-Specific Autoencoders.</i> Next Position: Software Engineer at Google Cloud.	2020
Kangle Deng, Undergraduate Summer Intern (Junior) from Peking University. <i>Lead author of ICLR 2021 paper on Exemplar Autoencoders.</i> Next Position: Ph.D. student, Robotics, CMU.	2019
Yuan Dong, Undergraduate Summer Intern (Junior) from Tsinghua University. Next Position: MS student, Computer Vision, CMU.	2017
Deepthi Hegde, Undergraduate Summer Intern (Senior) from NIT Karnataka. Next Position: MS student, Computer Science, UW Seattle.	2017
Chunhui Liu, Undergraduate Summer Intern (Junior) from Peking University. Next Position: MS student, Computer Vision, CMU.	2017
Past: Artists or UI/UX Developers	
Zixuan Li, M.S. student at HCII, CMU.	2019
Jinyi Ye, M.S. student at ETC, CMU.	2019
Xiao He, M.S. student at ETC, CMU.	2019
Akash Pushkar, M.S. student at CEE, CMU.	2019

Long Talks

- *Playing with 2D Pixels for Detailed 3D Modeling*
Workshop on 3D Generative Modeling, Banff (Canada). July 2023
CMU RISS RoboLaunch. July 2023
- *Building Generative A.I. Systems with Minimal Resources*
National University of Singapore (NUS). March 2023
University of California, Merced. Jan 2023
- *Building Maximal Vision Systems with Minimal Resources*
Indraprastha Institute of Information Technology (IIIT), Delhi. Jan 2023
Netaji Subhas University of Technology, Delhi. Jan 2023

- Indian Institute of Science (IISc), Bangalore. Dec 2022
- International Institute of Information Technology (IIIT), Hyderabad. Dec 2022
- *Computer: The Ultimate Communication Device*
Hollywood Professional Association Tech Retreat. March 2021
- *Unsupervised Exemplar Representations: Beyond Task-based Optimization*
Vision Group, University of Bristol. May 2021
Vision Seminar, Massachusetts Institute of Technology. April 2021
brAIn Group, Carnegie Mellon University. April 2021
- *Unsupervised Learning of the 4D Audio-Visual World from Sparse Real-World Samples*
Vision Group, University of Illinois, Urbana Champaign. Feb 2021
Facebook Reality Labs, Pittsburgh. Jan 2021
Google VisCam Group, Cambridge. Jan 2021
Carnegie Mellon University, Pittsburgh. Dec 2020
- *Computational Studio: Towards audio-visual social communication*
Vision & Learning Group, Stanford University. Aug 2020
Vision Group, University of California, Berkeley. Aug 2020
- *Computational Studio: A computational machinery to enhance social communication*
Inria, Paris. May 2020
Weizmann Institute of Sciences, Israel. May 2020
Pixel Club, Technion, Israel Institute of Technology. May 2020
Artificial Intelligence Seminar Series, CMU. May 2020
Max Planck Institute for Intelligent Systems, Tübingen. May 2020
Graphics Lunch, Stanford University. May 2020
Max Planck Institute for Informatics, Saarbrücken. May 2020
- *Data-driven Computational Studio*
CS Colloquium, University of North Carolina, Chapel Hill. March 2020
CS Colloquium, Yale University, New Haven. Feb 2020
- *Data-driven & User-Controllable Audio-Visual Content Creation*
Google Daydream, Seattle. Feb 2020
Adobe Research, San Jose. Jan 2020
Vision Seminar, Massachusetts Institute of Technology. Jan 2020
Guest Lecture, Creative AI (10-737), Machine Learning Department, CMU. Oct 2019
University of California, Irvine. Oct 2019
Snap Research, Los Angeles. Oct 2019
- *Association & Imagination*
ECE Department, Virginia Tech, Blacksburg. Oct 2019
PIXL Lunch, Princeton University. Sep 2019
Artificial Intelligence Seminar Series, CMU. Aug 2019
- *Data-driven & User-Controllable Visual Content Creation*

- Guest Lecture, Advanced Media Creation Studio (54-476). School of Drama, CMU. Aug 2019
- Aurora Inc, Pittsburgh. May 2019
- Loom AI, San Francisco. May 2019
- Guest Lecture, Ethics & Policies in Computing at CMU. Feb 2019
- *My experiments with simple nearest neighbors*
- Vision Group, University of California, Berkeley. May 2019
- Machine Learning Department (Lunch Seminar), CMU. May 2019
- *Virtual Time-Travel: 4D space-time visualization of dynamic events*
- Qualcomm Inc, San Diego. March 2019
- *Two Tales about Image Synthesis*
- University of British Columbia, Vancouver. May 2018
- Simon Fraser University, Vancouver. April 2018
- *Pixels: Synthesis, Inference, and Interpretation*
- Indian Institute of Science (IISc), Bangalore. Feb 2018
- Google DeepMind, London. Nov 2017
- *PixelNet*
- Intel ISTC, Pittsburgh. April 2017
- Qualcomm Inc, San Diego. April 2017
- National Robotics Engineering Consortium (NREC), Pittsburgh. March 2017
- Computer Vision & Robotics Group, Stanford University. Jan 2017
- Loom AI, San Francisco. Jan 2017
- Machine Learning Department (Lunch Seminar), CMU. Feb 2017
- *2D-3D Model Alignment via Surface Normal Prediction*
- Guest Lecture, Undergraduate computer vision at CMU. Sep 2016
- Fyusion Inc, San Francisco. Aug 2016

Short Talks

- *Data-driven & User-Controllable Audio-Visual Content Creation*
- Discovery Analytics Center, Virginia Tech. Oct 2019
- *Data-driven Visual Content Creation*
- CyLab/Software Engineering Institute, CMU. Sep 2019
- *Shapes & Context: In-the-wild Image Synthesis & Manipulation*
- Oral talk at CVPR 2019. June 2019
- *Marr Revisited: 2D-3D Model Alignment via Surface Normal Prediction*
- Mid-Atlantic Vision (MACV) Workshop, 2016 at Johns Hopkins. March 2017