Krishna Kumar Singh

Senior Research Scientist 2 and Manager
Adobe Research

7248718714
kkanshul@gmail.com
kkanshul.github.io

Summary

As a Senior Research Scientist 2 and Manager at Adobe Research, my work sits at the crossroads of generative models, visual understanding, and creative expression. I've spearheaded core technologies like image inpainting/outpainting and editing, now powering Generative Fill and Expand in Photoshop, Lightroom, and Firefly — tools used by millions, and recognized by TIME as one of the Best Inventions of 2023. Today, my focus is on designing unified multimodal models that can both generate and comprehend — bridging the gap between input and intent. My research has been published in 50+ top-tier papers (CVPR, ICCV, ECCV, NeurIPS, SIGGRAPH, etc.), and I filed 50+ patents across generative AI, computer vision, and multimodal systems.

Education

2015 – 2020 **PhD in Computer Science**, *University of California, Davis*
GPA: 3.93
Advisor: Prof. Yong Jae Lee

2013 – 2014 **Masters in Robotics**, *Carnegie Mellon University*
QPA: 4.05
Advisors: Prof. Alexei Efros, Prof. Kayvon Fatahalian

2009 – 2013 **B.Tech (Honours) in Computer Science and Engineering**, *IIIT Hyderabad*
GPA: 9.07/10
Advisor: Prof. P. J. Narayanan

Publications

2026 Streaming Multi-Agent Autoregressive Diffusion Model with World State Registers. Sicheng Mo, Yuheng Li, Ziyang Leng, Krishna Kumar Singh, Bolei Zhou. *arXiv 2026*.

2026 Tev: Scaling Long-Horizon Agentic Environments with Dense Progress Rewards. Wenbo Hu, Shraman Pramanick, Krishna Kumar Singh, Nanyun Peng, Kai-Wei Chang, Yong Jae Lee. *arXiv 2026*.

2026 Personal AI Agent for Camera Roll VQA. Thao Nguyen, Krishna Kumar Singh, Donghyun Kim, Yong Jae Lee, Yuheng Li. *arXiv 2026*.

2026 Pristine Models. Anirudh Sundara Rajan, Krishna Kumar Singh, Richard Zhang, Yong Jae Lee. *arXiv 2026*.

2026 Segment Where You Look: Attention as the Interface Between Language Models and Segmenters. Eslam Abdelrahman, Hamza Omar Alkaf, Yuheng Li, Siddharth Srinivasan Iyer, Nanxuan Zhao, Eli Shechtman, Krishna Kumar Singh, Mohamed Elhoseiny. *arXiv 2026*.

2026 AdvantageFlow: Advantage-Weighted Least Squares for RL in Flow Models. Branislav Kveton, Anup Rao, Subhojyoti Mukherjee, Krishna Kumar Singh, Viet Dac Lai. *arXiv 2026*.

2026 From Plans to Pixels: Learning to Plan and Orchestrate for Open-Ended Image Editing. Anirudh Sundara Rajan, Krishna Kumar Singh, Yong Jae Lee. *arXiv 2026*.

2026 UniTemp: Unlocking Video Generation in Any Temporal Order via Bidirectional Distillation. Lin Zhang, Sicheng Mo, Zefan Cai, Jinhong Lin, Zihao Lin, Jiuxiang Gu, Krishna Kumar Singh, Yuheng Li, Yin Li. *ECCV 2026*.

2026 MAOAM: Unified Object & Material Selection with Vision-Language Models. Jaden Park, Valentin Deschaintre, Jason Kuen, Kangning Liu, Iliyan Georgiev, Krishna Kumar Singh, Yong Jae Lee, Michael Fischer. *SIGGRAPH 2026*.

2026 MultiCOIN: Multi-Modal Controllable INbetweening. Maham Tanveer, Yang Zhou, Simon Niklaus, Ali Mahdavi Amiri, Hao Zhang, Krishna Kumar Singh, Nanxuan Zhao. *Eurographics 2026*.

2026 Controllable Cinemagraph Generation from A Still Image. Van Thanh Le, Daichi Ito, Aniruddha Mahapatra, Long Mai, Krishna Kumar Singh, Kuldeep Kulkarni, Feng Liu, Yun Fu, Jae Shin Yoon. *Eurographics 2026 (Short Paper)*.

2026 Stepwise Credit Assignment for GRPO on Flow-Matching Models. Yash Savani, Branislav Kveton, Yuchen Liu, Yilin Wang, Jing Shi, Subhojyoti Mukherjee, Nikos Vlassis, Krishna Kumar Singh. *CVPR 2026*.

