

Gengshan Yang

✉ y.gengshan@gmail.com
🌐 gengshan-y.github.io
🔗 [gengshan-y](https://github.com/gengshan-y)
🔗 [yRaFnrA4AAAJ](https://www.linkedin.com/in/yRaFnrA4AAAJ)

Education

- Jul 2023 **Ph.D. in Robotics**, *Carnegie Mellon University*, Advisor: Deva Ramanan
Thesis: Building 4D Models of Objects and Scenes from Monocular Videos
- Aug 2019 **M.S. in Robotics**, *Carnegie Mellon University*, Advisor: Deva Ramanan
Thesis: Volumetric Correspondence Networks for Stereo Matching and Optical Flow
- Jun 2017 **B.E. in Automation**, *Xi'an Jiaotong University*

Employment

- Sep 2024– **World Labs**, *San Francisco*
Member of Technical Staff
- 2023–2024 **Codec Avatar Labs at Meta**, *Pittsburgh*
Research Scientist
- 2021 **Facebook AI Research**, *Menlo Park*
Research Intern, Hosted by Hanbyul Joo
- 2020–2021 **Google Research**, *Cambridge*
Research Intern, Hosted by Deqing Sun and Varun Jampani
- 2018 **Argo AI**, *Pittsburgh*
Research Intern

Publications

- [1] Agent-to-Sim: Learning Interactive Behavior from Casual Videos.
Gengshan Yang, Andrea Bajcsy, Angjoo Kanazawa*, Shunsuke Saito*.
International Conference on Learning Representations (ICLR), 2025.
- [2] DressRecon: Freeform 4D Human Reconstruction from Monocular Video.
Jeff Tan, Donglai Xiang, Shubham Tulsiani, Deva Ramanan, **Gengshan Yang**.
International Conference on 3D Vision (3DV), 2025. **Oral Presentation.**
- [3] Tactile DreamFusion: Exploiting Tactile Sensing for 3D Generation.
Ruihan Gao, Kangle Deng, **Gengshan Yang**, Wenzhen Yuan, Jun-Yan Zhu.
Conference on Neural Information Processing Systems (NeurIPS), 2024.
- [4] SplatAM: Splat, Track & Map 3D Gaussians for Dense RGB-D SLAM.
Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, **Gengshan Yang**, Sebastian Scherer, Deva Ramanan, Jonathon Luiten.

IEEE Computer Vision and Pattern Recognition (CVPR), 2024.

- [5] SLoMo: A General System for Legged Robot Motion Imitation from Casual Videos.
John Z. Zhang, Shuo Yang, **Gengshan Yang**, Arun Bishop, Swaminathan Gurumurthy, Deva Ramanan, Zachary Manchester.
IEEE Robotics and Automation Letters (RA-L), 2023.
- [6] Physically Plausible Reconstruction from Monocular Videos.
Gengshan Yang, Shuo Yang, John Zhang, Zachary Manchester, Deva Ramanan.
International Conference on Computer Vision (ICCV), 2023. Oral Presentation.
- [7] Total-Recon: Deformable Scene Reconstruction for Embodied View Synthesis.
Chonghyuk Song, **Gengshan Yang**, Kangle Deng, Jun-Yan Zhu, Deva Ramanan.
International Conference on Computer Vision (ICCV), 2023.
- [8] Reconstructing Animatable Categories from Videos.
Gengshan Yang, Chaoyang Wang, N Dinesh Reddy, Deva Ramanan.
IEEE Computer Vision and Pattern Recognition (CVPR), 2023.
- [9] Distilling Neural Fields for Real-time Articulated Shape Reconstruction.
Jeff Tan, **Gengshan Yang**, Deva Ramanan.
IEEE Computer Vision and Pattern Recognition (CVPR), 2023.
- [10] 3D-aware Conditional Image Synthesis.
Kangle Deng, **Gengshan Yang**, Deva Ramanan, Jun-Yan Zhu.
IEEE Computer Vision and Pattern Recognition (CVPR), 2023.
- [11] BANMo: Building Animatable 3D Neural Models from Many Casual Videos.
Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, Hanbyul Joo.
IEEE Computer Vision and Pattern Recognition (CVPR), 2022. Oral Presentation.
- [12] ViSER: Video Surface Embeddings for Articulated 3D Shape Reconstruction.
Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Ce Liu, Deva Ramanan.
Conference on Neural Information Processing Systems (NeurIPS), 2021. Spotlight.
- [13] NeRS: Neural Reflectance Surfaces for Sparse-view 3D Reconstruction in the Wild.
Jason Y. Zhang, **Gengshan Yang**, Shubham Tulsiani*, Deva Ramanan*.
Conference on Neural Information Processing Systems (NeurIPS), 2021.
- [14] LASR: Learning Articulated Shape Reconstruction from a Monocular Video.
Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Huiwen Chang, Deva Ramanan, William T. Freeman, Ce Liu.
IEEE Computer Vision and Pattern Recognition (CVPR), 2021.
- [15] Learning to Segment Rigid Motions from Two Frames.
Gengshan Yang, Deva Ramanan.
IEEE Computer Vision and Pattern Recognition (CVPR), 2021.
- [16] Upgrading Optical Flow to 3D Scene Flow through Optical Expansion.
Gengshan Yang, Deva Ramanan.
IEEE Computer Vision and Pattern Recognition (CVPR), 2021. Oral Presentation.