2026 Relational Visual Similarity. Thao Nguyen, Sicheng Mo, Krishna Kumar Singh, Yilin Wang, Jing Shi, Nicholas Kolkin, Eli Shechtman, Yong Jae Lee, Yuheng Li. *CVPR 2026*.

2026 Group Diffusion: Enhancing Image Generation by Unlocking Cross-Sample Collaboration. Sicheng Mo, Thao Nguyen, Richard Zhang, Nick Kolkin, Siddharth Srinivasan Iyer, Eli Shechtman, Krishna Kumar Singh, Yong Jae Lee, Bolei Zhou, Yuheng Li. *CVPR 2026*.

2026 NP-Edit: Learning an Image Editing Model without Image Editing Pairs. Nupur Kumari, Sheng-Yu Wang, Nanxuan Zhao, Yotam Nitzan, Yuheng Li, Krishna Kumar Singh, Richard Zhang, Eli Shechtman, Jun-Yan Zhu, Xun Huang. *ICLR 2026*.

2026 Beyond Simple Edits: X-Planner for Complex Instruction-Based Image Editing. Chun-Hsiao Yeh, Yilin Wang, Nanxuan Zhao, Richard Zhang, Yuheng Li, Yi Ma, Krishna Kumar Singh. *AAAI 2026*.

2025 X-Fusion: Introducing New Modality to Frozen Large Language Models. Sicheng Mo, Thao Nguyen, Xun Huang, Siddharth S. Iyer, Yijun Li, Yuchen Liu, Abhishek Tandon, Eli Shechtman, Krishna Kumar Singh, Yong Jae Lee, Bolei Zhou, Yuheng Li. *ICCV 2025*.

2025 Generating, Fast and Slow: Scalable Parallel Video Generation with Video Interface Networks. Bhishma Dedhia, David Bourgin, Krishna Kumar Singh, Yuheng Li, Yan Kang, Zhan Xu, Niraj K. Jha, Yuchen Liu. *ICCV 2025*.

2025 DOLLAR: Few-Step Video Generation via Distillation and Latent Reward Optimization. Zihan Ding, Chi Jin, Difan Liu, Haitian Zheng, Krishna Kumar Singh, Qiang Zhang, Yan Kang, Zhe Lin, Yuchen Liu. *ICCV 2025*.

2025 Can3Tok: Canonical 3D Tokenization and Latent Modeling of Scene-Level 3D Gaussians. Quankai Gao, Iliyan Georgiev, Tuanfeng Yang Wang, Krishna Kumar Singh, Ulrich Neumann, Jae Shin Yoon. *ICCV 2025*.

2025 3D-Fixup: Advancing Photo Editing with 3D Priors. Yen-Chi Cheng, Krishna Kumar Singh, Jae Shin Yoon, Alex Schwing, Liangyan Gui, Matheus Gadelha, Paul Guerrero, Nanxuan Zhao. *SIGGRAPH 2025*.

2025 Yo'Chameleon: Personalized Vision and Language Generation. Thao Nguyen, Krishna Kumar Singh, Jing Shi, Trung Bui, Yong Jae Lee, Yuheng Li. *CVPR 2025*.

2025 ShotAdapter: Text-to-Multi-Shot Video Generation with Diffusion Models. Ozgur Kara, Krishna Kumar Singh, Feng Liu, Duygu Ceylan, James M. Rehg, and Tobias Hinz. *CVPR 2025*.

2025 Comprehensive Relighting: Generalizable and Consistent Monocular Human Relighting and Harmonization. Junying Wang, Jingyuan Liu, Xin Sun, Krishna Kumar Singh, Zhixin Shu, He Zhang, Jimei Yang, Nanxuan Zhao, Tuanfeng Y. Wang, Simon S. Chen, Ulrich Neumann, Jae Shin Yoon. *CVPR 2025*.

2024 ActAnywhere: Subject-Aware Video Background Generation. Boxiao Pan, Zhan Xu, Chun-Hao Paul Huang, Krishna Kumar Singh, Yang Zhou, Leonidas J. Guibas, Jimei Yang. *NeurIPS 2024*.

2024 Generative Portrait Shadow Removal. Jae Shin Yoon, Zhixin Shu, Mengwei Ren, Xuaner Zhang, Yannick Hold-Geoffroy, Krishna Kumar Singh, He Zhang. *SIGGRAPH Asia 2024*.

2024 Removing Distributional Discrepancies in Captions Improves Image-Text Alignment. Yuheng Li, Haotian Liu, Mu Cai, Yijun Li, Eli Shechtman, Zhe Lin, Yong Jae Lee, and Krishna Kumar Singh. *ECCV 2024*.