- [17] Volumetric Correspondence Networks for Optical Flow.
Gengshan Yang, Deva Ramanan.
Conference on Neural Information Processing Systems (NeurIPS), 2021.
- [18] Hierarchical Deep Stereo Matching on High-resolution Images.
Gengshan Yang, Joshua Manela, Michael Happold, Deva Ramanan.
IEEE Computer Vision and Pattern Recognition (CVPR), 2021.
- [19] Inferring Distributions Over Depth from a Single Image.
Gengshan Yang, Peiyun Hu, Deva Ramanan.
International Conference on Intelligent Robots and Systems (IROS), 2019.
- [20] Towards Hand-dominated Activity Recognition Systems with Wristband-interaction Behavior Analysis.
 Chao Shen, Yufei Chen, **Gengshan Yang**, and Xiaohong Guan.
IEEE Transactions on Systems, Man, and Cybernetics: Systems, 2018.

█ Awards and Honors

- 2023 AI Rising Stars, Michigan AI Symposium 2023
- 2023 Doctoral Consortium, CVPR 2023
- 2021 Qualcomm Innovation Fellowship
- 2014 – 2016 National Scholarship of China

█ Teaching

- Spring 2021 Graduate Computer Vision *Teaching Assistant, CMU*
- Fall 2020 Geometry-based Methods for Computer Vision *Teaching Assistant, CMU*

█ Service

- 2025 Area Chair *ICCV*
- 2024 Guest Editor *IJCV Special Issue on Computer Vision for Animal Tracking and Modeling*
- 2019– Conference Reviewer *CVPR, ECCV, ICCV, NeurIPS, ICLR, IROS, ICRA*
- 2019– Journal Reviewer *T-PAMI, IJCV, TMLR, RA-L, CVIU*
- 2020–2021 Undergrad AI Mentoring Program *School of Computer Science, CMU*

█ Workshop

- 2025 **Workshop on 4D Vision: Modeling the Dynamic World, CVPR 2025**
- 2025 **5th Workshop on CV4Animals, CVPR 2025**
- 2024 **4th Workshop on CV4Animals, CVPR 2024**

Talks

- 2024 **Stability AI**, *Hosted by Varun Jampani*
Towards Generative 3D World Models from Videos
- 2024 **UIUC**, *External Speaker Series, Hosted by Yu-Xiong Wang*
Towards 4D Reconstruction in the Wild
- 2024 **University of Maryland**, *CS Seminar*
Towards 4D Reconstruction in the Wild
- 2024 **UMich**, *CSE Seminar*
Towards 4D Reconstruction in the Wild
- 2024 **TTIC**, *Talks at TTIC*
Towards 4D Reconstruction in the Wild
- 2023 **UMich AI Symposium**, *Lightning Talk*
Towards 4D Reconstruction in the Wild
- 2023 **AIT Lab, ETH Zürich**, *Hosted by Jie Song*
Building 4D Models of Objects and Scenes from Videos
- 2023 **Stanford**, *Hosted by Jiajun Wu*
Capture A Dynamic 3D World from Casual Videos
- 2023 **BAIR, UC Berkeley**, *Hosted by Angjoo Kanazawa*
Capture A Dynamic 3D World from Casual Videos
- 2023 **NVIDIA**, *Hosted by Chen-Hsuan Lin*
Capture A Dynamic 3D World from Casual Videos
- 2023 **Epics Games**, *Hosted by Christoph Lassner*
Capture A Dynamic 3D World from Casual Videos
- 2023 **Boston Dynamics AI Institute**, *Hosted by Jiuguang Wang*
Capture A Dynamic 3D World from Casual Videos
- 2023 **Reality Labs at Meta**, *External Speaker Series, Hosted by Giljoo Nam*
Capture A Dynamic 3D World from Casual Videos
- 2022 **RPAD, Carnegie Mellon University**, *Hosted by Wenxuan Zhou, David Held*
Building Animatable 3D Neural Models from Many Casual Videos
- 2022 **Cornell Tech**, *Hosted by Qianqian Wang, Noah Snavley*
Building Animatable 3D Neural Models from Many Casual Videos
- 2022 **KAIR, UC Berkeley**, *Hosted by Angjoo Kanazawa*
Building Animatable 3D Neural Models from Many Casual Videos
- 2022 **VGG, Oxford**, *Hosted by Shangzhe (Elliott) Wu*
Building Animatable 3D Neural Models from Many Casual Videos

- 2022 **Baidu**, *Hosted by Errui Ding*
Building Animatable 3D Neural Models from Many Casual Videos
- 2021 **Carnegie Mellon University**, *Hosted by Adam Harley, Katerina Fragkiadaki*
Video-Specific Surface Embeddings for Articulated 3D Shape Reconstruction
- 2020 **Sun Yat-Sen University**, *Hosted by Yulan Guo*
Upgrading Optical Flow to 3D for Monocular Dynamic Scene Perception