2024 GroupDiff: Diffusion-based Group Portrait Editing. Yuming Jiang, Nanxuan Zhao, Qing Liu, Krishna Kumar Singh, Shuai Yang, Chen Change Loy, Ziwei Liu. *ECCV 2024*.

2024 Enhanced Controllability of Diffusion Models via Feature Disentanglement and Realism-Enhanced Sampling Methods. Wonwoong Cho, Hareesh Ravi, Midhun Harikumar, Vinh Khuc, Krishna Kumar Singh, Jingwan Lu, David I Inouye, Ajinkya Kale. *ECCV 2024*.

2024 Separate-and-Enhance: Compositional Finetuning for Text-to-Image Diffusion Models. Zhipeng Bao, Yijun Li, Krishna Kumar Singh, Yu-Xiong Wang, Martial Hebert. *SIGGRAPH 2024*.

2024 UniHuman: A Unified Model for Editing Human Images in the Wild. Nannan Li, Qing Liu, Krishna Kumar Singh, Yilin Wang, Jianming Zhang, Bryan A. Plummer, Zhe Lin. *CVPR 2024*.

2024 P2D: Plug and Play Discriminator for accelerating GAN frameworks. Min Jin Chong, Krishna Kumar Singh, Yijun Li, Jingwan Lu, David Forsyth. *WACV 2024*.

2024 Discovering and Mitigating Biases in CLIP-based Image Editing. Md Mehrab Tanjim, Krishna Kumar Singh, Kushal Kafle, Ritwik Sinha, Garrison W. Cottrell. *WACV 2024*.

2024 Consistent Multimodal Generation via A Unified GAN Framework. Zhen Zhu, Yijun Li, Weijie Lyu, Krishna Kumar Singh , Zhixin Shu, Sören Pirk, Derek Hoiem. *WACV 2024*.

2023 UMFuse: Unified Multi View Fusion for Human Editing applications. Rishabh Jain, Mayur Hemani, Duygu Ceylan, Krishna Kumar Singh, Jingwan Lu, Mausoom Sarkar, Balaji Krishnamurthy. *ICCV 2023*.

2023 Zero-shot Image-to-Image Translation. Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, Jun-Yan Zhu. *SIGGRAPH 2023*.

2023 Modulating Pretrained Diffusion Models for Multimodal Image Synthesis. Cusuh Ham, James Hays, Jingwan Lu, Krishna Kumar Singh, Zhifei Zhang, Tobias Hinz. *SIGGRAPH 2023*.

2023 Putting People in Their Place: Affordance-Aware Human Insertion into Scenes. Sumith Kulal, Tim Brooks, Alex Aiken, Jiajun Wu, Jimei Yang, Jingwan Lu, Alexei A Efros, Krishna Kumar Singh. *CVPR 2023*.

2023 VGFlow: Visibility guided Flow Network for Human Reposing. Rishabh Jain, Krishna Kumar Singh, Mayur Hemani, Jingwan Lu, Mausoom Sarkar, Duygu Ceylan, Balaji Krishnamurthy. *CVPR 2023*.

2023 Complete 3D Human Reconstruction from a Single Incomplete Image. Junying Wang, Jae Shin Yoon, Tuanfeng Y. Wang, Krishna Kumar Singh, Ulrich Neumann. *CVPR 2023*.

2022 Debiasing Image-to-Image Translation Models. Md Mehrab Tanjim, Krishna Kumar Singh, Kushal Kafle, Ritwik Sinha, Garrison W. Cottrell. *ECCV 2022*.

2022 Contrastive Learning for Diverse Disentangled Foreground Generation. Yuheng Li, Yijun Li, Jingwan Lu, Eli Shechtman, Yong Jae Lee, Krishna Kumar Singh. *ECCV 2022*.

2022 Discovering and Mitigating Biases in CLIP-based Text-to-Image Generation. Md Mehrab Tanjim, Krishna Kumar Singh, Kushal Kafle, Ritwik Sinha, Garrison W. Cottrell. *ECCV Workshop 2022*.

2022 Spatially-Adaptive Multilayer Selection for GAN Inversion and Editing. Gaurav Parmar, Yijun Li, Jingwan Lu, Richard Zhang, Jun-Yan Zhu, Krishna Kumar Singh. *CVPR 2022*.

2022 GIRAFFE HD: A High-Resolution 3D-aware Generative Model. Yang Xue, Yuheng Li, Krishna Kumar Singh, Yong Jae Lee. *CVPR 2022*.

2022 InsetGAN for Full-Body Image Generation. Anna Frühstück, Krishna Kumar Singh, Eli Shechtman, Niloy J. Mitra, Peter Wonka, Jingwan Lu. *CVPR 2022*.

2022 Generating and Controlling Diversity in Image Search. Md Mehrab Tanjim, Ritwik Sinha, Krishna Kumar Singh, Sridhar Mahadevan, David Arbour, Moumita Sinha, Garrison W. Cottrell. *WACV 2022*.

2021 Dance In the Wild: Monocular Human Animation with Neural Dynamic Appearance Synthesis. Tuanfeng Y. Wang, Duygu Ceylan, Krishna Kumar Singh, Niloy J. Mitra. *3DV 2021 (oral)*.

2021 PartGAN: Weakly-supervised Part Decomposition for Image Generation and Segmentation. Yuheng Li, Krishna Kumar Singh, Yang Xue, Yong Jae Lee. *BMVC 2021*.

2021 Collaging Class-specific GANs for Semantic Image Synthesis. Yuheng Li, Yijun Li, Jingwan Lu, Eli Shechtman, Yong Jae Lee, Krishna Kumar Singh. *ICCV 2021*.

2021 Seeing the Unseen: Predicting the First-Person Camera Wearer's Location and Pose in Third-Person Scenes. Yangming Wen, Krishna Kumar Singh, Markham Anderson, Wei-Pang Jan, Yong Jae Lee. *ICCV Workshop 2021*.

2021 IMAGINE: Image Synthesis by Image-Guided Model Inversion. Pei Wang, Yijun Li, Krishna Kumar Singh, Jingwan Lu, Nuno Vasconcelos. *CVPR 2021*.

2021 Generating Furry Cars: Disentangling Object Shape and Appearance across Multiple Domains. Utkarsh Ojha, Krishna Kumar Singh, Yong Jae Lee. *ICLR 2021*.

2020 Elastic-InfoGAN: Unsupervised Disentangled Representation Learning in Class-Imbalanced Data. Utkarsh Ojha, Krishna Kumar Singh, Cho-Jui Hsieh, Yong Jae Lee. *NeurIPS 2020*.

2020 Don't Judge an Object by Its Context: Learning to Overcome Contextual Bias. Krishna Kumar Singh, Dhruv Mahajan, Kristen Grauman, Yong Jae Lee, Matt Feiszli, Deepti Ghadiyaram. *CVPR 2020 (oral)*.

2020 MixNMatch: Multifactor Disentanglement and Encoding for Conditional Image Generation. Yuheng Li, Krishna Kumar Singh, Utkarsh Ojha, Yong Jae Lee. *CVPR 2020*.

2019 FineGAN: Unsupervised Hierarchical Disentanglement for Fine-Grained Object Generation and Discovery. Krishna Kumar Singh, Utkarsh Ojha, Yong Jae Lee. *CVPR 2019 (oral)*.

2019 You Reap What You Sow: Using Videos to Generate High Precision Object Proposals for Weakly-supervised Object Detection. Krishna Kumar Singh, Yong Jae Lee. *CVPR 2019*.

2018 DOCK: Detecting Objects by Transferring Common-sense Knowledge. Krishna Kumar Singh, Santosh Divvala, Ali Farhadi, Yong Jae Lee. *ECCV 2018*.

2018 Who Will Share My Image? Predicting the Content Diffusion Path in Online Social Networks. Wenjian Hu, Krishna Kumar Singh, Fanyi Xiao, Jinyoung Han, Chen-Nee Chuah, Yong Jae Lee. *WSDM 2018*.

2018 Hide-and-Seek: A Data Augmentation Technique for Weakly-Supervised Localization and Beyond. Krishna Kumar Singh, Hao Yu, Aron Sarmasi, Gautam Pradeep, Yong Jae Lee. *Arxiv 2018*.

2017 Hide-and-Seek: Forcing a Network to be Meticulous for Weakly-supervised Object and Action Localization. Krishna Kumar Singh, Yong Jae Lee. *ICCV 2017*.

2017 Identifying First-Person Camera Wearers in Third-Person Videos. Chenyou Fan, Jangwon Lee, Mingze Xu, Krishna Kumar Singh, Yong Jae Lee, David J. Crandall, Michael S. Ryoo. *CVPR 2017*.

2016 End-to-End Localization and Ranking for Relative Attributes. Krishna Kumar Singh, Yong Jae Lee. *ECCV 2016*.

2016 Track and Transfer: Watching Videos to Simulate Strong Human Supervision for Weakly-Supervised Object Detection. Krishna Kumar Singh, Fanyi Xiao, Yong Jae Lee. *CVPR 2016*.

2016 KrishnaCam: Using a Longitudinal, Single-Person, Egocentric Dataset for Scene Understanding Tasks. Krishna Kumar Singh, Kayvon Fatahalian, Alexei A. Efros. *WACV 2016*.

2014 Storytelling Patches: Predicting Tourist Spots in a City. Aayush Bansal, Krishna Kumar Singh. *ECCV 2014 Workshop*.

2012 Geometry Directed Browser for Personal Photographs. Aditya Deshpande, Siddharth Choudhary, P.J. Narayanan, Krishna Kumar Singh, Kaustav Kundu, Aditya Singh, Apurva Kumar. *ICVGIP 2012 (oral)*.

2012 Hybrid Multi-Core Algorithms for Regular Image Filtering Applications. Shrenik Lad, Krishna Kumar Singh, Kishore Kothapalli, P.J. Narayanan. *HiPC 2012 – SRS*.

Industry/Tech Transfers

Instruct Edit with Masks in Adobe Photoshop.

Prompt based Image Editing in Adobe Firefly and Photoshop.

Generative Remove in Adobe Lightroom.

Next version of Generative Fill and Expand in Adobe Photoshop and Firefly.

Generative Fill and Expand in Adobe Photoshop and Firefly (core researcher for modeling).

Skin Smoothing Neural Filter in Adobe Photoshop

Project Strike a Pose at Adobe Max Sneak

Experience

Present **Senior Research Scientist 2 and Manager**, *Adobe Research*

Summer 2019 **Research Intern**, *Facebook AI*
Worked on object and attribute classification bias mitigation.

Summer 2017 **Research Intern**, *Allen Institute for AI*
Worked on improving object detection via common-sense knowledge transfer.

Summer 2015 **Computer Vision Intern**, *Intel Labs*
Worked on video summarization and retrieval.

Spring 2020 **Co-instructor**, *Computer Vision, UC Davis*

2015 – 2020 **Research and Teaching Assistant**, *UC Davis*

2014 – 2015 **Graduate Research Assistant**, *RI, CMU*

Summer 2012 **Research Scholar**, *RI, CMU*

Summer 2011 **Research Assistant**, *IIIT Hyderabad*
Developed DLD virtual lab.

2011-2013 **Teaching Assistant**, *IIIT Hyderabad*
Courses: Information Retrieval, Data Structures, Cloud Computing.

Awards

Best Paper at CVPR 2026 CogVL Workshop

Best Paper at CVPR 2025 T4V Workshop

Best Poster at CVPR 2023 AI4CC Workshop

TIME Best Invention of 2023: Generative Fill and Expand

Adobe Max Sneaks Presenter

ICCV 2019 Doctoral Consortium

UC Davis Best Graduate Researcher Award (Hon. Mention), 2019

Microsoft Azure Research Award, 2017

AWS Research Grant, 2016

IIIT-H Research Award, 2012

Best Poster (NVIDIA), HiPC 2012

All India Rank 1902 in AIEEE 2009 (99.8 percentile)

Relevant Coursework

Graduate Computer Vision, ML, Big Data in CV, Visual Recognition, Visual Computing

Undergrad Computer Vision, Image Processing, Machine Learning, AI, Graphics, Info Extraction, Statistical AI

Skills

Languages Python, C, C++, Java (basic), Lua

Frameworks PyTorch, Pytorch Lightning, Caffe, Torch, Matlab, OpenGL, OpenCV

Tools FFmpeg, Vim, Git, Eclipse, Hadoop, EC2, Web2py, CUDA (basic), OpenMP, HTML

Service

Lead Area Chair for CVPR 2026

Co-orgnaizer for Workshop on Closing the Loop Between Vision and Language (CLVL), ICCV 2025

Area Chair for ICCV 2025

Co-orgnaizer for AI for Content Creation Workshop (AI4CC), CVPR 2025 and 2026

Area Chair for CVPR 2025

Area Chair for ECCV 2024

SPC for AAAI 2023

Reviewer: CVPR, ICCV, ECCV, NeurIPS, ICML, SIGGRAPH, TPAMI, IJCV, AAAI, IJCAI, WACV, BMVC

Program Committee: EPIC 2016, Assistive Vision ACCV 2016, Intention Understanding ACCV 2018

Mentor: Aggiementor UC Davis – student won national prize for fruit detection project

Talks/Panels

Panelist for AI for Creative Visual Content Generation Editing and Understanding (CVEU), CVPR 2025

Guest lecture at KAUST, 2024

Guest lecture at UW-Madison, 2022

Invited talk at Netflix, 2022

Panelist for Deep Learning 2.0 Virtual Summit, 2021